

PolyGym: Polyhedral Optimizations as an Environment for Reinforcement Learning

Alexander Brauckmann*

Chair for Compiler Construction

TU Dresden, Germany

alexander.brauckmann@tu-dresden.de

Andrés Goens*†

Barkhausen Institut

Dresden, Germany

andres.goens@barkhauseninstitut.org

Jerónimo Castrillon

Chair for Compiler Construction

Center for Advancing Electronics Dresden

TU Dresden, Germany

jeronimo.castrillon@tu-dresden.de

Abstract—The polyhedral model allows a structured way of defining semantics-preserving transformations to improve the performance of a large class of loops. Finding profitable points in this space is a hard problem which is usually approached by heuristics that generalize from domain-expert knowledge. Existing search space formulations in state-of-the-art heuristics depend on the shape of particular loops, making it hard to leverage generic and more powerful optimization techniques from the machine learning domain. In this paper, we propose a shape-agnostic formulation for the space of legal transformations in the polyhedral model as a Markov Decision Process (MDP). Instead of using transformations, the formulation is based on an abstract space of possible schedules. In this formulation, *states* model partial schedules, which are constructed by *actions* that are reusable across different loops. With a simple heuristic to traverse the space, we demonstrate that our formulation is powerful enough to match and outperform state-of-the-art heuristics. On the Polybench benchmark suite, we found the search space to contain transformations that lead to a speedup of 3.39x over LLVM O3, which is 1.34x better than the best transformations found in the search space of *isl*, and 1.83x better than the speedup achieved by the default heuristics of *isl*. Our generic MDP formulation enables future work to use reinforcement learning to learn optimization heuristics over a wide range of loops. This also contributes to the emerging field of machine learning in compilers, as it exposes a novel problem formulation that can push the limits of existing methods.

Index Terms—polyhedral optimization, loop scheduling, machine learning, reinforcement learning, PolyGym

I. INTRODUCTION

In most compute-intensive applications, a significant amount of the execution time is spent in loops. Loop optimization has thus received plenty of attention in the compiler community, where one finds as diverse strategies for optimization as are the loops themselves. Among them, polyhedral compilation is an important class of strategies for loop optimization which focuses on a particular class of (nested) loops with regular properties [1]–[3]. While this significantly restricts the family of loops that can be optimized, polyhedral compilation methods are very relevant in practice. Compilers for important domain-specific languages like TensorComprehensions [4] and Tiramisu [5] build atop these methods. Polyhedral compilation methods have also been used for efforts on recognizing *motifs*, or well-known high-level computations [6]. The models have

been extended beyond explicit loops to support recursive functions calls [7], or used to target emerging memory technologies by exploiting the detailed information in the model [8], [9].

With the polyhedral model we can find semantics-preserving transformations that significantly improve the loops’ performance. We do this by exploring constraints exposed by geometric properties of the dependencies between statements in different iterations of the loops. However, this exploration requires considering a prohibitively large space of possible code transformations. The two main methods to explore this space are heuristics, which use domain-specific models to directly find a good candidate transformation, and iterative meta-heuristics, which iteratively explore the space by evaluating multiple points and adapting the solutions. Model-based heuristics, like Pluto [10] and *isl* [2], are more applicable in practice, but the resulting code is consistently outperformed by significantly more time-costly iterative meta-heuristics, like [11], [12].




Fig. 1. PolyGym in the context of Reinforcement Learning

A promising strategy to bridge this gap is reinforcement learning (RL). RL has proven to be very successful for navigating vast discrete spaces, while training a decision policy that can later be leveraged to find good solutions directly [13], [14]. Precisely for this reason, RL is well-suited for learning heuristics in the spaces of compiler optimizations, seeking a sweet spot between applicable heuristics and iterative methods: While iterating the search space, a heuristic model is learned that can later be re-used in future iterations to lead to a profitable transformation.

The main challenge of using polyhedral scheduling as an RL environment is the lack of a suitable search space formulation. High-level loop transformations that result from polyhedral scheduling, such as loop fusion, fission or tiling, are ill-suited as actions in an RL environment. The first disadvantage concerns correctness: not every transformation can be applied to every schedule. This happens often and is computationally costly. The greater problem is that the

* Equal contribution

† Work done in part at the Chair for Compiler Construction, TU Dresden.

transformations themselves depend on the concrete instance of the problem. For example, loop fission needs to select the concrete statements and loop bounds, which depend on the concrete loop. If we encode these transformations as actions in a markov decision process (MDP) for RL, any policy for finding profitable schedules would be specific to that loop or kernel by the very structure of the formulation.

In this paper we propose a formulation that overcomes the limitations of transformations for an MDP to enable development of heuristics with optimization techniques. We utilize a known formulation [15], [16] that leverages Farkas’ lemma to construct an abstract space of exclusively valid schedules. We present a clear overview of the background from discrete geometry, emphasizing why the different steps in the construction are important for this use-case (Section II). Our MDP navigates this abstract space in an instance-independent fashion by leveraging its algebraic properties (Section IV). The resulting search space boasts the potential for achieving great performance and learning, a fact we demonstrate by traversing it with simple heuristics (Section V). Even without sophisticated optimization algorithms, simple heuristics that guide the iterative process outperform the state-of-the-art heuristics in isl [2] in the Polybench benchmark suite [17] in a small number of iterations.

Concretely, we make the following contributions:

- We define the search space of affine loop transformations in polytopes as a loop-agnostic MDP, exposing this problem to new learning algorithms and heuristic models in the future. Figure 1 shows PolyGym, our environment, in the context of an RL method.
- We show the performance potential of this search space by traversing it with simple heuristics.

Importantly, as well, the field of machine learning for compilers is an emerging field of study on its own [18]. RL algorithms are a particularly relevant class of methods for compiler optimization [19], and have successfully been applied to certain subproblems already [20]–[22]. Even in the industrial LLVM compiler, an RL-trained heuristic for inlining was recently integrated into the mainline version [23].¹ However, defining RL environments for problems in compiler optimization is challenging in general, for many of the same reasons it is challenging in the case of polyhedral optimizations. Defining the problem in a way that action policies can be learned in a uniform fashion, generalizing across multiple samples, is an important step to advance the field of machine learning in compilers as well.

We release our formulation as PolyGym, an open-source environment that can be readily used in future work to learn heuristic models.²

II. BACKGROUND: POLYHEDRAL COMPILATION

This section introduces the basics of the polyhedral model, based on the original work of Feautrier [15], [16], and the

¹<https://github.com/llvm/llvm-project/blob/main/llvm/lib/Analysis/InlineAdvisor.cpp>

²Source code available at: <https://github.com/tud-ccc/polygym>

algorithms for schedule space construction due to Pouchet and others [11], [24]. Consider the following C kernel for the matrix-vector multiplication $y = Ax$:

```

int i, j;
for (i = 0; i < N; i++)
S:   y[i] = 0;
for (i = 0; i < N; i++)
for (j = 0; j < N; j++)
T:   y[i] += A[i][j]*x[j];

```

The iterations of the loops in this kernel exhibit a regular structure. The set of points i, j of the values the variables i, j take during the execution of the loops is $I := \{i, j \in \mathbb{Z} \mid 0 \leq i, j \leq N - 1\}$. This kernel is what we call a Static Control Part (SCoP). A detailed description of SCoPs and how they can be constructed can be found in [3]. For the purposes of this paper however, it suffices to know that these programs have regular iteration spaces. Concretely, if we consider the iteration variables in the loops of the SCoP, they have a geometric interpretation as a lattice polytope. A polytope is the n -dimensional generalization of a polygon in 2 dimensions and polyhedral in 3 dimensions. The lattice contains a set of discrete points (e.g. integer values) in the interior of the polytope. For example, the set I of the example kernel can be seen as the discrete points in the lattice \mathbb{Z}^2 , restricted to an $N \times N$ square starting at the origin, as depicted in Figure 2.




Fig. 2. The lattice polytope interpretation of the iteration space of the matrix-vector multiplication kernel for $N=5$.

In the kernel, there are two distinct statements that are executed in the loop, which we labeled by S and T . The first statement $S : y[i] = 0$ does not actually depend on j , and is executed N times, which we denote by $S(i)$ for $i = 0, \dots, N - 1$. The statement $T : y[i] += A[i, j]*x[j]$ depends on both variables i, j and is executed N^2 times, which we denote by $T(i, j)$ for $i, j = 0, \dots, N - 1$. In particular, executing this kernel amounts to executing these $N^2 + N$ statements. Analyzing the statements we can see which statements have to be executed before which, and where it does not matter. Indeed, we know that for every i , the instance $S(i)$ has to be executed before all $T(i, j)$. Similarly, we need to execute $T(i, j)$ before we execute $T(i, j + 1)$ for all valid i, j .³ We can use the geometric interpretation

³If we know that addition is commutative, we can relax these conditions on the $T(i, j)$, but we ignore the commutativity of addition for this discussion.

as polytopes (in this example, polygons) to visualize these dependencies, as in Figure 3.

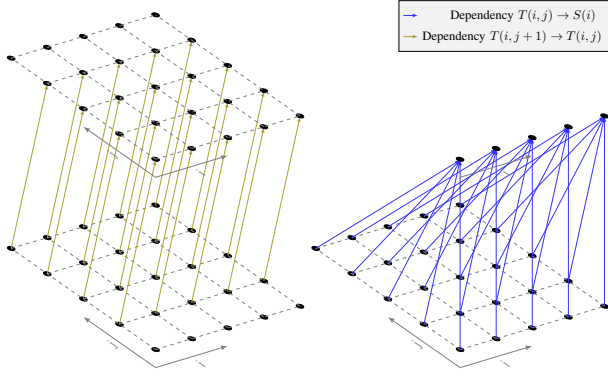


Fig. 3. The geometric interpretation of dependencies in the matrix-vector multiplication kernel for $N=5$.

Any permutation of the $N^2 + N$ statements $S(i), T(i, j)$ that respects the dependencies visualized in Figure 3 will produce the same (correct) result. They are all semantics-preserving permutations. For example, the original code executes the statements in the order $S(0), S(1), \dots, S(N-1), T(0,0), T(0,1), \dots, T(N-1, N-1)$. However, we could also execute them in the order $S(0), T(0,0), S(1), T(0,1), T(1,0), S(2), \dots, S(N-1), T(0, N-1), \dots, T(N-1,0), T(1, N-1), \dots, T(N-1, N-2), T(N-1, N-1)$, as this respects all dependencies. This order would translate to the code:

```

int i, j;
for (i = 0; i < N; i++) {
    y[i] = 0;
    for (j = 0; j <= i; j++)
        y[j] += A[j][i-j]*x[i-j]; }
for (i = 0; i < N; i++)
    for (j = i+1; j < N; j++)
        y[j] += A[j][i-j+N]*x[i-j+N];

```

The semantics-preserving permutations of the statements can be represented as *schedules*. Using the lexicographic ordering, we can encode schedules as vectors. For example, the original statement order for $N = 5$ could be encoded as the 1-dimensional vectors for $S(i) = (i)$, and $T(i, j) = 5 \cdot i + j + 5$ respectively, which results in the order $(0), \dots, (4)$ for S and $(5), \dots, (29)$ for T , encoding the order of the statement executions. The crucial observation is that these vectors are affine functions

$$\Theta_S : \mathbb{Z}^2 \rightarrow \mathbb{Z}^1, (i, j) \mapsto (1, 0, 0)(i, j, 1)^T,$$

$$\Theta_T : \mathbb{Z}^2 \rightarrow \mathbb{Z}^1, (i, j) \mapsto (5, 1, 5)(i, j, 1)^T.$$

This schedule is called the *identity schedule*, as it represents the original untransformed code.

In general, when the schedule transformations are affine linear functions, we can use a powerful result of discrete geometry called Farkas' lemma in its affine form [25] to characterize all such transformations [15]. The transformed schedule we

showed above, for example, has the order $0, 2, 5, 9, 14$ for the instances $S(i), i = 1, \dots, 4$, which is not an affine function of i . We can express the schedule as an affine function, however, if we use two-dimensional vectors $S(i) = (i, 0)$ and $T(i, j) = (i+j, i+1)$. These vectors, sorted lexicographically, give the order of the schedule above and are affine functions as well. When scheduling SCoPs, this characterization results in a set of linear inequalities describing the valid schedules [24]. This set of inequalities comes about by encoding each dependency between statements as a precedence constraint on the iteration spaces. Farkas' lemma then allows us to transform affine functions respecting these dependencies into this set of inequalities. For example, for the matrix-vector multiplication kernel, the valid 1-dimensional schedules are integer solutions to the inequalities on the 11 variables i_0, \dots, i_{10}

$$\begin{aligned} & i_2 > 0 \quad \text{and} \quad i_6 \geq i_3 \quad \text{and} \quad i_7 \geq i_4 \\ & \text{and} \quad i_8 \geq i_5 \quad \text{and} \quad i_7 \geq i_0 - i_1 + i_3 + i_4 - i_6 \quad (1) \\ & \text{and} \quad i_{10} > i_3 + i_4 + i_5 - i_6 - i_7 - i_8 + i_9. \end{aligned}$$

Geometrically, this set of inequalities can be equivalently considered as another lattice polytope, where each point in this polytope represents a valid schedule. Importantly, this polytope should not be confused with those for the iteration space or the dependencies, where the points represent the values of the iteration variables. Figure 4 depicts this set for one-dimensional schedules. The space of valid schedules is itself a polytope in 11 dimensions, corresponding to the 11 variables in Equation 1. As such, we depict it with a parallel plot, where each horizontal line in the grid is a dimension and each point has a different color. In fact, this lattice polytope has an infinite number of points; we show only the meet of the polytope with the fundamental region of the lattice ($\cong \mathbb{Z}^{11}$) in this plot. Each point represents a schedule, like the original schedule which is depicted explicitly in the figure. Note that one-dimensional refers to the schedule vectors, i.e. the images of $\Theta_S : \mathbb{Z} \rightarrow \mathbb{Z}^1$ and $\Theta_T : \mathbb{Z}^2 \rightarrow \mathbb{Z}^1$. This dimension is different from the points depicted, which encode the constraints as characterized by Farkas' lemma. The polytope corresponding to the constraints in Equation 1 includes both schedule functions, Θ_S and Θ_T .




Fig. 4. The geometric interpretation of the space of valid one-dimensional schedules in the matrix-vector multiplication kernel for $N=5$. The space is a lattice polytope on 11 dimensions, visualized as a parallel plot.

Any valid schedule can be expressed as a one-dimensional transformation function. The caveat is, however, that this transformation is not necessarily an affine function. As such, Farkas' lemma does not apply, and we cannot capture these schedules in the lattice polytope of valid one-dimensional

transformations. We can use more dimensions to describe more schedules as affine functions [16]. To define a k -dimensional schedule, we define k different polytopes P_1, \dots, P_k using Farkas' lemma, as described above for the one-dimensional case (cf. Figure 4). Each polytope P_d encodes the functions for the corresponding component, i.e. $(\Theta_S)_d, (\Theta_T)_d$. In the example of the two-dimensional schedule for the matrix-vector multiplication example, we have $(\Theta_S)_1(i, j) = i, (\Theta_S)_2(i, j) = 0$ and similarly, $(\Theta_T)_1(i, j) = i + 1, (\Theta_T)_2(i, j) = i + 1$. Each of these is an affine function of the inputs.

The characterization of the space of *all* possible multi-dimensional schedules is not as straightforward as it is in the one-dimensional case. For this we distinguish between two types of precedence constraints. For a dependency $S \rightarrow T$ we say that the polytope P_i carries the dependency *strongly* if it encodes that the value of $\Theta_S < \Theta_T$. If instead it only encodes the weaker constraint $\Theta_S \leq \Theta_T$, we say that P_i carries this dependency *weakly*. While this does not enforce the constraint, it ensures that this constraint is not violated either. A proper definition of these concepts can be found in [11]. For an example, consider the dependency $S \rightarrow T$ in the matrix vector multiplication. In the the two-dimensional schedule defined above, this dependency is carried weakly by the polytope P_1 corresponding to the first dimension, since $i \leq i + j$ but $i \not\leq i + j$ for $(i, j) \in I$. The dependency $T \rightarrow T$, on the other hand, is carried strongly by P_1 , since $i + j < i + (j + 1)$ for all $(i, j) \in I$. The polytope P_2 for second dimension then carries the dependency $S \rightarrow T$ strongly, since $0 < i + 1$ for all $(i, j) \in I$.

If we consider a polytope carrying all dimensions strongly, we leave out some possible schedules. In some cases, the polytope carrying all dimensions strongly might even be empty, an example of this is given in [16]. Instead, we can iteratively construct a multi-dimensional schedule by carrying only some dimensions strongly and ensuring all the other dependencies are carried weakly. Since we interpret schedules in lexicographical order, once a dependency is carried strongly by a dimension, all subsequent dimensions need not carry it weakly anymore. If we iteratively add dependencies to dimensions in this way, we will always find a multi-dimensional schedule for every SCoP [16]. However, the order in which we let dimensions carry the dependencies strongly affects the space of possible schedules. In Section IV we define a general algorithm with free parameters to explore all possible such constructions.

To find profitable schedules for a SCoP, we need to choose a schedule after having constructed the space of all possible (multi-dimensional) schedules. For this, we can independently consider the polytopes P_d corresponding to each dimension d of the schedule space. Points in this schedule space are integer solutions to a set of inequalities, but equivalently, they are points with integer values in a special linear combination of terms:

$$p = \sum_i^s \lambda_i v_{i,d} + \sum_i^t \alpha_i r_{i,d}, \lambda_i, \alpha_i \in \mathbb{R}_{\geq 0}, \sum_i^{k_1} \alpha_i = 1. \quad (2)$$

Equation 2 describes a convex combination of terms $v_{i,d}$ called vertices, and a positive linear combination of terms $r_{i,d}$ called rays, where the index d corresponds to the dimension d . Every such combination with integer values, i.e., where $p \in \mathbb{Z}^n$, is a point in the lattice polytope of valid schedules. In other words,

$$P_d = \mathbb{Z}^n \cap \left\{ \sum_i \lambda_i v_{d,i} + \sum_i \alpha_i r_{d,i} \mid \lambda_i, \alpha_i \geq 0, \sum_i \lambda_i = 1 \right\}.$$

This representation can be calculated with Chernikova's algorithm [26] and is equivalent to the constraint set representation (as e.g. in Equation 1). We call the set $G_d = \{v_{d,1}, v_{d,1}, \dots, r_{d,1}, \dots\}$ the set of *generators* of P_d . Choosing a point $p_d \in P_d$ in the corresponding polytope for every dimension thus defines a multi-dimensional schedule.

As we have seen, a confusing aspect of the polyhedral model is how there are several lattice polytopes associated to a SCoP region: the iteration space, geometric interpretations of the dependencies, and the spaces of valid schedules. Figure 5 gives an overview of the model and the different lattice polytopes, as well as the relationships between them, on the basis of the matrix-vector multiplication example.

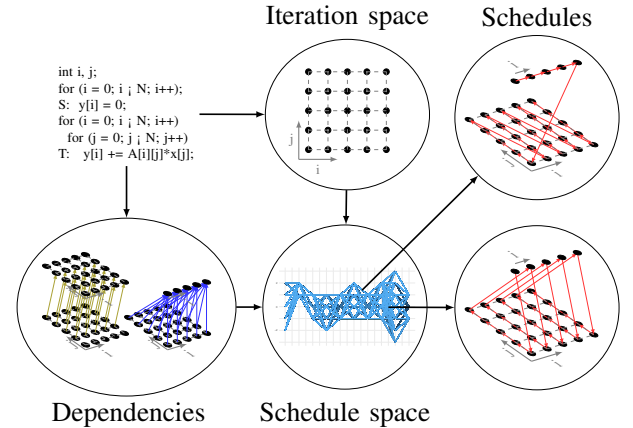


Fig. 5. An overview of schedules in the polytope model

Optimizing a loop amounts to finding a good schedule for the loop and a target architecture. Since the construction of a multi-dimensional space of valid schedules is not unique, this can be considered as two distinct steps: Constructing a schedule space (of valid schedules) and exploring it to find a profitable schedule. An important contribution of this paper are MDP models for finding such a schedule based on these two steps, as we describe in Section IV. In Section III we describe how state-of-the-art methods go about finding them.

III. RELATED WORK

There are two main families of methods for schedule optimization in the polyhedral model: iterative methods that are often guided by meta-heuristics and model-driven heuristics. Iterative models evaluate multiple schedules iteratively by generating code and benchmarking it, eventually returning the best schedule [24]. The iterative exploration of such large spaces is typically guided by meta heuristics such as genetic

algorithms, where the results of evaluations influence future sampling points in the space [11], [12], [27]. Since these methods require multiple executions of the program, they are time-consuming in practical compilation flows and need additional benchmark tooling. In contrast to this, model-driven heuristics like Pluto [10] and isl [2] propose a reasonably good solution directly without benchmarking. Using predefined assumptions about communication and locality, these approaches formulate the optimization problem as an Integer Linear Program. The assumptions require domain knowledge that may not generalize to other architectures, and are usually outperformed by their iterative counterparts.

Learned heuristics for schedule optimization have been investigated in the scope of domain-specific languages (DSLs) that do loop-heavy computations, like TensorComprehensions for machine learning [4] or Halide [28] for image processing. A recent extension to Halide [29] uses machine learning on a set of program features to build a cost model. The authors use the cost model to steer a beam search to find profitable schedules iteratively without benchmarking the programs. In similar fashion, TVM [30] uses a learned cost model to estimate the performance of a candidate schedule in its form of a program representation. TensorComprehensions, Halide and TVM work on a similar space of loop transformation primitives. Compared to the search space presented in this work, their space does not include affine transformations that lead to potentially fast loop variants. These methods are also restricted to programs expressed in the corresponding DSLs. In contrast, this work presents a search space that is applicable to general C code.

Recent work using Tiramisu [31] overcomes some of these limitations by considering polyhedral optimizations as well. In addition to that, they use machine learning to estimate the execution time of a loop after the transformation, but use an iterative monte-carlo tree search method to explore the possible transformations. Our contribution is orthogonal to their approach. Unlike our formulation, however, the search-space in [31] includes incorrect transformations and, more importantly, is specific to the loop being optimized. This makes it more challenging to design and learn a policy that generalizes across different loops.

RL has successfully been applied to certain subproblems of compiler optimizations. NeuroVectorizer learns a model that aims to predict a solution in the search space of vectorization factors for a given loop [20]. Other work has focused on maximizing speedup [22] and code size [21] on the search space of analysis and transformation passes in LLVM, which is enormously large. A smaller search space for code size optimization is used by [23] to predict whether or not to inline a function. More generally, machine learning has recently been applied successfully to multiple tasks in compilers [32]–[35]. While these works have advanced the models of code, they have focused on rather coarse-grained tasks, like mapping OpenCL kernels to a CPU or GPU. An important reason for this is that formulating problems in compilers in an instance-independent fashion is challenging. Finding optimal

loop schedules is a more subtle problem than, for instance, selecting a particular GPU or CPU for the execution of a kernel. This paper contributes to the field of machine learning in compilers by providing more challenging, yet lucrative task that can better challenge the models of code.

IV. POLYGYM’S MARKOV DECISION PROCESS

In this section we present our formulation of the problem of finding a profitable schedule as a markov decision process (MDP). This enables us to explore the space of possible schedules for a SCoP in the context of RL. In particular, the actions in our formulation are independent of the particular instance, i.e. the SCoP being optimized. If, instead, the actions depended on the particular instance, then a different model needs to be used for every instance. This property is thus a necessary requirement for a heuristic to learn to generalize over different instances.

As explained in Section II, the problem of finding a profitable schedule can be divided in two stages: the construction of the schedule space and the selection of a schedule. We define two different MDPs, one for each subproblem. From the point of view of the formulation, these are two distinct MDPs, even if the space of the state space of the second MDP depends on the options chosen in the first. Thus, a single combined RL agent can operate on the combined space, or alternatively, two distinct agents can be trained to explore them.

A. Schedule space construction

We look for a schedule space in its general form, and thus define it as a multi-dimensional space [11], [16]. As illustrated in Algorithm 1, we construct a k -dimensional schedule space iteratively, by going through the dimensions of the space. For each dimension, we decide which dependencies to include as strong dependencies, until we have selected all dimensions. Crucially, the function that decides this, `select_dependency`, is left as an unspecified, free function. It then calculates the polytope of possible schedules strongly satisfying these dependencies and weakly satisfying the remaining unselected dependencies. Using Chernikova’s algorithm [26] it calculates generators as vertices and rays for this polytope. This algorithm goes back to the principles outlined by Feautrier [16], which is the same basis for the heuristic in Section 3.2 in [11] and Algorithm 1 in [12]. Algorithm 1 generalizes these principles by leaving the decision function `select_dependency` unspecified, instead of proposing a concrete heuristic. Note that we write `select_dependency(d)` to specify that this function depends on the representation of the dependency d . In general, this decision could also depend on other parameters like properties of the SCoP.

This simple change, making `select_dependency` a free function, has profound consequences. The greedy heuristics of [11], [16] construct one deterministic schedule space, which fundamentally limits the space of possible schedules that can be found with the method. On the other hand, the randomized heuristic of [12] could in principle produce any schedule space.

`next_dim`, `next_dep`, `select_dep`. The first action, `next_dim`, increases the dimension, which is the first component $(1, \dots) \rightarrow (2, \dots)$. Then, the `next_dep` action changes the second index, indicating the current dependency: $(2, 1, 0, 0) \rightarrow (2, 2, 0, 0)$. Since no dependency has been selected so far, the value of j is 2, corresponding to the first case outlined above, yielding $n = 1$. After increasing the dimension with a `next_dim` action, the next `next_dep` action again changes the dependency. This time there is no dependency with index > 2 , and thus $j = 1$, by the second case, corresponding to $n = -1$. A second `next_dep` returns the second index, marking the dependency, to 2, returning to the state $(3, 2, 0, 0)$. At this point the `select_dep` adds an index to the current dependency (2), marking the current dimension to carry this dependency strongly: $(3, 2, 0, 3)$. The dependencies $S \rightarrow T$ and $T \rightarrow T$ are sorted, in that order. This process continues until it reaches the state $(4, 1, 4, 3)$. This corresponds to a 4-dimensional space, as indicated by the first component. As indicated by the third component⁵, the third dimension strongly carries the first dependency, $S \rightarrow T$, and the fourth dimension strongly carries the second dependency $T \rightarrow T$, as indicated by the last component. In particular, the accepting state itself uniquely defines the schedule space, as the order in which the dependencies are selected within a dimension is irrelevant to the construction, it only matters which dimensions specifically carry the dependencies strongly.

B. Schedule space exploration

Once the schedule space is generated, we look into how to find profitable schedules in this space. From Algorithm 1, the schedule space is given by a vector $S = (G_1, G_2, \dots, G_k)$, where each G_d corresponds to the generators of a lattice polytope for the d -th dimension. A valid schedule consists precisely of a point in each of the lattice polytopes corresponding to the dimensions. To construct each of these points we need to consider the generators G_d . These are given as a set of vertices v_1, \dots, v_s and rays r_1, \dots, r_t as explained in Section II. An arbitrary point p is in the polytope iff it can be written as a convex combination of vertices and a positive linear combination of rays, i.e.

$$p = \sum_{i=1}^s \lambda_i v_i + \sum_{i=1}^t \alpha_i r_i,$$

where $\lambda_i \geq 0$ for all $i = 1, \dots, s$ and $\sum_{i=1}^s \lambda_i = 1$ and $\alpha_i \geq 0$ for all $i = 1, \dots, t$. Note that some work uses a third generator type, lines. This can always be converted to an equivalent set of generators as we define here. We choose to have only rays instead of rays and lines as it is more uniform this way, making this representation more amenable for RL. Schedules correspond to the points in the lattice polytope, which is a subset of this general polytope. This means that p has to have integer coefficients.

⁵Since the first two indicate the current dimension and dependency.

To formulate this as an MDP, we introduce Algorithm 2 which generates points in the polytopes following the same principles as [11], [12], [15], [16], [24], [27]. Once again, our algorithm leaves an unspecified, free function `select_coeff`. The goal of this function is to find the values of the coefficients $\lambda_i, i = 1, \dots, s$ and $\alpha_i, i = 1, \dots, t$. We do this by iterating over all vertices and rays and selecting a coefficient in this iteration. Like the authors in [12], we use a correction step, multiplying by the least common denominator (LCD) in the end to ensure the point is in the lattice polytope, i.e. has integer coefficients. While this can be avoided through the design of the `select_coeff` function, this correction step can simplify the design of the function.

Algorithm 2 General exploration of the schedule space

input: A schedule space $S = (G_1, G_2, \dots, G_k)$ as a vector, where each G_i corresponds to the generators of a lattice polytope for the i -th dimension.
output: A point in the schedule polytope for each dimension
 $p = (p_1, \dots, p_k)$
1: **for** $i \in \{1, \dots, k\}$ **do**
2: $v_1, \dots, v_s, r_1, \dots, r_t \leftarrow G_i$
3: $p_i \leftarrow 0$
4: **for** $x \in \{v_1, \dots, v_s, r_1, \dots, r_t\}$ **do**
5: $p_i \leftarrow p_i + \text{select_coeff}(x) \cdot x$
6: **if** $p_i \in \mathbb{Q}^n \setminus \mathbb{Z}^n$ **then**
7: $p_i \leftarrow \text{LCD}(p_i) \cdot p_i$
return (p_1, \dots, p_k)

To define a corresponding MDP, we start by defining the state space. For an integer $N > 0$, we define the space as the finite set⁶ of the coefficients

$$S_{\text{expl}} = \{(\lambda_{1,1}, \dots, \lambda_{1,s_1}, \alpha_{1,1}, \alpha_{1,t_1}, \dots, \lambda_{k,1}, \dots, \alpha_{k,t_k}) \mid \lambda_{i,j}, \alpha_{i,j} \in \{0, \dots, N, \perp\} \text{ for all } i, j\}$$

Note that since we have multiple polytopes P_1, \dots, P_k , corresponding to the multiple dimensions, we use two indices for the generators, where the first index d corresponds to the polytope P_d and the second index iterates over the generators in G_d . This can also be conceptually understood as unrolling the two loops in Algorithm 2 in the state space. We use the symbol \perp to mark coefficients that have not yet been selected, as this is distinct from selecting 0 as a coefficient. Choosing the coefficients $\lambda_i \in \{0, \dots, N\}$ means that in most cases, $\sum_{i=1}^s \lambda_i > 1$. In that case we norm the coefficients by building the convex combination as $\sum_{i=1}^s \lambda_i v_i / \sum_{i=1}^s \lambda_i$. Additionally, when $s = 1$ it is clear that $\lambda_1 = 1$, since for a single point there is only one possible convex combination. We thus remove the corresponding coefficients entirely from the state space⁷.

Since we need to choose a coefficient for each term, we do not include actions to steer the exploration as we did for the schedule space construction. Thus, the action

⁶This is thus parameterized by N . Parameters like these are sometimes also called hyper-parameters, especially in the context of machine learning.

⁷This was the case for almost all examples we evaluated in this paper.

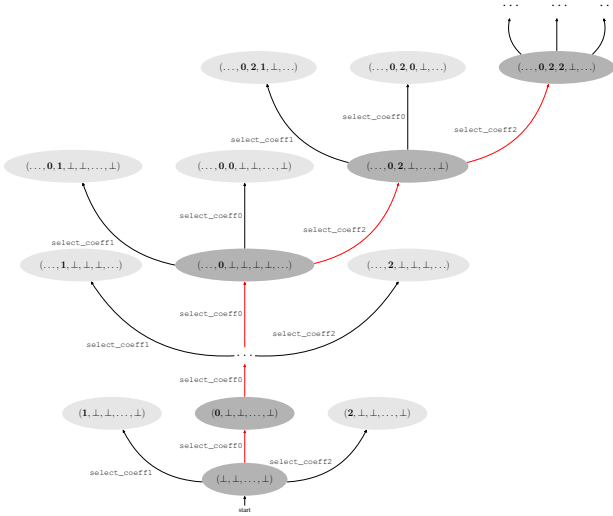


Fig. 7. An example of the schedule space exploration MDP and a concrete action sequence for the matrix-vector multiplication example.

space corresponds directly to the function `select_coeff`. We define actions `select_coeff0, ..., select_coeffN` accordingly. We set

$$\mathbf{P}_{\text{expl}}((a_1, \dots, a_i, \perp, \dots, \perp), \text{select_coeff}X, (a_1, \dots, a_i, X, \perp, \dots, \perp)) = 1,$$

and for all other states $s, s' \in S_{\text{expl}}$ and actions $a \in \text{Act}_{\text{expl}}$, we set $\mathbf{P}_{\text{expl}}(s, a, s') = 0$. The initial state corresponds to the starting configuration with no coefficients defined, i.e., (\perp, \dots, \perp) .

Similar to the space construction, a crucial aspect of this formulation is that the actions are independent of the SCoP. This way, a heuristic model can learn to find good schedules by selecting the correct coefficients.

Consider again the example of the matrix-vector multiplication kernel. With the schedule space constructed as depicted in Figure 6, the schedule space has four dimensions, where the third carries the dependency $S \rightarrow T$ strongly and the fourth one $T \rightarrow T$. This yields four polytopes, one for each dimension, which after using Chernikova’s algorithm can be generated each by a single vertex and 10, 10, 10, 13 rays respectively. These all live in a 7 dimensional vector space. Note that neither the 4 dimensions of the schedule space nor the 7 dimensions of the polytopes for each schedule space dimension have a direct interpretation in terms of the loop bounds. They represent loop schedules according to Farkas’ lemma (cf. Section II) and they are not intuitively easy to understand in terms of the loops’ abstract syntax tree (AST). As mentioned above, the coefficient selection for the vertex is not part of the coefficient selection in the state space of the MDP, since it is a single vertex in all four dimensions.

Figure 7 shows the beginning of a sample run through this space, given by the following action sequence (omitting the `select_coeff` part of the name): 0, 0, 0, 0, 0, 2, 2, 1, 0, 2, 0, 0, 0, 0, 2, 0, 0, 1, 0, 2, 0, 0, 0, 1, 0, 2, 0, 0, 0, 0, 0,

0, 2, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0. The coefficients of the rays are selected dimension by dimension: $\alpha_{1,1} = 0, \dots, \alpha_{1,6} = 2, \alpha_{1,7} = 2, \alpha_{1,8} = 1, \dots, \alpha_{1,10} = 2, \alpha_{2,1} = 0, \dots, \alpha_{4,13} = 0$. After selecting all coefficients with that sequence, the resulting points are given by:

$$\begin{aligned} p_1 &= (0, 0, 0, 0, 0, 0, 0) + (0, 0, 0, 0, 0, -1, 0) \\ &\quad + 2 \cdot (0, 0, 0, 0, 0, -1, -1) + 2 \cdot (0, 0, 0, -1, -1, 0, 0) \\ &\quad + (-1, -1, 0, 0, 0, 0, 0) = (-1, -1, 0, -2, -2, -3, -1) \\ &\quad \vdots \\ p_4 &= (0, 0, 1, 0, 0, 0, 0) + 2 \cdot (0, 0, 1, 0, 0, 0, 0) \\ &= (0, 0, 3, 0, 0, 0, 0). \end{aligned}$$

We see how this MDP can produce concrete schedules. Finally, this points can be translated into a valid transformation $(i, j) \rightarrow (k, l)$ as explained in Section II. This translates to the following transformed C kernel:

```

if (N >= 1)
  for (int k = -N; k <= 0; k += 1) {
    if (N + k >= 1)
      for (int l = 0; l < N; l += 1)
T:      y[-k] += y[-k][l]*x[l];
    if (k <= -1)
S:      y[-k - 1] = 0;
  }

```

Note again that the correspondence between the representation as points in the schedule space and the transformed kernel is not directly visible. The dimensions of a multidimensional schedule define how the individual points in the schedule are ordered (lexicographically), it does not directly translate e.g. to loop nesting.

This concrete example was in fact found by a simple heuristic, with a bias towards defining coefficients to be 0. We describe the heuristic in Section V.

C. Rewards

The rewards of an MDP define a feedback loop – the quality of an action taken in a given state. This is needed for a learning algorithm so it can decide about what action is favorable to take in a given state and adjust the agent’s model in case the action was unfavorable. Reinforcement Learning can then be used to train an agent that navigates an MDP. Recent work has shown that it is possible to use Deep Neural Networks as agent models that are able to generalize over the large search space and able to learn to make good decisions on in previously unseen states [13], [14].

Figure 8 shows the flow with two MDP models in PolyGym in a RL context. Note that PolyGym does not define an agent, it only defines the environment to learn with such an agent. The two separate MDPs produce a single schedule, which is then compiled with LLVM/Polly and executed, to determine the reward. There is no immediate reward to the construction of the schedule space, nor to the partial states in the exploration step, since these do not define an actual schedule. As such, we define the rewards for both individual MDPs uniformly, based on the final constructed schedule, as follows. If the action leads

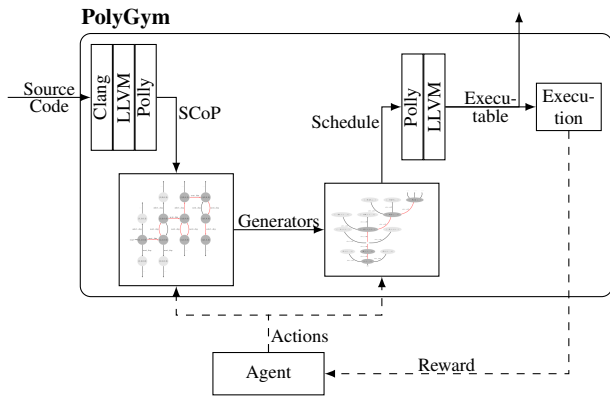


Fig. 8. The PolyGym flow

to a complete schedule, which is a terminal state, then we give the speedup over a traditional optimizing compiler, e.g., using LLVM with the `-O3` flag, as a sample-independent metric as reward. If it leads to an incomplete schedule, we give 0 as reward. Finally, if it leads to any invalid state, we give negative reward. Invalid states can occur when constructing an empty polytope in the MDP of the schedule space construction. They can also occur in the exploration MDP, when there are multiple vertices in a polytope and all vertex coefficients are selected to be 0 (leading to no convex combination).

In this way, (positive) rewards are only given for the final construction which obtains a valid schedule. This constitutes what is called a sparse reward setting. For example, the schedule for the matrix-vector multiplication kernel⁸ found by the examples described in Figures 6 and 7, the resulting kernel runs an impressive $140\times$ faster than the original kernel compiled with `-O3`⁹. This level of improvement is thus not easy to obtain in more complex kernels and larger input sizes. Thus, the final action `select_coeff0` (not shown in Figure 7 for space reasons) has a reward of 140, while all other actions have a reward of 0. State-of-the-art methods in RL trickle down this reward to the other actions [13], [14]. Importantly, this reward should also be trickled down to the schedule space construction in a learning setting, even if the MDP is conceptually a different one.

D. Limitations

While our formulation of the schedule search problem as an MDP with an instance-independent action set enables reinforcement learning for this problem, it still has some limitations.

An important limitation is that some accepting states do not yield schedules, as discussed previously. There are two possibilities of this happening. The first corresponds to the schedule space construction. Not every schedule space constructed with Algorithm 1 actually has schedules. If too many dependencies are carried strongly in a dimension, the intersection of the

corresponding polytopes can be empty, leading to a space without any valid schedules. The randomized construction of [12], similarly, does not provide such a (formal) guarantee. In practice, however, biasing the construction towards many dimensions, as they do, yields non-empty spaces. We are confident this fact can easily be learned by pertinent methods.

The second possibility has to do with the vertices in a generating set for a polytope. If all vertices have a coefficient of 0, Algorithm 2 does not produce a schedule, since the point constructed is not in the lattice polytope. The authors in [12] choose the vertex at random. We cannot do the same, as it breaks the MDP abstraction, where the reward of a state is non-deterministic.¹⁰ We did not run into this problem in any example considered in this paper. Having more than one vertex seems to be an extremely rare occurrence in practice. This is both a benefit and a problem: While we mostly do not have to deal with this problem, if at some point we ever do, we will probably not have enough samples to learn to deal with it.

An additional limitation corresponds to the exploration phase. There are two ways in which we exclude some points. By choosing a finite N as a hyper-parameter, we technically limit the possible coefficients of the schedule space. This is necessary, however, since otherwise there is an infinite amount of possible schedule polytopes¹¹. Similarly, the fact that we use integer coefficients and calculate the LCD in Algorithm 2 might exclude some points that would otherwise be found with rational coefficients. With an unbounded N this problem would not exist. Consequently, a larger N mitigates it.

While the limitations discussed here technically exclude some solutions or may make algorithms 1 and 2 fail, they do so only in rare corner-cases. Moreover, our reward design and the hyper-parameter N allow us to avoid or at least mitigate these limitations in those rare corner-cases.

Finally, as is commonplace in polyhedral compilation, our method suffers from important scalability limitations. Polyhedral methods are known to struggle with SCoPs that have a large number of statements and dependencies. There are different established methods to deal with scalability issues, at different levels. Incremental scheduling approaches work by splitting the connected components of the dependency graph of statements, splitting the scheduling problem into subproblems [36], [37]. While these approaches work on the linear-programming-based heuristic of schedulers, some of the core principles could be modified and applied to PolyGym in future work. Grouping approaches, on the other hand, aim to reduce the number of statements by grouping them into super-statements [3], [38], [39]. While, strictly speaking, this limits the schedule space, clever groupings can improve the scalability without removing the most profitable schedules. These heuristics are orthogonal to PolyGym and could be leveraged to improve the scalability.

⁸For the vector dimension $N = 2000$, with data generated procedurally using the Polybench tools.

⁹The full setup is described in Section V.

¹⁰Since the execution times of the SCoPs are non-deterministic as well, this is the case for all rewards. However, if the statistical variance of execution times is not negligible, then the whole problem of selecting an optimal schedule is ill-posed in the first place. This is not the case in practice.

¹¹Many of these are equivalent, since the number of orderings is finite

TABLE I
SEARCH SPACE OF ISL-POLLY12

Configuration	Values
-polly-opt-optimize-only	{all, raw}
-polly-opt-fusion	{min, max}
-polly-opt-maximize-bands	{yes, no}
-polly-opt-outer-coincidence	{yes, no}
-polly-opt-max-constant-term	{-1, ..., 499, 500}
-polly-opt-max-coefficient	{-1, ..., 499, 500}

V. EVALUATION

In this section we evaluate the performance potential of the search space of our MDP formulation by exploring it with simple heuristics. We compare the best schedules we find in a fixed number of explorations with the ones from the state-of-the-art heuristic implemented in isl and Polly in recent versions. In the following, we describe the experimental setup we used and discuss the results obtained.

A. Experimental setup

We implemented the MDPs of the schedule space as an OpenAI Gym Environment, so it is conveniently usable for future machine learning heuristics [40]. The algorithms are implemented in Python, relying on the isl binding `islpy`¹² for operations on integer sets. To transform the polytopes into the generator representation, we used Polyite’s implementation of the Chernikova algorithm¹³. Finally, we use LLVM and Polly in version 3.9 to extract polyhedral representations of SCoPs and to transform them according to the computed schedules.

All schedules were evaluated on the same system, consisting of an AMD 3960X with 24 cores and 48 threads, organized in 8 units, each with 3 cores. The units have a dedicated 16 MB L3 cache. Each core has a dedicated 32 KB L1 and a 512 KB L2 cache. The benchmarks are executed in a single unit, that is, using at most 3 cores (and 6 threads). We do so because of the low interference between units, which allows us to run up to four experiments in parallel without adding too much noise to the observations. The configuration of four units displayed negligible side-effects in preliminary experiments.

For our evaluation we use the Polybench benchmark suite [17] in version 4.1, which consists of 30 numerical computations from various domains, such as image processing, statistics, and linear algebra. Each of the benchmarks contains a SCoP kernel, i.e. a loop nest that satisfies the conditions of the polyhedral model. We had to exclude the `ludcmp`, `heat-3d`, and `adi` kernels because of performance issues, which we will discuss in Section V-C. We use the default dataset size LARGE. Since programs generated by unfavorable schedules could potentially run for a long time, we set a timeout that is proportional to the LLVM O3 runtime by a factor of 10 when executing.

We evaluate the search space of PolyGym by considering the quality of the schedules and use the performance obtained by optimizing compilers and isl

as reference points. As standard optimizing compiler we use LLVM in version 12 with the `-O3` flag (LLVM O3). In the experiments, we used 4 measurements and report the minimum to eliminate measurement inaccuracies. All remaining experiments were run with the polly configuration flags `-polly-parallel=true` `-polly-vectorizer=none` `-polly-tiling=true` `-polly-default-tile-size=64`. We also compared our search space to the search space of isl, as given by the parameters shown in Table I.

B. Result analysis

To analyze the profitable schedules included in the space spawned by our MDP formulation, we generate 1000 schedules per benchmark kernel, biasing the exploration of the schedule space towards the `select_coeff0` action. This bias results in overall less complex schedules. Figure 9 shows the distributions of speedups by individual samples, along with the performance of isl and LLVM O3 for reference. We can observe that the search space contains many profitable schedules. Many benchmarks, like `2mm`, `doitgen`, `gemm`, `jacobi-2d` or `seidel-2d`, show a significant distribution of points better than those found by isl. This suggests that an agent could learn to achieve this better performance results without iteratively executing the kernel.

Figure 10 shows the best (maximum) measured speedups for individual kernels of the polybench suite comparing the search space of PolyGym with that of the isl heuristic tuning parameters. We also include the results of isl with the default parameters for reference. We find schedules with an overall speedup of $3.39\times$ over O3-clang12, which is $1.34\times$ faster than best schedules found in the isl-polly12 space, that have an overall speedup of $2.52\times$ over O3-clang12. The results are not directly comparable to Polyite [12], [27], because we use a different hardware system. More importantly, we do not propose these random searches as an iterative heuristic. Instead, the goal of this evaluation is to show the potential of the formulation for an RL agent in the future. This can enable an agent to learn to navigate this space without requiring an iterative execution.

We further analyze the influence that a potential heuristic has on the sampling process. We employ different trivial heuristics for demonstration.

- In *bias_select_dep*, we bias the schedule space construction phase towards the `select_dep` action, which results in less schedule dimensions.
- In *bias_coeff_0*, we bias towards the `select_coeff0` action in the schedule space exploration phase.
- In *uniform*, we select actions uniformly at random.

Figure 11 shows the performance of the PolyGym search space with different heuristics on the Polybench suite. The plot shows the geometric mean across the different benchmarks as the number of schedules sampled increases. We see that different heuristics impact the space exploration achieving different performances overall. All evaluated heuristics outperform the schedules found by isl-polly3.9 and isl-polly12 in less than 40

¹²<https://github.com/inducer/islpy>

¹³<https://github.com/stganser/chernikova>

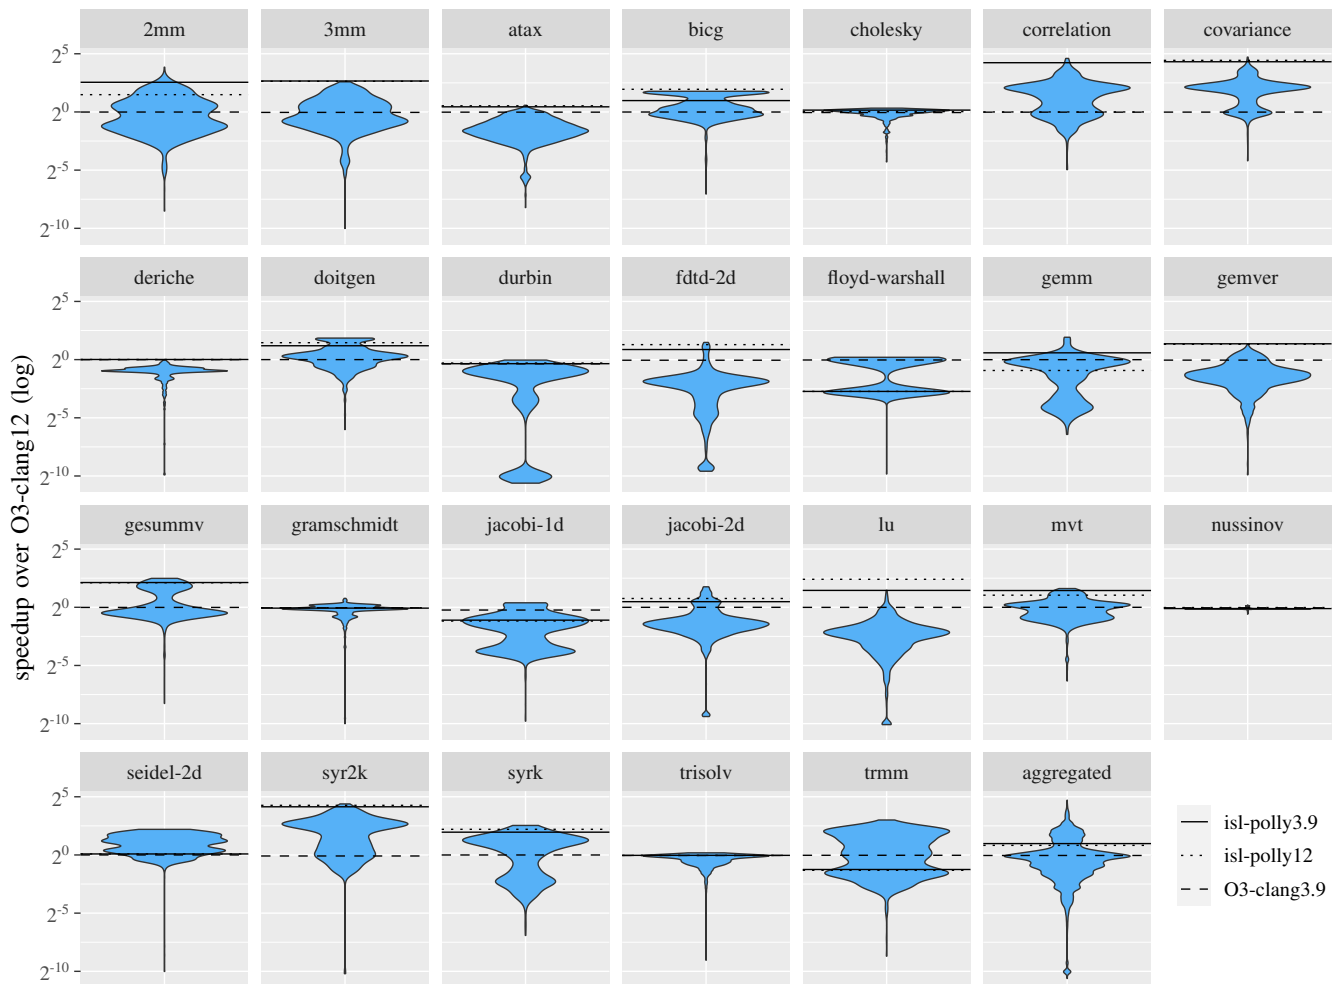


Fig. 9. Violin plot of the distribution of speedups by benchmark, with speedups as reference lines (geometric mean, 4 executions).

sampling iterations. Most notably, when biasing the selection towards the coefficient 0 ($bias_coeff_0$), the cross-over is at 11 iterations for `isl-polly12` and after 13 iterations for `isl-polly3.9`. Crucially, this shows that there is potential to learn to explore this space efficiently. All of the evaluated heuristics, however, are very simple and make not use of the MDP’s states. In future work, this heuristic can be a Deep Neural Network that is optimized to take the best actions using Reinforcement Learning algorithms.

C. Scalability

As pointed out in Section IV-D, an important limitation of these methods is their scalability. We evaluate this by looking at the different kernels and compare the times needed to construct the search spaces while controlling for kernel complexity (i.e. the number of dependencies). Concretely, we measure the construction time for around 200 random schedules for each kernel and report the median value. To evaluate the scalability while factoring out implementation-specific details, we norm the values, using the kernel with the fastest median construction time, `gesummv`, as a baseline. The

results can be seen in Figure 12, where the colors of the bars represent the number of dependencies of the kernel. Moreover, the kernels are sorted by the number of dependencies. We can see that while other aspects obviously play a role, the number of dependencies and the overall complexity of the kernel significantly impact the construction time. Some kernels, like `deriche` or `durbin` have a median time that is over 1000 times larger than the baseline `gesummv`. In these two cases, the median value reported is the timeout value, meaning the actual construction time is even larger.

While the scalability is a limitation of this method, this is not unique to PolyGym, and is more generally a feature of Polyhedral compilation. As discussed in Section IV-D, methods to improve this are the focus of state-of-the-art research in Polyhedral compilation.

VI. CONCLUSIONS

In this paper we presented PolyGym – an environment for polyhedral schedule optimizations for reinforcement learning (RL), based on a generic, SCoP-independent markov decision process (MDP). This formalization enables future RL-based




Fig. 10. Speedups over O3-clang12 of individual kernels of Polybench (execution time of O3-clang12, in seconds, in brackets). PolyGym-best and isl-polly12-best show the best-performing schedule found in a search of 1000 iterations.

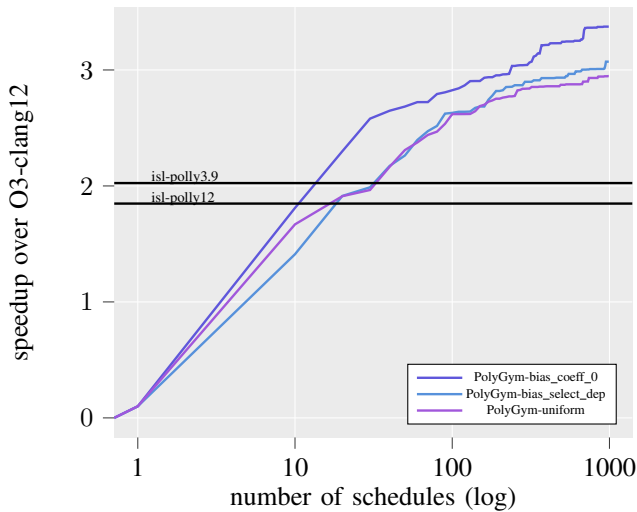


Fig. 11. Maximum speedup over LLVM O3 achieved by the different heuristics for action selection with increasing number of explored schedules. The plot depicts the geometric mean over all selected benchmarks of the Polybench suite. As reference, the aggregated speedups of isl over LLVM O3 in different version (isl-polly3.9, isl-polly12) is also included.

heuristics that learn models that produce profitable schedules directly, without iterative benchmarking. Using the PolyBench suite, we have shown that the search space contains learnable schedules that significantly outperform state-of-the-art heuristics. PolyGym is usable as an environment conforming to the Gym interface, which is widely-used by RL algorithms, enabling its integration with minimal effort. PolyGym thus enables research on RL-based optimizations without diving into the theory of polyhedral models, which imposes an otherwise high entry barrier. We believe that this will expose polyhedral optimizations as a learning problem, advancing the field of machine learning for compilers. In future work, we plan to investigate models and RL algorithms to learn profitable

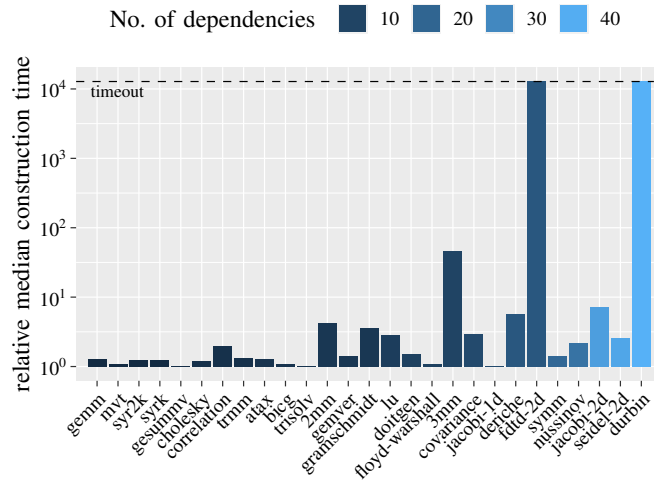


Fig. 12. Search-space construction time as measured on the different kernels, relative to `gesummv`.

schedules, designing reward functions for performance and other objectives such as energy consumption.

ACKNOWLEDGMENTS

We thank the anonymous reviewers and the shepherd whose comments and suggestions helped to improve this paper. This work was funded in part by the German Federal Ministry of Education and Research (BMBF) within the project ScaDS AI (BMBF 01IS18026B), the German Research Council (DFG) through the TraceSymm project (number 366764507) and the Studienstiftung des deutschen Volkes.

REFERENCES

- [1] S. Pop, A. Cohen, C. Bastoul, S. Girbal, G.-A. Silber, and N. Vasilache, "Graphite: Polyhedral analyses and optimizations for gcc," in *Proceedings of the 2006 GCC Developers Summit*, 2006, p. 2006.

- [2] S. Verdoolaege, “isl: An integer set library for the polyhedral model,” in *International Congress on Mathematical Software*. Springer, 2010, pp. 299–302.
- [3] T. Grosser, H. Zheng, R. Aloor, A. Simbürger, A. Größlinger, and L.-N. Pouchet, “Polly-polyhedral optimization in llvm,” in *Proceedings of the First International Workshop on Polyhedral Compilation Techniques (IMPACT)*, vol. 2011, 2011, p. 1.
- [4] N. Vasilache, O. Zinenko, T. Theodoridis, P. Goyal, Z. Devito, W. S. Moses, S. Verdoolaege, A. Adams, and A. Cohen, “The next 700 accelerated layers: From mathematical expressions of network computation graphs to accelerated gpu kernels, automatically,” *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 16, no. 4, pp. 1–26, 2019.
- [5] R. Baghdadi, J. Ray, M. B. Romdhane, E. D. Sozzo, A. Akkas, Y. Zhang, P. Suriana, S. Kamil, and S. Amarasinghe, “Tiramisu: A polyhedral compiler for expressing fast and portable code,” in *2019 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, 2019, pp. 193–205.
- [6] R. Gareev, T. Grosser, and M. Kruse, “High-performance generalized tensor operations: A compiler-oriented approach,” *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 15, no. 3, pp. 1–27, 2018.
- [7] S. Kobeissi, A. Ketterlin, and P. Clauss, “Rec2poly: Converting recursions to polyhedral optimized loops using an inspector-executor strategy,” in *International Conference on Embedded Computer Systems*. Springer, 2020, pp. 96–109.
- [8] K. Vadivel, L. Chelini, A. BanaGozar, G. Singh, S. Corda, R. Jordans, and H. Corporaal, “Tdo-cim: Transparent detection and offloading for computation in-memory,” in *2020 Design, Automation Test in Europe Conference Exhibition (DATE)*, 2020, pp. 1602–1605.
- [9] A. A. Khan, H. Mewes, T. Grosser, T. Hoefler, and J. Castrillon, “Polyhedral compilation for racetrack memories,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD). Special issue on Compilers, Architecture, and Synthesis of Embedded Systems (CASES’20)*, vol. 39, no. 11, pp. 3968–3980, Oct. 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/9216560>
- [10] U. Bondhugula and J. Ramanujam, “Pluto: A practical and fully automatic polyhedral parallelizer and locality optimizer,” 2007.
- [11] L.-N. Pouchet, C. Bastoul, A. Cohen, and J. Cavazos, “Iterative optimization in the polyhedral model: Part ii, multidimensional time,” *ACM SIGPLAN Notices*, vol. 43, no. 6, pp. 90–100, 2008.
- [12] S. Ganser, A. Größlinger, N. Siegmund, S. Apel, and C. Lengauer, “Iterative schedule optimization for parallelization in the polyhedron model,” *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 14, no. 3, pp. 1–26, 2017.
- [13] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, “Human-level control through deep reinforcement learning,” *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [14] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, “Mastering the game of go with deep neural networks and tree search,” *nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [15] P. Feautrier, “Some efficient solutions to the affine scheduling problem. i. one-dimensional time,” *International journal of parallel programming*, vol. 21, no. 5, pp. 313–347, 1992.
- [16] —, “Some efficient solutions to the affine scheduling problem. part ii. multidimensional time,” *International journal of parallel programming*, vol. 21, no. 6, pp. 389–420, 1992.
- [17] L.-N. Pouchet *et al.* (2012) Polybench: The polyhedral benchmark suite. [Online]. Available: <http://web.cs.ucla.edu/~pouchet/software/polybench/>
- [18] Z. Wang and M. O’Boyle, “Machine learning in compiler optimization,” *Proceedings of the IEEE*, vol. 106, no. 11, pp. 1879–1901, 2018.
- [19] H. Leather and C. Cummins, “Machine learning in compilers: Past, present and future,” in *2020 Forum for Specification and Design Languages (FDL)*. IEEE, 2020, pp. 1–8.
- [20] A. Haj-Ali, N. K. Ahmed, T. Willke, Y. S. Shao, K. Asanovic, and I. Stoica, “Neurovectorizer: End-to-end vectorization with deep reinforcement learning,” in *Proceedings of the 18th ACM/IEEE International Symposium on Code Generation and Optimization*, 2020, pp. 242–255.
- [21] C. Cummins, H. Leather, B. Steiner, H. He, and S. Chintala, “CompilerGym: A reinforcement learning toolkit for compilers,” <https://github.com/facebookresearch/CompilerGym/>, 2020.
- [22] R. Mammadli, A. Jannesari, and F. Wolf, “Static neural compiler optimization via deep reinforcement learning,” in *2020 IEEE/ACM 6th Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC) and Workshop on Hierarchical Parallelism for Exascale Computing (HiPar)*. IEEE, 2020, pp. 1–11.
- [23] M. Trofin, Y. Qian, E. Brevdo, Z. Lin, K. Choromanski, and D. Li, “Mlgo: a machine learning guided compiler optimizations framework,” *arXiv preprint arXiv:2101.04808*, 2021.
- [24] L.-N. Pouchet, C. Bastoul, A. Cohen, and N. Vasilache, “Iterative optimization in the polyhedral model: Part i, one-dimensional time,” in *International Symposium on Code Generation and Optimization (CGO’07)*. IEEE, 2007, pp. 144–156.
- [25] A. Schrijver, *Theory of linear and integer programming*. John Wiley & Sons, 1998.
- [26] H. Le Verge, “A Note on Chernikova’s algorithm,” INRIA, Research Report RR-1662, 1992. [Online]. Available: <https://hal.inria.fr/inria-00074895>
- [27] S. Ganser, A. Größlinger, N. Siegmund, S. Apel, and C. Lengauer, “Speeding up iterative polyhedral schedule optimization with surrogate performance models,” *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 15, no. 4, pp. 1–27, 2018.
- [28] J. Ragan-Kelley, C. Barnes, A. Adams, S. Paris, F. Durand, and S. Amarasinghe, “Halide: a language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines,” *Acm Sigplan Notices*, vol. 48, no. 6, pp. 519–530, 2013.
- [29] A. Adams, K. Ma, L. Anderson, R. Baghdadi, T.-M. Li, M. Gharbi, B. Steiner, S. Johnson, K. Fatahalian, F. Durand *et al.*, “Learning to optimize halide with tree search and random programs,” *ACM Transactions on Graphics (TOG)*, vol. 38, no. 4, pp. 1–12, 2019.
- [30] T. Chen, T. Moreau, Z. Jiang, L. Zheng, E. Yan, H. Shen, M. Cowan, L. Wang, Y. Hu, L. Ceze *et al.*, “{TVM}: An automated end-to-end optimizing compiler for deep learning,” in *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, 2018, pp. 578–594.
- [31] R. Baghdadi, M. Merouani, M.-H. Leghettas, K. Abdous, T. Arbaoui, K. Benatchba *et al.*, “A deep learning based cost model for automatic code optimization,” *Proceedings of Machine Learning and Systems*, vol. 3, 2021.
- [32] C. Cummins, P. Petoumenos, Z. Wang, and H. Leather, “End-to-end deep learning of optimization heuristics,” in *2017 26th International Conference on Parallel Architectures and Compilation Techniques (PACT)*. IEEE, 2017, pp. 219–232.
- [33] A. Brauckmann, A. Goens, S. Ertel, and J. Castrillon, “Compiler-based graph representations for deep learning models of code,” in *Proceedings of the 29th ACM SIGPLAN International Conference on Compiler Construction (CC 2020)*, ser. CC 2020. New York, NY, USA: Association for Computing Machinery, Feb. 2020, p. 201–211. [Online]. Available: <https://doi.org/10.1145/3377555.3377894>
- [34] G. Ye, Z. Tang, H. Wang, D. Fang, J. Fang, S. Huang, and Z. Wang, “Deep program structure modeling through multi-relational graph-based learning,” in *Proceedings of the ACM International Conference on Parallel Architectures and Compilation Techniques*, 2020, pp. 111–123.
- [35] C. Cummins, Z. V. Fisches, T. Ben-Nun, T. Hoefler, and H. Leather, “Programl: Graph-based deep learning for program optimization and analysis,” *arXiv preprint arXiv:2003.10536*, 2020.
- [36] S. Mehta, P.-H. Lin, and P.-C. Yew, “Revisiting loop fusion in the polyhedral framework,” in *Proceedings of the 19th ACM SIGPLAN symposium on Principles and practice of parallel programming*, 2014, pp. 233–246.
- [37] S. Verdoolaege and G. Janssens, “Scheduling for ppcg,” *Report CW*, vol. 706, 2017.
- [38] S. Mehta and P.-C. Yew, “Improving compiler scalability: Optimizing large programs at small price,” *ACM SIGPLAN Notices*, vol. 50, no. 6, pp. 143–152, 2015.
- [39] R. Baghdadi and A. Cohen, “Scalable polyhedral compilation, syntax vs. semantics: 1–0 in the first round,” in *IMPACT 2020 workshop (associated with HIPEAC 2020)*, 2020, informal proceedings.
- [40] G. Brockman, V. Cheung, L. Pettey, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “Openai gym,” *arXiv preprint arXiv:1606.01540*, 2016.