RL4ReAl: Reinforcement Learning for Register Allocation

S. VenkataKeerthy[1], **Siddharth Jain**[1], **Rohit Aggarwal**[1], **Albert Cohen**[2] and **Ramakrishna Upadrasta**[1]

[1]Indian Institute of Technology Hyderabad

[2]Google

{cs17m20p100001, cs20mtech12003, cs18mtech11030}@iith.ac.in, albertcohen@google.com, ramakrishna@cse.iith.ac.in

Abstract

We propose a novel solution for the Register Allocation problem, leveraging multi-agent hierarchical Reinforcement Learning. We formalize the constraints that precisely define the problem for a given instruction-set architecture, while ensuring that the generated code preserves semantic correctness. We also develop a gRPC based framework providing a modular and efficient compiler interface for training and inference. Experimental results match or outperform the LLVM register allocators, targeting Intel x86 and ARM AArch64.

1 Introduction

Register allocation is one of the well-studied and important compiler optimization problems. It involves assigning a finite set of registers to an unbounded set of variables. Its decision problem is reducible to graph coloring, which is one of the classical NP-Complete problems [Garey and Johnson, 1979; Bouchez *et al.*, 2006]. Register allocation as an optimization involves additional sub-tasks, upstream and downstream from graph coloring itself. Several formulations have been proposed that return exact, or heuristic-based solutions.

Broadly, the optimal solutions are often formulated as constraint-based optimizations [Lozano *et al.*, 2012; Kuchcinski, 2003], ILP formulations [Appel and George, 2001; Barik *et al.*, 2006; Chang *et al.*, 1997; Nagarakatte and Govindarajan, 2007], PBQP formulations [Kim *et al.*, 2021], and are fed to a variety of solvers. These approaches are known to have scalability issues.

On the other hand, several heuristic-based approaches also exist. They have been widely used in modern compilers owing to their scalability: good solutions for practical benchmarks in near linear time. However, developing good heuristics is highly non-trivial and requires very specialized domain expertise, on compiler construction as well as on hardware architecture. Various heuristics have been proposed over the past 40 years [Chaitin *et al.*, 1981; Chow and Hennessy, 1990; Briggs *et al.*, 1994], extending to recent times [Chen *et*

al., 2018]. They are often fine-tuned for a particular architecture and give non-optimal performance.

Recently, with the success of Machine Learning (ML) in several domains, ML-based approaches are being proposed to solve compiler optimization problems that have been known to be computationally expensive. Analogous to that of natural language representations like word2vec [Mikolov *et al.*, 2013], and GloVe [Pennington *et al.*, 2014], several approaches have been proposed to represent programs as information-rich embeddings: inst2vec [Ben-Nun *et al.*, 2018], IR2Vec [VenkataKeerthy *et al.*, 2020], Flow2Vec [Sui *et al.*, 2020]. These approaches have been applied to classical optimizations like phase ordering of optimization passes [Fursin and others, 2011; Ashouri *et al.*, 2017], function inlining decisions [Simon *et al.*, 2013], throughput prediction [Mendis *et al.*, 2019a], etc., *where correctness is not an issue.* However, they remain to be applied to compiler optimizations under hard semantic constraints, such as register allocation. We believe this did not happen yet because of some of the following reasons.

- Register allocation is a complex problem, composed of multiple sub-tasks, including splitting, coalescing, spilling. These sub-tasks all have to considered in addition to modeling hardware complexities.
- ML-based allocation schemes should ensure correctness; no two variables in the same live range be assigned to the same register, and the register types should be respected. Such semantic constraints should not suffer any approximation, unlike forgetful optimizations like function inlining decisions.
- On the practical side, it is hard to integrate ML solutions with modern compiler frameworks where the ML solutions are often written in Python and compilers are written in C++ and are among the most complex software engineering architectures. It is a challenging engineering problem.

In this work, we propose a retargetable Reinforcement Learning (RL) register allocator addressing the above mentioned challenges. Our approach finds a middle ground between the heuristic-based and exact solutions in terms of scalability and performance. We formulate a *multi-agent hierarchical reinforcement learning* approach (1) to model the sub-tasks of register allocation

like coloring, live range splitting and spilling, and (2) to encode the correctness constraints for semantically preserving and hardware-compatible register assignments. The legality of the register allocations and assignments is preserved by imposing constraints on the action space, or outcome of each agent. Formulating register allocation as an RL approach is more natural, as it results in a complex combinatorial problem for which establishing the ground truth is hard. Also, imposing correctness constraints would otherwise be not possible.

We leverage the LLVM infrastructure [Lattner and Adve, 2004]. The interference graph of a function is extracted from the Machine Intermediate Representation (MIR) of LLVM, representing the instructions within each node as IR2Vec vectors [VenkataKeerthy *et al.*, 2020]. For this, we extend IR2Vec to generate embeddings for MIR entities. These embeddings represent vertices of the interference graph which is traversed by RL agents. Finally, we propose a generic gRPC-based framework (https://grpc.io) to interface the RL model with the compiler.

Earlier works [Huang *et al.*, 2019a; Lemos *et al.*, 2019; Musliu and Schwengerer, 2013] addressed the conventional graph coloring problem using ML. However, these approaches may not result in a practical register allocation scheme: they only solve one sub-task of the overall problem, and do not exploit semantic information such as live ranges to arrive at the final coloring and spilling decisions.

Das et al. [2020] investigated the usage of LSTM for register allocation. They use ML to construct an initial graph coloring, which then undergoes a correction phase to rectify inconsistencies in coloring interfering nodes. Their work focuses on graph coloring alone, and to our understanding it was not integrated in a compiler to produce actual spilling decisions and register assignments.

Contributions
- We propose the *first end-to-end application* of RL for solving the register allocation problem.
- Formalizing the constraints to restrict the action space and preserve semantic correctness.
- MIR2VEC encodings to generate representations at Machine IR (MIR) level.
- Experimental evaluation on two hardware architectures (x86-64, AArch64) and reference benchmarks.

2 Background and Mathematical Model

In this section, we mathematically formulate the register allocation problem by defining the constraints that arise out of hardware and the program. We also give an overview of the available register allocators in LLVM.

2.1 Register constraints

Optimizing compilers convert the source code into an Intermediate Representation (IR), where target-independent optimizations take place. In the backend compiler, this IR is incrementally lowered to a machine-specific one. In LLVM this representation is called Ma-

chine IR (MIR). The MIR at the stage of register allocation is very close to machine instructions, as instruction selection and other low-level optimizations have already been performed. After instruction selection, certain physical registers that are mandated by the architecture are immediately assigned.

For instance, x86 processors mandate the output of 32-bit division to be stored only in $eax and $edx registers.

As it can be seen from Fig. 1, the IDIV32 instruction divides the contents of $eax and $edx by %x and stores the result in $eax. Such mandatory assignments–including calling conventions—are assigned. The register allocation problem can now be reduced to assigning physical registers to the other *left-out* virtual registers ($\mathcal{V}$) while respecting the following constraints.

Type constraint The register file ($\mathbf{R}$) of a machine consists of collection of registers R^t belonging to different *types* (t): $\mathbf{R} = \bigcup_t R^t$. Assigning a physical register r to a virtual register v of type t, $v^t \blacktriangleright r$, should satisfy the register *type* constraint $\chi^{\mathcal{T}}(v^t) = \{v^t \blacktriangleright r : r \in R^t\}$. In Fig. 1(b), each virtual register is associated with a particular register type. For instance, %x is of gr32 type, which means that it belongs to a 32-bit wide general-purpose register type. Consequently, only registers belonging to that type (like $eax) can be assigned.

Congruence constraint Real-world instruction set architectures like x86 and AArch64 have a hierarchy of register *classes*. For instance, 32 bit type of registers (like $eax, $ebx) are physically *part of* the 64 bit ones (like $rax, $rbx). We consider the registers that adhere to this *part of* relation as a congruent class $\mathcal{C}(r)$. For example, registers $al, $ah, $ax, $eax, $rax of x86, which are "chunks" of the *same physical register* belong to the same congruent class, satisfying $al, $ah $\sqsubseteq$ $ax $\sqsubseteq$ $eax $\sqsubseteq$ $rax. So, the register assignments for virtual register v should be among the set of registers that satisfy the following *congruence* constraint:

$$\chi^{\mathcal{C}}(v^t) = \{v_i^t \blacktriangleright r : \forall v_i, v_j \in \mathcal{V}, v_i \neq v_j, \nexists v_j \blacktriangleright \mathcal{C}(r) \in L(v_i)\}$$

Here $L(v)$ corresponds to the live range of variable v, and is computed as $L(v) = [P_v^{def}, P_v^{end}]$; the definition of v occurs at program point P^{def}, and its last use is in P^{end}. Fig. 1(a) gives an example. The live ranges of the corresponding variables are shown in Fig. 2(b).

Interference constraint Register allocation has been modeled as a graph coloring problem [Chaitin *et al.*, 1981]. For each function in the program, it involves creating an *Interference graph* $G(V, E)$ defined as follows: the vertices of the graph are mapped to virtual registers (v) or physical registers (R_a), meaning $V \in (\mathcal{V} \cup R_a)$; the edges E are computed as $\{(v_i, v_j) : v_i, v_j \in V \wedge L(v_i) \cap L(v_j)\}$. The interference graph corresponding to the example in Fig. 1 is shown in Fig. 2(b). The *interference* constraint says that no two adjacent nodes in G should be allocated the same color. The set of registers that satisfy this constraint is given by:

$$\chi^{\mathcal{I}}(v^t) = \{R^t \setminus \{r : \forall u((u, v) \in E \wedge u \blacktriangleright r)\}\}$$

In summary, for a given virtual register v of type t, the set of available registers for allocation $\chi(v^t)$ is defined as the set of registers that satisfy the (i) type, (ii) congruence, and (iii) interference constraints:

$$\chi(v^t) = \chi^{\mathcal{T}}(v^t) \cap \chi^{\mathcal{C}}(v^t) \cap \chi^{\mathcal{I}}(v^t)$$

2.2 Live range splitting and spilling

The above formulation of register allocation in terms of graph coloring is well known and natural, as a *decision problem*. Yet register allocation as an *optimization problem* is actually *much more* than graph coloring. For instance, when there are not enough physical registers available, deciding which variable (virtual register) has to be spilled to memory is important, as memory accesses take far more time than register accesses. Spilling a variable, $\mu(v)$ involves writing/reading it to/from a memory location on access. A trivial example is the loop induction variable: it would incur high cost to read/write from/to memory, if a decision is made to spill it. Hence, register allocators try to reduce the spill cost $M(v)$, in addition to minimizing the number of spills. For a machine with 3 registers, the example code shown in Fig. 1 is not *3-colorable*, and results in spilling a variable.

A live range of a virtual register can be split. Let $k \in K$ denote a program point among the uses of v. Splitting live range of v at k is defined as $\varphi(v, k) : L(v) \rightsquigarrow (L(v'), L(v''))$; $L(v') = [P_v^{def}, P_v^k]$, $L(v'') = [P_v^{k+1}, P_v^{end}]$. Determining which variable to split, and at which point is non trivial. In Fig. 2(b), splitting i at 5 into (i', i'') makes the interference graph *3-colorable*.

2.3 Register allocators in LLVM

Register allocators are implemented in compilers as *passes*. LLVM currently has four different register allocators: fast, basic, greedy, and PBQP, ranked according to the complexity of implementation. They are strictly intraprocedural, operating one function at a time.

The fast register allocator which operates at the basic block level is an improved version of the linear scan algorithm [Poletto and Sarkar, 1999]. The basic register allocator is an improved variation of the fast register allocator and operates at the function level [Xavier *et al.*, 2012]. The greedy allocator was developed to address the shortcomings of the basic allocator [Jakob, 2011]; it combines four strategies iteratively: splitting, spilling, coalescing (merging of live ranges), eviction (de-allocating the already-allocated physical register). Each of these strategies is driven by greedy heuristics. This allocator iterates over the virtual registers and obtains a legal physical register allocation if possible. In this process it applies node-splitting and eviction. The PBQP register allocator is the only solver based mechanism in LLVM and models it as a Partitioned Boolean Quadratic Problem to obtain allocations [Hames and Scholz, 2006].

Not all allocators implement all strategies; live range splitting only takes place in greedy, whereas coalescing is present in greedy and PBQP, and eviction in iterative allocators like greedy and not PBQP. Our framework is non-iterative, and we currently model splitting, spilling and coloring; we leave coalescing for future work.

3 Representing Interference Graphs

We represent nodes of the interference graph as embeddings obtained from LLVM's MIR instructions. This forms the input to a Gated Graph Neural Network (GGNN) that learns to generate the representation of the state space. Here, we describe this process.

The IR2Vec framework is based on the generic IR of LLVM. We use the same approach on MIR to generate MIR2Vec representations. This involves generating triplets by forming relations between entities, training TransE [Bordes *et al.*, 2013] to obtain the seed embedding vocabulary, and using it to create instruction-level representations. As MIR is target-specific, the generated embeddings are also specific to the architecture.

MIR Entities Opcodes and MIR instruction arguments form the entities. The arguments to MIR instructions primarily include physical and virtual registers, and immediate values. Like IR2Vec, we abstract out these arguments with generic identifiers.

We create 2 different relations. (i) *NextInst*: Captures the relation between the current opcode and the next instruction opcode, (ii) Arg_i: Captures the relation between the opcode and the arguments of the instruction. Once the triplets are generated, we train the TransE model to obtain the embeddings for each of the entities.

Grouping of opcodes Unlike LLVM IR, MIR contains more specialized opcodes, in terms of the operating width, among other factors. LLVM exposes about 15.3K of different possible opcodes in X86 and about 5.4K in Aarch64[1]. In comparison, LLVM IR only has about 64 entities (used by IR2Vec) in total.

Obtaining a dataset to cover all such specialized operands would be highly infeasible, and in turn, would not generate good representations. Hence we mask out the opcodes based on their operating width, the source and destination locations (immediate, register, and memory) and group them together. For example, there are about 200 different `MOV` instructions operating on different bit width, sources, and destinations, like `MOV32r0`, `MOVZX64rr16`, `MOVAPDrr`, etc. All such opcodes are grouped together as a generic `MOV` token while forming the triplets. The generated triplets are used as an input to the TransE model to generate the embeddings for each entity. This forms the seed embedding vocabulary for our problem.

Representing instructions For a given MIR instruction with opcode O and n arguments $A_1, A_2, \ldots, A_n$, its representation is computed as

$$W_o \cdot [\![\mathbf{O}]\!] + W_a \cdot ([\![\mathbf{A_1}]\!] + [\![\mathbf{A_2}]\!] + \cdots + [\![\mathbf{A_n}]\!]), W_o > W_a$$

---

[1]From the inc files - {build_dir}/lib/Target/X86/X86GenInstrInfo.inc and {build_dir}/lib/Target/AArch64/AArch64GenInstrInfo.inc
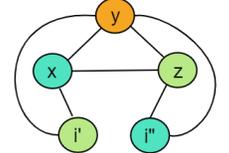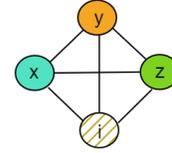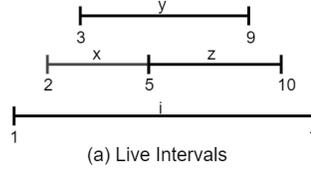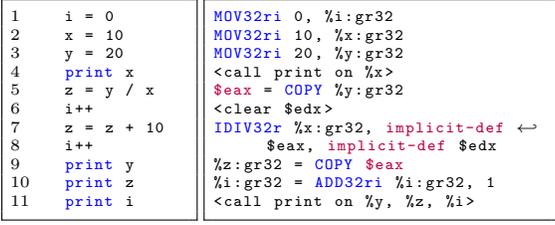
Figure 1: (a) Example source code and (b) its Machine IR

Figure 2: Register allocation with and without splitting
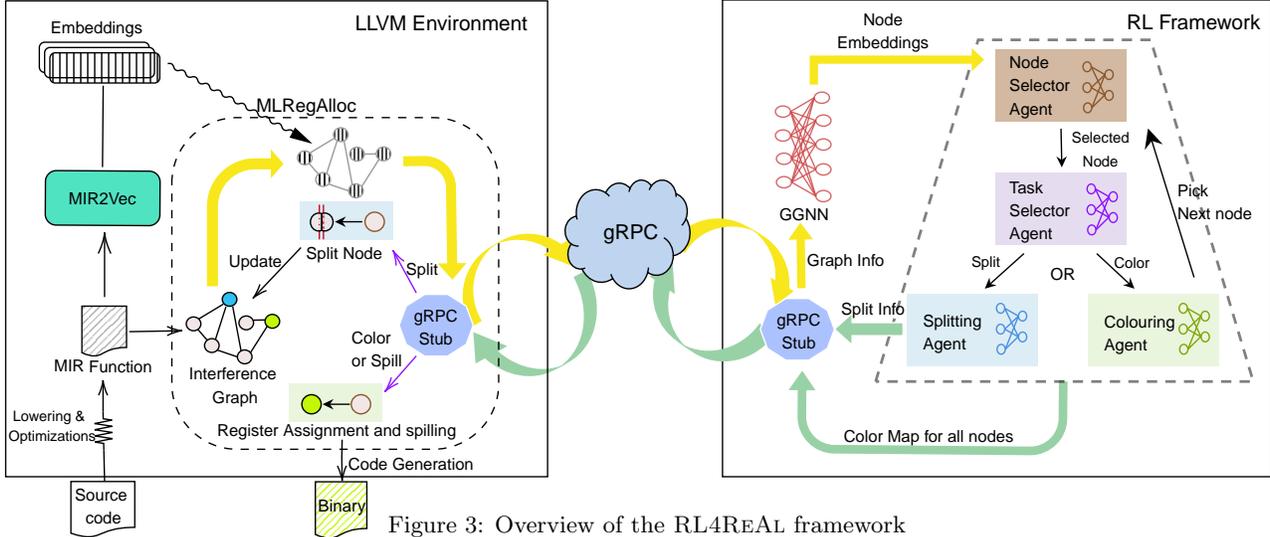
Figure 3: Overview of the RL4ReAl framework

where $[\![\cdot]\!]$ denotes the embedding of the entity from seed-embedding vocabulary, as proposed by the symbolic encodings of IR2Vec.

3.1 Interference Graphs

We use the information collected from MIR to obtain the interference graphs in LLVM. As mentioned earlier, the MIR at the stage of register allocation contains partial physical register assignments and virtual registers. The physical registers are assigned for the instruction operands that have restrictions on the particular register to be used. Virtual registers are used in all other places.

Consequently, we need to take into account the edges corresponding to both virtual and physical registers. Virtual registers are marked with the register class as explained earlier. Hence the assignments can only be one among the physical registers in that class.

For computing the interference graphs, considering the interferences between the (physical register, virtual register) and (virtual register, virtual register) would be sufficient as the graph is bidirectional and that we do not need to worry about the physical registers that are already assigned. Hence the interference graph G is computed as shown in Algorithm 1.

We use a collection of instructions in the live-range of a variable to represent a vertex of the interference graph. Each instruction is *represented* in $\mathbb{R}^n$ using MIR2Vec

---

Algorithm 1: Creating Interference Graph

Parameter: MachineFunction F, Graph &G
for *each physical register $p_i \in F$* **do**
  for *each virtual register $v_i \in F$* **do**
    if *Interference(p_i, v_i)* **then**
      G.addEdge(p_i, v_i)
for *each virtual register $v_i \in F$* **do**
  for *each virtual register $v_j \in F$ and $v_i \neq v_j$*
  do
    if *Interference(v_i, v_j)* **then**
      G.addEdge(v_i, v_j)

---

embeddings. Consequently, a vertex v is represented as a matrix of embeddings $[\![v]\!]$ in $\mathbb{R}^{m \times n}$, where m denotes the number of instructions in its live range.

In the recent times, Gated Graph Neural Networks (GGNNs) have found wide applications in programming language problems that are modeled as graphs [Cummins and others, 2021b; Mendis *et al.*, 2019b]. GGNNs involve message passing between the nodes of the graph where the information is propagated across the nodes multiple times to arrive at the representation for a node. Also, they allow annotating the nodes and edges based on their types and properties, and consider them while learning the representations. We use GGNNs to process the embedded interference graph to get the final representation. This network transforms $\mathbb{R}^{m \times n} \to \mathbb{R}^k$, where

k is a hyperparameter. We set $k = n$ in our experiments, while considering the following node types.

- **Not visited** — Nodes that are not visited yet.
- **Spill** — Nodes that are marked as spill.
- **Colored** — Nodes that are assigned a register.

Nodes in the graph are marked with these annotations, along with the spill cost. Such node representations are propagated through a GGNN by means of message passing. Messages received from adjacent nodes are aggregated and passed through a Gated Recurrent Unit [Cho *et al.*, 2014] to yield a final representation.

4 Hierarchical Reinforcement Learning

We formulate register allocation as a Markov Decision Process (MDP) using hierarchical Reinforcement Learning (RL). We model the sub-tasks of register allocation as lower level tasks controlled by multiple agents. Fig. 3 provides an overview of our approach. It involves interactions between the LLVM compiler and the RL model for both training and inference.

4.1 Environment

We implemented a new MLRegAlloc pass in LLVM, to generate an interference graph (G), allocate, split and spill registers as predicted by the agents. This pass also generates a representation of G using MIR2VEC.

4.2 Agents

The task of allocating registers is split into multiple sub-tasks across the horizon. Each of these tasks are modeled as agents $\{\omega_v, \omega_\tau, \omega_\varphi, \omega_\xi\} \in \Omega$, that learn their respective policies π_ω to optimally solve the low level tasks. We formulate hierarchical agents for four sub-tasks as shown in Fig. 3:

- Node selector (ω_v): Top level agent that learns to pick a node $v \in G$.
- Task selector (ω_τ): Mid level agent that learns to select a task among $\{\chi, \varphi\}$ on v picked by ω_v.
- Splitter (ω_φ): Low level agent that learns to identify a split point k for v.
- Coloring Agent (ω_ξ): Low level agent that learns to pick a valid $\chi_i \in \chi$ or μ.

As it can be seen, each high level agent invokes a low level agent while following the timeline: $\omega_v \prec \omega_\tau \prec \{\omega_\varphi, \omega_\xi\}$. Each agent ω has its own state space S_ω, action space A_ω, and reward R_ω to learn a policy π_ω.

Coloring Agent (ω_ξ) For a graph of V nodes, given a set of registers $\chi(v)$ available at the instant, the state space of ω_ξ is given as a tuple ($[\![v]\!], |\chi(v)|, |V_{\text{nclr}}|$), where $V_{\text{nclr}} = V \setminus V_{\text{clrd}}$ are the nodes to be colored, v is the node that is picked by ω_v, and $[\![.]\!]$ denotes its embedding. Meaning, the coloring agent uses the following information to decide the register to be assigned: the embedding of the vertex v, the number of registers that satisfy the constraints (defined in Sec. 2.1), and the number of uncolored nodes in the interference graph. If no registers are available, the coloring agent spills v. Hence the legal action space of ω_ξ is:

$$A(\omega_\xi) = \begin{cases} \chi(v), & |\chi(v)| > 0 \\ \mu(v), & otherwise \end{cases}$$

$\chi(v)$ gives the set of legal registers for v (Sec. 2.1). To improve performance, the agent should maximize the use of registers for vertices with higher spill cost. And, spill cost roughly corresponds to the importance of the node v. Hence the reward for the coloring agent is given as:

$$R(\omega_\xi) = \begin{cases} +M(v), & \text{if } \chi(v) \\ -M(v), & \text{if } \mu(v) \end{cases}$$

Splitter (ω_φ) For predicting where to split the live range of a variable v, the node splitter considers the spill weights at each use of the variable $\mathcal{M}(v) = \{M(v_k) : \forall k \in K\}$, the distances between each successive use $D_v = \{D(v_i, v_{i+1}), \forall i \in K\}$, and the embedding of $[\![v]\!]$. The use distance is the number of program points between two uses of v. Hence, the state space is given as a tuple ($[\![v]\!], \mathcal{M}(v), D_v$). For a given state, the agent learns an optimal program point $p \in K$ where v can be split. Hence the action space $A(\omega_\varphi) = K$. The reward for the agent on splitting v into (v', v'') is dependent on two factors: (i) the variation of use distance between the chosen split point and the successor use, and (ii) the variation in the number of interferences ($\delta(v)$), the degree of v before and after the split. If the variation of use distances is higher than that of the interferences, the split is beneficial. The corresponding reward is:

$$R(\omega_\varphi) = D(v', v'') + \delta(v) - (\delta(v') + \delta(v''))$$

Task Selector (ω_τ) For selecting a task (τ) among coloring and splitting, the agent ω_τ considers the parameters specific to each of the tasks: the representation of v, the number of available registers, the number of interferences, its life-time, spill cost. Hence the state space is formulated as the tuple: ($[\![v]\!], |\chi(v)|, \delta(v), |K(v)|, M(v)$). The action space of ω_τ is defined as:

$$A(\omega_\tau) = \begin{cases} \{\varphi, \chi\}, & |K(v)| \geq k \\ \chi, & otherwise \end{cases}$$

Here $|K(v)| \geq k$ indicates that v should have at least k uses to be considered for splitting. We define k as a hyper-parameter. We set $k = 2$ (1 definition and 1 use) in our experiments. We model the reward for this agent based on the outcome of the low level tasks.

$$R(\omega_\tau) = \begin{cases} R(\omega_\xi), & \tau = \chi \\ R(\omega_\varphi) - (1.001 \times \#Splits)/10, & \tau = \varphi \end{cases}$$

Clearly, $R(\omega_\varphi)$ is always positive hence ω_τ would always favor splitting. But too many splits are also not desirable because they may induce extraneous move instructions in the generated assembly code. So, while computing $R(\omega_\tau)$, a discount factor is applied that decreases the reward with increase in number of splits.

Figure 4: Reward Vs. step

Arch.	Considered registers
x86	[A-D]L, [A-D]X, [E,R][A-D]X, [SI,DI]L, SI, DI, [E,R][SI,DI], R[8-15][B,W,D]
Aarch64	X[0-30], W[0-30]

Table 1: Registers considered
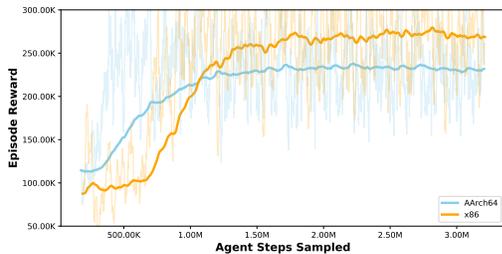
Node Selector (ω_v) The state space of ω_v comprises the embedding of each vertex in G obtained from a GGNN. Along with these embeddings, the agent uses the spill weights of the nodes M to characterize the state. Hence, the state space is given as a tuple $([\![G]\!], M(V))$ Its legal action space is $A(\omega_v) = V_{\text{nclr}}$. The learned policy is deemed good based on the final outcome of the node. Hence the reward for this agent is also modeled based on the rewards of the low-level agents.

$$R(\omega_v) = \begin{cases} R(\omega_\xi), & \tau = \chi \\ R(\omega_\varphi), & \tau = \varphi \end{cases}$$

5 Compiler Integration

Developing frameworks to ease the integration of compilers with RL environments is important to make the solution practically viable. CompilerGym [2021a] makes it possible to leverage Python libraries for solving compiler optimization problems; it exposes RL environments and datasets by training RL models.

In this work, we developed a framework to support *both training and inference*, where the deployment of the model is transparent to the end user. It is particularly critical to demonstrate the practicality of our approach to representative benchmarks. The design of RL4ReAl involves *to-and-fro communication* between the compiler and Python model (Fig. 3). To our knowledge, such a facility is *not available* in other frameworks; it allows to integrate ML models easily with the compiler, and give RL4ReAl the power to operate on *any* module—basic-block, loop, function—of the input program.

For example, a splitting decision by the model is communicated to the compiler, which then applies it and responds back with the update containing new interferences and live ranges. The model then updates the interference graph using the received information and continues the traversal. After processing all vertices of G, all the coloring decisions are communicated to the compiler as a color map. The backend compiler uses the allocation decisions made by the model to generate code.

6 Experimental Evaluation

Methodology   For training MIR2Vec representations, we randomly select 2000 source code files from SPEC CPU 2017 benchmarks and C++ Boost library. MIR triplets are generated by applying `-O3` optimizations. The seed embedding vocabulary is obtained by

training a TransE model on the generated triplets, with the same hyperparameters as IR2Vec, running an SGD optimizer over 1000 epochs to obtain an embedding vector of 100 dimensions. We evaluate performance on a complex x86 microarchitecture (Intel Xeon W2265, 16 cores, 64GB RAM) and a simpler mobile AArch64 processor (ARM Cortex A72, 2 cores, 4GB RAM).

We obtain 1 billion MIR triplets from which $\{675, 315\}$ entities and $\{25, 17\}$ relations are generated for $\{$x86, AArch64$\}$ respectively. For training the RL model, we generate the interference graphs corresponding to the selected functions using our MLRegAlloc pass. We randomly choose ≈ 2000 functions having a number of variables between 70–120. Functions with less than 70 variables are skipped as the number of registers is large enough to color the vast majority of them without splitting or spilling. We also skipped those with more than 120 variables as they only form $\approx 1\%$ of the dataset. For the purpose of experimentation, we consider allocations of general purpose registers for both x86 and AArch64, as these registers contribute to major portion of the code. The list of registers is detailed in Tab. 1. Training took place on a P100 GPU and sampling was done using 12 threads of an Intel Xeon server.

Our framework controls the set of registers to be considered for allocation through a configuration file. For the virtual registers that belong to other register classes that are not specified in this file, allocation is done by the existing greedy register allocator of LLVM. We disable allocations to vector registers while generating the code for the purpose of comparison. Our framework is implemented as a compiler pass—MLRegAlloc in LLVM 10.0.1. We integrated gRPC v1.34, and protobuff v3.13, and designed a client-server model for a seamless communication between LLVM and the Python model. We train the RL models using the PPO policy with the standard set of hyperparameters on the training set of functions, until convergence of reward graph. The reward curve is shown in the Fig. 4.

Results   We study the performance on the standard benchmarks from SPEC CPU 2017 and 2006 in terms of runtimes and number of reloads (memory access) per spill, and compare the results with the standard register allocators of LLVM. By design, basic, greedy and PBQP allocators generally outperform better the fast allocator. However, it is not possible to identify a single allocator among these that perform the best for all programs. Hence, we compare our results with these three allocators. The runtimes of the code generated by using our register allocator RL4ReAl and other allocators are shown in Fig. 5(I). These runtimes are obtained
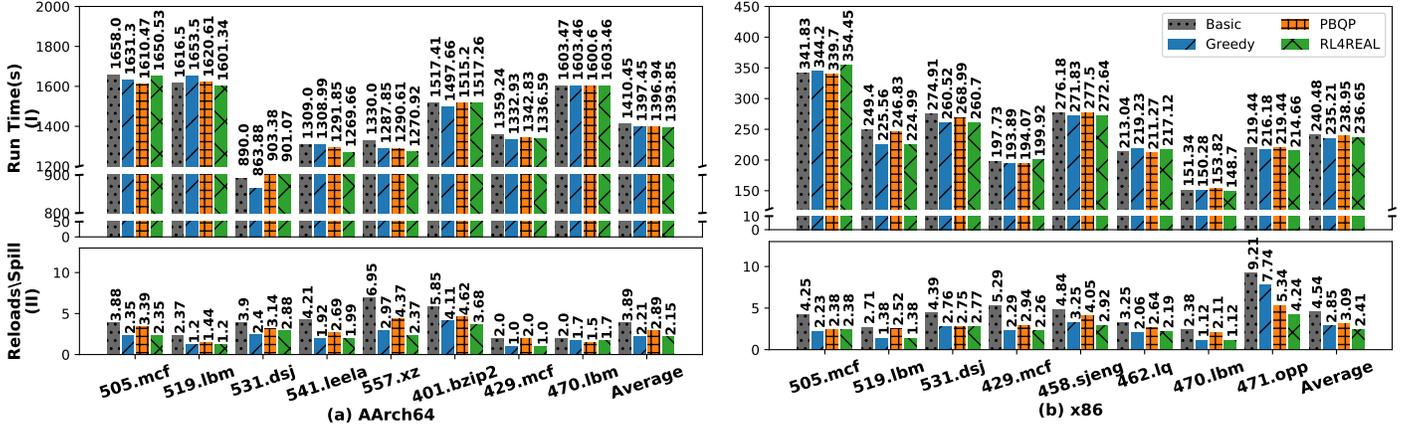
Figure 5: Comparison of runtimes and #reloads/spill obtained using various register allocators

Benchmark	Basic	Greedy	PBQP	RL4REAL
505.mcf_r	1658.00	1631.30	**1610.47**	1650.53
519.lbm_r	1616.50	1653.50	1620.61	**1601.34**
531.deepsjeng_r	890.00	**863.88**	903.38	901.07
541.leela_r	1309.00	1308.99	1291.85	**1269.66**
557.xz_r	1330.00	1287.85	1290.61	**1270.92**
401.bzip2	1517.41	**1497.66**	1515.20	1517.26
429.mcf	1359.24	**1332.93**	1342.83	1336.59
470.lbm	1603.47	1603.46	**1600.60**	1603.46
Average	1410.45	1397.45	1396.94	**1393.85**

Table 2: Runtimes (in s) with different register allocators on AArch64

Benchmarks	Basic	Greedy	PBQP	RL4REAL
505.mcf_r	341.83	344.20	**339.70**	354.45
519.lbm_r	249.40	225.56	246.83	**224.99**
531.deepsjeng_r	274.91	**260.52**	268.99	**260.70**
429.mcf	197.73	**193.89**	194.07	199.92
458.sjeng	276.18	**271.83**	277.50	272.64
462.libquantum	213.04	219.23	**211.27**	217.12
470.lbm	151.34	150.28	153.82	**148.70**
471.omnetpp	219.44	216.18	219.44	**214.66**
Average	240.48	**235.21**	238.95	236.65

Table 3: Runtimes (in s) with different register allocators on x86

by taking the median of three executions. On average, the runtimes obtained by RL4REAL are *comparable* or perform *better than the best* allocator of LLVM.

Among the standard allocators, greedy is best on x86 and results in 2.19% and 1.57% improvement over basic and PBQP allocators. On AArch64, PBQP is best and results an improvement of 0.96% and 0.04% over basic and greedy allocators. RL4REAL itself improves the runtime over basic and PBQP by 1.59% and 0.96% respectively, and nearly matches greedy at -0.61% On the other hand, in AArch64, RL4REAL results in an improvement of 1.18%, 0.22% and 0.26% on the runtimes obtained by basic, PBQP, and greedy respectively. *It can be noted that these strong results have been obtained fully automatically, against production-grade allocators tuned over many man-decades of experience and effort.* The obtained runtimes corresponding to AArch64 and x86 architectures are shown in Tab. 2 and Tab. 3 respectively.

Another metric which can be used to evaluate the register allocator is based on effective spilling strategy. It is important to note that the absolute number of spills is not a sufficient metric for comparison. Indeed, allocators that perform live range splitting may create multiple live ranges that could result in multiple spilling instances, yet delivering higher overall performance. Hence, to study spilling, we use the number of times a spilled virtual reg-

ister is reloaded from memory as a metric. Any register allocator should try to minimize the number of memory accesses (reloads) by spilling a virtual register which is used infrequently. The number of reloads performed per a spilled virtual register is shown in Fig. 5(II). As it can be observed, RL4REAL results in less reloads per spill among all the allocators in both the architectures. This shows the RL agents have learned a good policy that chooses an appropriate vertex in the interference graph to spill. We believe that the spilling improvments result from a better coordination among all the agents in our model; this helps picking an appropriate node, splitting its live range to ease coloring, and efficient spilling that minimize the number of reloads. The number of spills, reloads and reloads per spill ratio for AArch64 and x86 architectures are shown in Tab. 4 and Tab. 5 respectively.

It is also important to note that minimizing the number of reloads per spill alone may not necessarily lead to improved runtime. Several other factors, like other compiler optimizations post register allocations, link time optimizations, and the complexity of underlying micro-architecture etc. will play a major role in the performance of the final code that is generated.

7 Related Work

ML for compiler optimizations There is a growing interest in applying machine learning techniques

Benchmarks	Basic			Greedy			PBQP			RL4REAL		
	S	R	R/S	S	R	R/S	S	R	R/S	S	R	R/S
505.mcf_r	24	93	3.88	23	54	2.35	23	78	3.39	23	54	2.35
519.lbm_r	30	71	2.37	35	42	1.2	32	46	1.44	35	42	1.2
531.dsj_r	188	734	3.9	196	471	2.4	198	621	3.14	215	619	2.88
541.leela_r	144	606	4.21	137	263	1.92	132	355	2.69	140	279	1.99
557.xz_r	192	1334	6.95	187	555	2.97	187	818	4.37	171	405	2.37
401.bzip2	71	415	5.85	85	349	4.11	76	351	4.62	102	375	3.68
429.mcf	3	6	2	1	1	1	1	2	2	9	9	1
470.lbm	3	6	2	10	17	1.7	2	3	1.5	10	17	1.7
Average	81.87	408.12	3.89	84.25	219	2.21	81.37	284.25	2.89	88.12	225	2.15

Table 4: Number of spills (S) and reloads (R) and reloads per spill ratio (R/S) induced by different allocators in AArch64

Benchmarks	Basic			Greedy			PBQP			RL4REAL		
	S	R	R/S	S	R	R/S	S	R	R/S	S	R	R/S
505.mcf_r	53	225	4.25	56	125	2.23	56	133	2.38	60	143	2.38
519.lbm_r	52	141	2.71	56	77	1.38	54	136	2.52	56	77	1.38
531.deepsjeng_r	265	1163	4.39	269	743	2.76	277	763	2.75	276	764	2.77
429.mcf	14	74	5.29	14	32	2.29	17	50	2.94	19	43	2.26
458.sjeng	58	281	4.84	61	198	3.25	59	239	4.05	72	210	2.92
462.libquantum	128	416	3.25	135	278	2.06	126	333	2.64	135	295	2.19
470.lbm	24	57	2.38	25	28	1.12	27	57	2.11	100	112	1.12
471.omnetpp	143	1317	9.21	143	1107	7.74	174	929	5.34	83	352	4.24
Average	92.12	459.25	4.54	94.87	323.5	2.85	98.75	330	3.09	100.12	249.5	2.41

Table 5: Number of spills (S) and reloads (R) and reloads per spill ratio (R/S) induced by different allocators in x86

for compiler optimizations. They include prediction of vectorization factor [Haj-Ali *et al.*, 2020], unroll factors [Stephenson and Amarasinghe, 2005], making inlining decisions [Simon *et al.*, 2013], predicting thread coarsening factor [Magni *et al.*, 2013], estimation of throughput of a code [Mendis *et al.*, 2019a], etc. Such recent works use learned embeddings for representing the input programs to the ML model. Several ways of representing programs have been proposed [Alon *et al.*, 2019; Ben-Nun *et al.*, 2018; Sui *et al.*, 2020]. In this work, we use an extension of IR2Vec [VenkataKeerthy *et al.*, 2020] to represent programs.

Reinforcement Learning solutions for optimizations   Recently, several works have been proposed that use machine learning reinforcement learning based solutions for compilation. This can possibly be because such problems are computationally hard to determine the ground truth, which would otherwise be necessary. Mammadli et al [2020] and Huang et al [2019b] proposed a phase ordering solution using reinforcement learning that would result in better optimization characteristics when compared to O3. Khadka et al [2021] formulate memory placement in neural network accelerators as a reinforcement learning problem to optimize memory movement. Haj-Ali et al [2020] model loop vector factor prediction using reinforcement learning to obtain better performance when compared to the native LLVM's vectorizer. Mendis et al used a variation of reinforcement learning, called imitation learning to model SLP vector-

ization decisions by imitating a solver [2019b]. Other optimization works that use reinforcement learning can be found in the survey paper by Wang et al [2018].

ML and Register Allocation   Ours is the first end-to-end application of reinforcement learning for solving the register allocation problem. An initial attempt to solve this problem using ML models was by Das et al [2020], where they use an LSTM network to come up with an initial graph coloring scheme, which then undergoes a correction phase to rectify the inconsistency in coloring interfering nodes. Their work focuses on solving the graph coloring problem, and to our understanding, the solution was not integrated to obtain the final register assignments so as to study the performance.

Several works [Huang *et al.*, 2019a; Lemos *et al.*, 2019; Santos, 2020; Musliu and Schwengerer, 2013] have been proposed to solve the conventional graph coloring problem using machine learning. These however may not be applied directly to obtain an optimal register allocation scheme, as such approaches do not consider program-specific information like live ranges to arrive at the final coloring.

RL Framework to support compiler optimizations   With the growing interest in applying machine learning in compilation, a couple of frameworks have been proposed to ease the integration of compiler APIs and RL environments for training and inference.

Cummins et al. proposed CompilerGym [2021a], a python library for compiler optimization problems. It

exposes RL environments and datasets, integrated with popular compilers and tools.

Another framework, MLGO [Trofin *et al.*, 2021] integrates trained machine learning models within the LLVM compiler. For this purpose, the compiler loads a trained model and accesses it via C++ APIs of Tensorflow or ahead-of-time generated code (release mode). The framework is used in production, with improved decisions for inlining for size, and live range eviction (in register allocation) when compared to the compiler's default heuristics.[2]

The problem of supporting multiple stages of register allocation involves a two-way communication between the model and the compiler. To our understanding, these frameworks do not support such a communication, while our framework is designed to handle it.

Register Allocation Traditionally several approaches [Chaitin *et al.*, 1981; Briggs *et al.*, 1994; Chow and Hennessy, 1990; Chen *et al.*, 2018; Huang *et al.*, 2012] have been proposed to solve the register allocation problem using heuristics. Often the register allocation problem is understood to interfere with other optimizations like instruction selection. Hence, different works have been proposed to model them as a combinatorial problem of these two using solvers [Wilson *et al.*, 1994; Gebotys, 1997; Bashford and Leupers, 1999; Nagarakatte and Govindarajan, 2007; Lozano *et al.*, 2019].

8 Conclusion and Future Work

We propose an architecture-independent Reinforcement Learning solution to the Register Allocation problem. We use a multi-agent hierarchical approach to learn an optimal policy for the different sub-tasks of register allocation, including coloring, live range splitting, and spilling. Semantic correctness is ensured by the constraints encoded as the action masks for the agents. Our method exhibits better allocations, improved runtime, and less memory accesses when compared to the standard register allocators of LLVM.

We plan to address other sub-tasks of register allocation: coalescing, multi-allocation and register packing. We will open-source the framework in the near future.

References

[Alon *et al.*, 2019] U Alon, M Zilberstein, O Levy, and E Yahav. Code2vec: Learning distributed representations of code. POPL, pages 40:1–40:29, 2019.

[Appel and George, 2001] A W. Appel and L George. Optimal spilling for cisc machines with few registers. PLDI '01, page 243–253, 2001.

[Ashouri *et al.*, 2017] A. H. Ashouri, A Bignoli, G Palermo, C Silvano, S Kulkarni, and J Cavazos. Micomp: Mitigating the compiler phase-ordering problem using optimization sub-sequences and machine learning. 14(3), 2017.

---

[2]https://github.com/google/ml-compiler-opt

[Barik *et al.*, 2006] R Barik, C Grothoff, R Gupta, V Pandit, and R Udupa. Optimal bitwise register allocation using integer linear programming. LCPC, pages 267–282, 2006.

[Bashford and Leupers, 1999] Steven Bashford and Rainer Leupers. Phase-coupled mapping of data flow graphs to irregular data paths. *Design automation for embedded systems*, 4(2):119–165, 1999.

[Ben-Nun *et al.*, 2018] T Ben-Nun, A S Jakobovits, and T Hoefler. Neural code comprehension: A learnable representation of code semantics. NIPS'18, pages 3589–3601, 2018.

[Bordes *et al.*, 2013] A Bordes, N Usunier, A Garcia-Durán, J Weston, and O Yakhnenko. Translating embeddings for modeling multi-relational data. NIPS'13, pages 2787–2795, 2013.

[Bouchez *et al.*, 2006] F Bouchez, A Darte, C Guillon, and F Rastello. Register allocation: What does the np-completeness proof of chaitin et al. really prove? or revisiting register allocation: Why and how. In *Int. Workshop on LCPC*, pages 283–298, 2006.

[Briggs *et al.*, 1994] P Briggs, K D. Cooper, and L Torczon. Improvements to graph coloring register allocation. *ACM Trans. Program. Lang. Syst.*, 16(3):428–455, may 1994.

[Chaitin *et al.*, 1981] G J Chaitin, M A Auslander, A K Chandra, J Cocke, M E Hopkins, and P W Markstein. Register allocation via coloring. *Computer languages*, 6(1):47–57, 1981.

[Chang *et al.*, 1997] CM Chang, CM Chen, and CT King. Using integer linear programming for instruction scheduling and register allocation in multi-issue processors. *Computers & Mathematics with Applications*, 34(9):1–14, 1997.

[Chen *et al.*, 2018] WY Chen, GY Lueh, P Ashar, K Chen, and B Cheng. Register allocation for intel processor graphics. CGO 2018, page 352–364, 2018.

[Cho *et al.*, 2014] K Cho, M B van, C Gulcehre, D Bahdanau, F Bougares, H Schwenk, and Y Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. EMNLP 2014, pages 1724–1734, October 2014.

[Chow and Hennessy, 1990] F C. Chow and J L. Hennessy. The priority-based coloring approach to register allocation. *ACM TOPLAS*, 12(4):501–536, oct 1990.

[Cummins and others, 2021a] C Cummins et al. Compilergym: Robust, performant compiler optimization environments for AI research, 2021.

[Cummins and others, 2021b] C Cummins et al. Programl: A graph-based program representation for data flow analysis and compiler optimizations. volume 139 of *ICML'21*, pages 2244–2253, 18–24 Jul 2021.

[Das *et al.*, 2020] D Das, S A Ahmad, and V Kumar. Deep learning-based approximate graph-coloring algorithm for register allocation. Workshop on the LLVM Compiler Infrastructure in HPC, pages 23–32, 2020.

[Fursin and others, 2011] G Fursin et al. Milepost gcc: Machine learning enabled self-tuning compiler. *IJPP*, 39(3):296–327, Jun 2011.

[Garey and Johnson, 1979] M R Garey and D S Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[Gebotys, 1997] C.H. Gebotys. An efficient model for dsp code generation: performance, code size, estimated energy. In *Proceedings. Tenth International Symposium on System Synthesis (Cat. No.97TB100114)*, pages 41–47, 1997.

[Haj-Ali *et al.*, 2020] 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. CGO'20, page 242–255, 2020.

[Hames and Scholz, 2006] L Hames and B Scholz. Nearly optimal register allocation with pbqp. In *Joint Modular Languages Conference*, pages 346–361. Springer, 2006.

[Huang *et al.*, 2012] Y Huang, M Zhao, and C J Xue. Wcet-aware re-scheduling register allocation for real-time embedded systems with clustered vliw architecture. LCTES '12, page 31–40, 2012.

[Huang *et al.*, 2019a] J Huang, M M A Patwary, and G F. Diamos. Coloring big graphs with alphagozero. *CoRR*, abs/1902.10162, 2019.

[Huang *et al.*, 2019b] Q Huang, A Haj-Ali, W S Moses, J Xiang, I Stoica, K Asanović, and J Wawrzynek. Autophase: Compiler phase-ordering for hls with deep reinforcement learning. International Symposium on Field-Programmable Custom Computing Machines (FCCM), pages 308–308, 2019.

[Jakob, 2011] S O Jakob. Greedy Register Allocation in LLVM 3.0. http://blog.llvm.org/2011/09/greedy-register-allocation-in-llvm-30.html, 2011.

[Khadka *et al.*, 2021] S Khadka, E Aflalo, M Mardar, A Ben-David, S Miret, S Mannor, T Hazan, H Tang, and S Majumdar. Optimizing memory placement using evolutionary graph reinforcement learning. ICLR'21, 2021.

[Kim *et al.*, 2021] M Kim, JK Park, and SM Moon. Irregular register allocation for translation of test-pattern programs. *ACM Trans. Archit. Code Optim.*, 18(1), dec 2021.

[Kuchcinski, 2003] K Kuchcinski. Constraints-driven scheduling and resource assignment. *ACM TODAES*, 8(3):355–383, 2003.

[Lattner and Adve, 2004] C Lattner and V Adve. Llvm: A compilation framework for lifelong program analysis & transformation. CGO'04, page 75, 2004.

[Lemos *et al.*, 2019] H Lemos, M Prates, P Avelar, and L Lamb. Graph colouring meets deep learning: Effective graph neural network models for combinatorial problems. ICTAI'19, pages 879–885, 2019.

[Lozano *et al.*, 2012] R Lozano, M Carlsson, F Drejhammar, and C Schulte. Constraint-based register allocation and instruction scheduling. In *ICPPCP*, pages 750–766, 2012.

[Lozano *et al.*, 2019] R Lozano, M Carlsson, G Hjort Blindell, and C Schulte. Combinatorial register allocation and instruction scheduling. *ACM TOPLAS*, 41(3), jul 2019.

[Magni *et al.*, 2013] A Magni, C Dubach, and M F P O'Boyle. A large-scale cross-architecture evaluation of thread-coarsening. SC'13, pages 11:1–11:11, 2013.

[Mammadli *et al.*, 2020] R Mammadli, A Jannesari, and F Wolf. Workshop on the LLVM Compiler Infrastructure in HPC, pages 1–11, 11 2020.

[Mendis *et al.*, 2019a] C Mendis, A Renda, S Amarasinghe, and M Carbin. Ithemal: Accurate, portable and fast basic block throughput estimation using deep neural networks. volume 97 of *ICML'19*, pages 4505–4515, 09–15 Jun 2019.

[Mendis *et al.*, 2019b] C Mendis, C Yang, Y Pu, S Amarasinghe, and M Carbin. Compiler auto-vectorization with imitation learning. volume 32 of *NeurIPS'19*, 2019.

[Mikolov *et al.*, 2013] T Mikolov, K Chen, G Corrado, and J Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.

[Musliu and Schwengerer, 2013] N Musliu and M Schwengerer. Algorithm selection for the graph coloring problem. In *ICLIO*, pages 389–403, 2013.

[Nagarakatte and Govindarajan, 2007] S G Nagarakatte and R Govindarajan. Register allocation and optimal spill code scheduling in software pipelined loops using 0-1 ilp formulation. CC'07, pages 126–140, 2007.

[Pennington *et al.*, 2014] J Pennington, R Socher, and C D Manning. Glove: Global vectors for word representation. EMNLP'14, pages 1532–1543, 2014.

[Poletto and Sarkar, 1999] M Poletto and V Sarkar. Linear scan register allocation. *ACM TOPLAS*, 21(5):895–913, sep 1999.

[Santos, 2020] H L Santos. Solving the decision version of the graph coloring problem: a neural-symbolic approach using graph neural networks. Master's thesis, Universidade Federal do Rio Grande do Sul, 2020.

[Simon *et al.*, 2013] D Simon, J Cavazos, C Wimmer, and S Kulkarni. Automatic construction of inlining heuristics using machine learning. CGO'13, page 1–12, 2013.

[Stephenson and Amarasinghe, 2005] M. Stephenson and S. Amarasinghe. Predicting unroll factors using supervised classification. CGO'05, pages 123–134, March 2005.

[Sui *et al.*, 2020] Y Sui, X Cheng, G Zhang, and H Wang. Flow2vec: Value-flow-based precise code embedding. *Proc. ACM Program. Lang.*, 4(OOPSLA), Nov 2020.

[Trofin *et al.*, 2021] Mircea Trofin, Yundi Qian, Eugene Brevdo, Zinan Lin, Krzysztof Choromanski, and David Li. MLGO: a machine learning guided compiler optimizations framework. *CoRR*, abs/2101.04808, 2021.

[VenkataKeerthy *et al.*, 2020] S. VenkataKeerthy, R Aggarwal, S Jain, M S Desarkar, R Upadrasta, and Y. N. Srikant. IR2Vec: LLVM IR Based Scalable Program Embeddings. *ACM Trans. Archit. Code Optim.*, 17(4), December 2020.

[Wang and O'Boyle, 2018] Z Wang and M O'Boyle. Machine learning in compiler optimization. *Proceedings of the IEEE*, 106(11):1879–1901, 2018.

[Wilson *et al.*, 1994] T. Wilson, G. Grewal, B. Halley, and D. Banerji. An integrated approach to retargetable code generation. In *Proceedings of 7th International Symposium on High-Level Synthesis*, pages 70–75, 1994.

[Xavier *et al.*, 2012] T Xavier, G Oliveira, E Lima, and A Silva. A detailed analysis of the llvm's register allocators. In *2012 31st International Conference of the Chilean Computer Science Society*, pages 190–198, 2012.