

# NOELLE Offers Empowering LLVM Extensions

Angelo Matni  
Northwestern University  
angelomatni2018@u.northwestern.edu

Enrico Armenio Deiana  
Northwestern University  
EnricoDeiana2020@u.northwestern.edu

Yian Su  
Northwestern University  
yiansu2018@u.northwestern.edu

Lukas Gross  
Northwestern University  
lukasgross@u.northwestern.edu

Souradip Ghosh  
Northwestern University  
souradipghosh2021@u.northwestern.edu

Sotiris Apostolakis  
Princeton University  
sa8@cs.princeton.edu

Ziyang Xu  
Princeton University  
ziyangx@cs.princeton.edu

Zujun Tan  
Princeton University  
zujunt@cs.princeton.edu

Ishita Chaturvedi  
Princeton University  
ishitac@cs.princeton.edu

David I. August  
Princeton University  
august@cs.princeton.edu

Simone Campanoni  
Northwestern University  
simonec@eecs.northwestern.edu

## Abstract

Modern and emerging architectures demand increasingly complex compiler analyses and transformations. As the emphasis on compiler infrastructure moves beyond support for peephole optimizations and the extraction of instruction-level parallelism, they should support custom tools designed to meet these demands with higher-level analysis-powered abstractions of wider program scope. This paper introduces NOELLE, a robust open-source domain-independent compilation layer built upon LLVM providing this support. NOELLE is modular and demand-driven, making it easy-to-extend and adaptable to custom-tool-specific needs without unduly wasting compile time and memory. This paper shows the power of NOELLE by presenting a diverse set of ten custom tools built upon it, with a 33.2% to 99.2% reduction in code size (LoC) compared to their counterparts without NOELLE.

## 1 Introduction

The compiler community is on the front lines to satisfy the continuous demand for computational performance and energy efficiency. The focus of compiler advancements is shifting beyond peephole optimizations and the extraction of instruction-level parallelism. More aggressive optimizations and more sophisticated analyses with wider scope are required to accommodate the needs of emerging architectures and applications.

Modern compilers use low-level intermediate representations (IR) to perform optimizations that are language-agnostic and architecture-independent, such as LLVM IR from the LLVM compiler framework [6, 37] and GIMPLE from GCC [4]. Low-level IR, along with a set of low-level abstractions built

around it, is designed to aid program analyses and optimizations and has shown its value for peephole optimizations and extraction of ILP. However, low-level abstractions are not enough for advanced code analyses and transformations. Consider automatic parallelization, one of the most powerful program optimization techniques, exists only in a basic form [1, 2, 7], or does not exist at all in most general-purpose compilers. This paper shows that with proper abstractions, a daunting automatic parallelization transformation can be implemented in fewer than a thousand lines of code.

Advanced code analyses and transformations go hand in hand with higher-level abstractions, as shown by many existing compilers or frameworks. Several compiler infrastructures that support automatic parallelization [3, 10, 20] all operate on high-level abstractions and perform source-to-source translation. The recent success of domain-specific compilers/frameworks also proves the importance of high-level abstractions for optimizations by uncovering optimization opportunities at a domain-specific graph or operator level [5, 12]. However, these compilers limit themselves to specific program languages or problem domains, and miss opportunities only presented in low-level IRs, including more fine-grained operations and more canonical code patterns.

The combination of higher-level abstractions and lower-level IR is the key to advanced program analysis and optimizations. The claim can be found in the SUIF compiler [11], which provides low-level IR as well as higher-level constructs [51]; and the IMPACT compiler [25], which provides hierarchical IRs to enable optimizations at different levels. Despite the claim, we are not aware of actively-maintained domain-independent compilers that fulfill this combination.

While LLVM has become the de-facto compiler infrastructure to build upon, it does not provide proper abstractions for

advanced analyses and transformations, including abstractions designed to describe properties of a wider code scope (e.g., program dependence graph, program call graph) or abstractions that provide the mechanisms for advanced code transformations (e.g., loop transformations, code scheduler, task creation). These abstractions can ease the implementation of new transformations and make the existing code transformations available in LLVM more powerful.

We propose a new open-source compilation layer called NOELLE that brings forth abstractions for advanced code analyses and transformations. To demonstrate the importance of NOELLE, we have implemented ten advanced code transformations, nine of which need only a few lines of code. Only one of these transformations is already available in LLVM (loop invariant code motion). We will show that our version is significantly more powerful, requires significantly fewer lines of code, and the implementation is more elegant than the LLVM counterpart. The other nine transformations are missing in LLVM because they are challenging to implement with the low-level abstractions LLVM provides.

We have implemented a variety of code transformations upon NOELLE: a few parallelizing compilers, a Pseudo-Random value generator selector, a comparison optimization for timing speculative micro-architectures, a dead function elimination, a memory guard optimization, a code analysis and transformation to replace hardware interrupts, and a loop invariant code motion. We call them NOELLE’s custom tools. It is a challenge to implement these custom tools only using the low-level abstractions provided by LLVM. Relying on NOELLE, instead, most tools are implemented in only a few hundred lines of code. We tested all these tools on 41 benchmarks from three benchmark suites (SPEC CPU2017, PARSEC 3.0, and MiBench). All these tools improve the quality of the code generated by LLVM with its highest level of optimization. Finally, the high heterogeneity between these ten custom tools suggests NOELLE provides general abstractions and support for a wide variety of advanced code analyses and transformations. Finally, we have released NOELLE publicly (<https://github.com/scampanoni/noelle>).

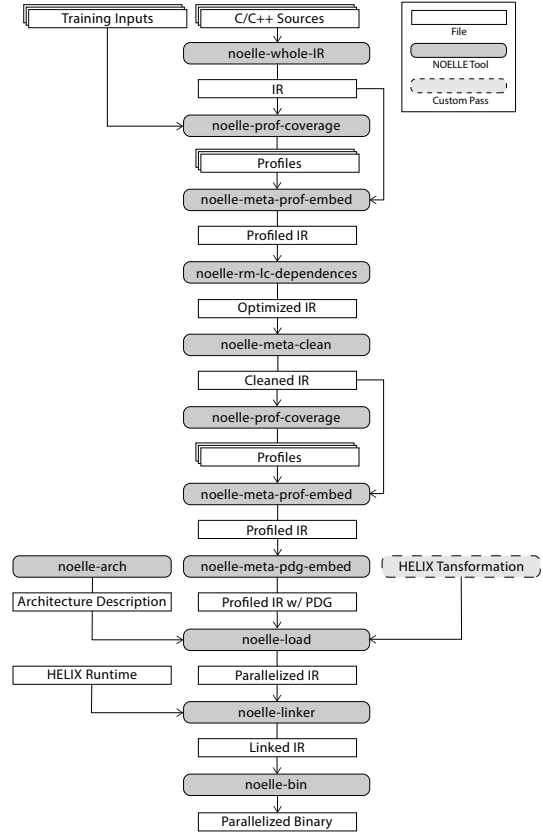
This paper:

- introduces NOELLE, a robust open-source domain-independent compilation layer built upon LLVM;
- describes the abstractions provided by NOELLE (Section 2.2) that ease the development of advanced code transformations and analyses;
- presents the tools provided by NOELLE (Section 2.3) that ease the deployment of custom compilation tool-chains;
- describes the testing infrastructure provided by NOELLE (Section 2.4) that enables automatic testing of custom tools;
- describes a diverse set of ten custom tools built upon NOELLE (Section 3) and highlights the benefits of NOELLE’s custom tools compared to vanilla LLVM (Section 4.2);
- evaluates the accuracy of NOELLE’s abstractions (Section 4.1); and,

- further motivates the need for NOELLE by comparing it with prior work (Section 5).

## 2 NOELLE and Its Abstractions

Next, we describe NOELLE, its abstractions, and its tools.

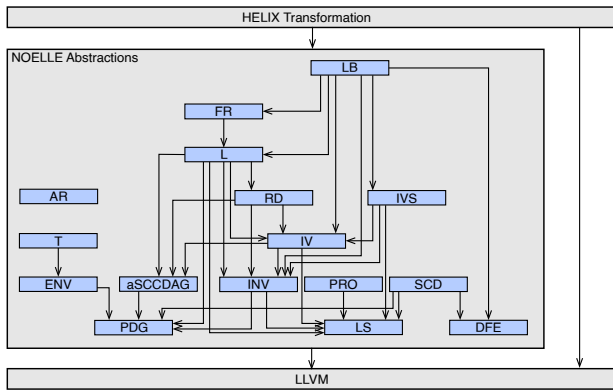


**Figure 1.** Compilation flow of the HELIX custom tool using NOELLE tools and a custom pass, HELIX Transformation. Figure 2 shows in detail how to build HELIX Transformation using NOELLE abstractions.

### 2.1 NOELLE in a Nutshell

The goal of NOELLE is to provide abstractions that enable a simple implementation of complex code analyses and transformations (we call them custom tools) that target wide program scopes. Custom tools built upon NOELLE include LLVM passes that work at the IR level to perform their code analyses and transformations. Allowing these custom tools to be easily implementable and maintainable requires simple and domain-independent abstractions powered by either accurate low-level code analyses or complex low-level code transformations. NOELLE provides such abstractions (Section 2.2) with a modular design allowing its users to pay only the cost of creating the abstractions requested.

NOELLE’s abstractions are powered by code analyses, some of which are provided by third parties. For example,



**Figure 2.** HELIX transformation, a custom pass written using NOELLE abstractions. Arrows in the graph describe the dependence between analyses. Refer to Table 1 for descriptions of all NOELLE abstractions. Table 4 and Table 2 describe abstractions used per custom and NOELLE tool respectively.

the PDG abstraction NOELLE provides is computed by running several alias analyses implemented by external codebases (SCAF [16] and SVF [47]). Moreover, NOELLE’s modular design makes it easy to extend the list of external code analyses that power NOELLE’s abstractions. NOELLE also provides tools (Section 2.3) for faster user-specific compilation flows. Finally, NOELLE provides a testing infrastructure (Section 2.4) to facilitate automatic testing of NOELLE itself as well as custom tools built upon it.

**Input and Output** The input of a compilation flow built upon NOELLE is the source code of a program and optionally, a set of training inputs that could be used for profile-guided or autotuning-based custom tools. The output is a binary for a target architecture supported by vanilla LLVM backends.

**An Example of Compilation Flow** NOELLE enables its users to deploy custom compilation flows by providing a set of tools, described in Section 2.3. Next, we describe an example of a compilation flow built using NOELLE’s tools (shown in Figure 1). This is the compilation flow used by the custom tool HELIX (further described in Section 3).

Each source file composing a program being compiled is consumed by `noelle-whole-IR`, which outputs a single LLVM IR file that includes the whole program’s code as well as options to use to generate the final binary (e.g., the libraries to link with). Then, using training inputs given to NOELLE, `noelle-prof-coverage` runs several profilers to collect statistics about the single IR file’s execution. These statistics include the hotness of code regions (e.g., a loop, a basic block), loop-specific information (e.g., the total number of iterations of a loop, the average number of iterations per invocation of a loop), and function-specific information (e.g., number of invocations of a function, the average number of recursive calls of a recursive function). The program’s profiles are then embedded into the IR file by `noelle-meta-prof-embed`. The generated IR is consumed

by `noelle-rm-lc-dependences`, which applies a set of code transformations that aim to reduce loop-carried data dependences in hot loops (the minimum hotness required to consider a loop). The generated IR is now more amenable to loop-centric code parallelization techniques. The tool `noelle-meta-clean` cleans all NOELLE-specific metadata from the IR file. Then, `noelle-prof-coverage` and the tool `noelle-meta-prof-embed` re-generate and embed the program’s profiles, respectively. Then, `noelle-meta-pdg-embed` computes the program dependence graph (PDG) and embeds it as metadata inside the IR file. The `noelle-arch` computes architecture-specific profiles (e.g., communication latency between cores). Its output is used by the HELIX transformation. Finally, the `noelle-load` tool is invoked, which loads in memory NOELLE’s compilation layer, to run the HELIX transformation. The HELIX transformation relies on NOELLE’s abstractions to parallelize hot loops. The generated parallelized IR file is then consumed by `noelle-linker`, which embeds the HELIX-specific runtime into the IR. Finally, `noelle-bin` generates the parallel binary.

## 2.2 NOELLE’s Abstractions

Next, we describe the abstractions that NOELLE provides to its users. NOELLE’s abstractions (summarized by Table 1) are demand-driven to preserve compilation time and memory. Hence, users only pay for the abstractions they need. In other words, if a user does not need the program dependence graph (PDG), then it will not pay the cost of analyzing the program to compute its dependences.

**PDG.** NOELLE provides the Program Dependence Graph (PDG) representation of a program [30]. This is obtained by extending NOELLE’s *dependence graph*, a templated class designed to represent a generic graph of directed dependences between nodes. What constitutes a node is decided when the class is instantiated. For example, the PDG instantiates this templated class with the LLVM instruction class. Hence, the nodes of the PDG are the instructions of a program. Each edge of the dependence graph contains attributes to differentiate between control and data dependences. Data dependences are further characterized based on the dependence type (Read-After-Write, Write-After-Write, Write-After-Read), whether it is loop-carried, dependence distance, whether it is memory or register dependence, and whether it is an apparent or actual dependence [26].

An analysis or transformation (i.e., pass) built upon NOELLE can use the PDG abstraction to create loop dependence graphs and function dependence graphs. The former is a dependence graph of a specific loop. The latter refers to dependences only between the instructions of a function. When a pass requests the loop dependence graph from a PDG, NOELLE runs loop-centric analyses to refine (and improve the precision about) the dependences that are included in the PDG for the specific loop in-question.

**Table 1.** Abstractions provided by NOELLE

Abstraction	Description	LoC	Depends on
PDG	All dependences between instructions of a program	6775	
aSCCDAG	SCCDAG of a loop with attributes on each SCC (e.g., an SCC has loop-carried data dependence, it is reducible)	4517	PDG
Call graph (CG)	Complete call graph of a program including indirect calls and their possible callees	620	PDG
Environment (ENV)	Variables needed by a task to execute (live-ins and live-outs)	991	PDG
Task (T)	Code region (and its inputs and outputs) executed by a thread	297	ENV
Data-flow engine (DFE)	Optimized engine to quickly evaluate data flow equations provided as inputs	332	
Loop structure (LS)	Describe the structure of a loop, its exits, latches, header, pre-header, basic blocks.	301	
Profiler (PRO)	Set of profilers at the IR level	1625	LS
Scheduler (SCD)	Mechanisms to change the schedule of instructions within and between basic blocks	1523	PDG, LS, DFE
Invariant (INV)	Instructions, values, or memory locations that are loop invariants for a given loop	137	PDG, LS
Induction variable (IV)	Induction variables of a loop including the identification of the governing one (if it exists)	352	LS, INV aSCCDAG
Induction variable stepper (IVS)	Modifies the code of a loop to implement a change in step value of its induction variables	425	LS, INV, IV
Reduction (RD)	Identification and capability of reducing variables of a loop	868	aSCCDAG, INV, IV
Loop (L)	Canonical loop with its dependence graph, its SCCDAG, its invariants, its induction variables, and its exits	1508	LS, PDG, IV, INV, aSCCDAG, RD
Forest (FR)	Forest of trees with the capability to adjust when a node is deleted to keep the connections between the parent and the children of the deleted node	202	L, CG
Loop builder (LB)	Set of loop transformations that modify a loop (e.g., split a loop, translate do-while loops to while form and vice versa)	4535	FR, L, DFE, IV, IVS, INV
Islands (ISL)	Capability to identify the disconnected sub-graphs of a graph (e.g., call graph, PDG)	56	PDG, CG
Architecture (AR)	Description of the underlying architecture in terms of logical/physical cores, NUMA nodes. It also provides the measured latencies and bandwidths between pairs of cores	381	
Others		691	
<b>LoC of NOELLE’s abstractions</b>		<b>26142</b>	

Users of this abstraction often want to know not only about the nodes of a dependence graph that belong to a related code region (e.g., instructions of a loop for a loop dependence graph) but also about the inputs, the outputs, or both of the graph. For example, a parallelizing code transformation of a loop needs to know the live-in and live-out sets of the target loop. Because of this need, the templated class *dependence graph* offers two sets of nodes, the internal and the external ones. The former belong to the related code region; the latter represents the live-ins, live-outs, or both of that code region. The computation of both sets of nodes is computed by NOELLE when a pass requests either a loop dependence graph or a function dependence graph.

**aSCCDAG.** Advanced code transformations like parallelization techniques can be implemented as different strategies to schedule instances of the nodes that compose the SCCDAG of a loop [43, 49]. For instance, HELIX distributes instances of a given SCCDAG node around the cores. DSWP instead distributes nodes of an SCCDAG between cores while keeping all instances of a given node within the same core. Hence, an important abstraction is the SCCDAG. To this end, we introduce the *augmented SCCDAG* abstraction or *aSCCDAG*. An aSCCDAG of a given loop is a complete description of loop dependences, including those with the rest of the program.

A node of an aSCCDAG can be *Independent*, *Sequential*, or *Reducible*. This categorization of a node  $n$  depends on the relation between the instructions’ dynamic instances

included in  $n$  for a given loop invocation. If all these instances are independent of each other, then  $n$  is tagged as *Independent*. If an instance of an instruction of  $n$  depends on another instance of an instruction of  $n$ , then this node is tagged as *Sequential*. Finally, if there are dependences between instances of  $n$ , but they are reducible by a reduction code transformation (e.g., by cloning the defined variable  $s$  in  $s += work(d)$ ), then  $n$  is tagged as *Reducible*, and the related reduction is described within the node.

**Call graph (CG).** NOELLE provides the *call graph* of a program where nodes are functions, and edges indicate a given function invokes another. This abstraction relies on the PDG to compute the possible callees of an indirect call. Edges of the NOELLE’s call graph can be *must* or *may* depend on whether a given caller-callee relation is proved to hold or not. Each edge has sub-edges to indicate with which specific instructions a caller invokes another function. Finally, CG can compute the set of disconnected islands of such a graph.

NOELLE’s call graph differentiates with LLVM’s one by being complete: the latter does not compute an indirect call’s possible callees. By being complete, NOELLE’s call graph enables custom tools to assume that the call graph’s lack of an edge means a function cannot invoke another. CG is used by the *DeadFunctionEliminator* custom tool built upon NOELLE, aiming to reduce the binary size of a program.

**Environment (ENV).** NOELLE offers the *Environment* abstraction, which is an array of pointers of variables. Variables

**Table 2.** NOELLE's tools

Tool	Description	LoC	Depends on
noelle-whole-IR	Generate a single IR file from C/C++ source files embedding the compilation options as metadata inside the generated IR file	1522	
noelle-rm-lc-dependences	Transform loops to remove as many loop-carried data dependences as possible	912	aSCCDAG, CG, L, PRO, FR, LB
noelle-prof-coverage	Inject code into the IR file given as input to profile IR instructions	1761	PRO, FR, CG
noelle-meta-prof-embed	Embed profiles into the IR file given as input	152	PRO, FR, CG
noelle-meta-pdg-embed	Compute and embed the PDG into the IR file given as input	451	PDG
noelle-load	Load the NOELLE abstractions into memory without computing them	12	
noelle-arch	Generate a file that describes the underlying architecture and its profiles (e.g., core-to-core latencies)	259	AR
noelle-linker	Links IR files together while preserving the semantic of metadata generated by NOELLE's tools	59	
noelle-bin	Generate a standalone binary from an IR file using the compilation options specified as metadata inside the IR file given as input	15	
	<b>LoC of NOELLE's tools</b>	<b>5143</b>	

within an Environment represent the incoming and outgoing values from and to a set of instructions. This set of instructions is described by a subset of the nodes of an aSCCDAG. An example pass that relies on the Environment abstraction is a parallelization technique that needs to propagate values explicitly between the cores. Finally, NOELLE provides *Environment Builder* to create, modify, and query environments.

**Task (T).** NOELLE offers the *Task* abstraction to describe a code region that runs sequentially. Parallelization techniques use the above abstraction in the following way. Nodes within an aSCCDAG are partitioned into tasks. An Environment is created for each task. At runtime, tasks are submitted to a thread-pool, which will run them in parallel across the cores. The explicit forwarding of data values between tasks is performed by loading/storing values from/to variables pointed by their environments.

**Data flow engine (DFE).** NOELLE provides a *data flow engine* that can be used to implement data flow analyses. DFE implements conventional optimizations like bitvectors, basic block granularity optimization, working list algorithm, and loop-based priority [17]. Finally, NOELLE provides a set of common data flow analyses that rely on DFE.

**Profiler (PRO).** NOELLE provides several code profilers, the ability to embed their results into IR files, and abstractions to facilitate high-level queries on such data. Examples of queries that can be performed are the hotness of a code region (e.g., a loop, an SCC of a dependence graph), loop-specific information (e.g., loop iteration count, average loop iteration count per invocation), and function-specific information (e.g., the average number of times that an invocation of a function invokes another).

**Scheduler (SCD).** NOELLE provides the *scheduler* abstraction that offers the capability of moving instructions within and among basic blocks while preserving the original code semantics. The scheduler relies on the PDG abstraction to guarantee semantic preservation. The abstraction provides a hierarchy of schedulers starting from a generic one and including loop-specific and within-basic-block schedulers.

Each scheduler augments the generic capabilities with specialized capabilities (e.g., reducing the header size of a loop).

**Loop Builder (LB).** NOELLE offers the *loop builder* abstraction that enables passes to modify/create/delete loops. LB is similar to the IRBuilder abstraction offered by LLVM, but instead of targeting instructions, LB targets loops.

**Induction variables (IV).** NOELLE provides the *induction variable* abstraction. Because LLVM's IR is in SSA form, the concept of the loop's induction variable is embodied by an SCC of the aSCCDAG of that loop. NOELLE's abstraction exposes such SCC, the starting and ending value of an induction variable, the step amount per loop iteration, and whether an induction variable controls the number of loop iterations that will be executed. We call governing induction variables those that control the number of loop iterations. Finally, IV exposes the potential relationship with other induction variables for those that are derived.

The main difference between LLVM's induction variable and NOELLE's version is that the former only provides the PHI instruction that composes the SCC of an induction variable that belongs to the loop header. Another difference is that NOELLE's version implements a more robust algorithm to detect governing induction variables based on the aSCCDAG abstraction. LLVM's implementation, instead, relies on the low-level def-use chains of the IR because of the lack of the SCCDAG abstraction within LLVM. NOELLE's IV, therefore, detects more governing induction variables.

**Induction Variable Stepper (IVS).** A common operation for modern and emerging code transformations is to modify the step of induction variables. For example, loop rotation needs to revert the step value of induction variables. Another example is an advanced DOALL parallelization, which needs to perform chunking between iterations to increase spatial locality. The NOELLE's abstraction *induction variable stepper* offers the capability to modify any step value of induction variables of a loop; users only need to specify the new step values, and the abstraction modifies the loop accordingly.

**Loop (L).** This abstraction includes a representation of the loop structure (called LS). The latter is equivalent to the loop

abstraction of LLVM. The abstraction L, instead, adds to LS the loop dependence graph (computed from the PDG) and the loop-specific instances of the abstractions IV and INV.

**Other abstractions.** Above, we have described the most important abstractions NOELLE provides. However, NOELLE provides additional abstractions used for simple compilation tasks such as *control equivalence*, *reduction operations*, *extendible metadata* attached to control structures like loops, *SCCDAG partitioner*, *forests(FR)*, and *graphs* designed to restore connections among remaining parts when a node is deleted, *architecture* to describe how logical cores are mapped to physical cores and NUMA nodes, and deterministic *IDs* for instructions, loops, functions, and basic blocks.

Furthermore, NOELLE offers a new implementation of the *loop structure (LS)*, *dominator*, and *scalar evolution* abstractions. This is because the LLVM abstractions computed by Function passes free their memory when they are invoked to analyze a different function. This generates subtle, but unfortunately common, bugs that affect module passes. The bug is generated when a module pass caches the pointers of the abstractions returned by a function pass applied to multiple functions. All previous pointers but the last one are invalid. This problem can become even more subtle because a function pass’s invocation can invalidate the abstraction returned by another function pass. To avoid this common bug, NOELLE offers implementations of these LLVM abstractions with the property that only their users can free these memory objects.

### 2.3 NOELLE’s Tools

NOELLE includes tools (Table 2) to help users deploy their compilation tool-chain. Next are the most important ones.

**noelle-whole-IR** generates a single IR file. Merging all bitcode into a single bitcode file is important for the analyses and transformations that span a wide code region (e.g., the whole program). Such an example is the alias analyses used by NOELLE to compute the PDG. This tool is based on `gllvm`.

**noelle-rm-lc-dependences** modifies an IR program to remove or reduce the impact of loop-carried data dependences (e.g., using Loop Builder to split a loop).

**noelle-prof-coverage** profiles IR code using representative program’s inputs. At the moment, NOELLE includes an instruction profiler, a branch profiler, and a loop profiler.

**noelle-meta-pdg-embed** computes the PDG of an IR file. This tool computes the PDG by invoking many time-consuming and accurate alias analyses that power NOELLE. Then, this tool embeds the computed PDG as metadata into the IR file so that NOELLE can re-construct the requested abstractions without requiring memory analyses. This tool relies on NOELLE’s PDG and IDs abstractions.

**noelle-load** loads the NOELLE’s layer in memory. Custom tools invoke NOELLE’s empowered LLVM pass by using `noelle-load` rather than the LLVM tool opt.

**noelle-arch** measures architecture-specific characteristics. At the moment, this tool measures the core-to-core latency and bandwidth. This tool also interacts with the tool `hwloc` [9] to find the number of physical and logical cores of the underlying platform, their mapping, and NUMA nodes.

### 2.4 NOELLE’s Testing

NOELLE provides a testing infrastructure composed of hundreds of regression tests, unit tests, and performance tests. These tests are micro C/C++ programs to illustrate corner cases or common code patterns found in popular benchmark suites such as SPEC CPU2017 and PARSEC 3.0. This testing infrastructure allows NOELLE’s users to quickly test their work with representative code patterns without paying the high compilation and profiling costs of the original codebase of the mentioned benchmark suites. Finally, this infrastructure is integrated with distributed systems, such as HTCondor and Slurm, to run tests in parallel across multiple machines. Optionally, NOELLE generates a bash file that executes all tests sequentially.

Tests are enabled by exposing NOELLE options and can be extended. This allows NOELLE’s users to surgically generate tests that stress a specific aspect of a specific code transformation. For example, a user can force a parallelizing custom tool to parallelize only a given loop.

### 2.5 Impact of NOELLE’s Abstractions

NOELLE’s abstractions may depend on each other to simplify design while keeping high precision. For example, the invariant abstraction (INV) uses the PDG abstraction to identify loop invariants. Next, we compare this NOELLE’s implementation with the LLVM’s to highlight the impact of building upon higher-level abstractions rather than lower-level ones.

Algorithm 1 shows the simplified logic of LLVM’s implementation that relies on low-level abstractions to decide whether a given instruction is a loop invariant. First, the algorithm checks if any operand of I is defined within loop L. If no operands are defined within L, it checks the type of the instruction I. If I is a load instruction, it checks if any other instruction of L can modify the same memory location accessed by I. If I is a store instruction, it checks if any memory use precedes I in L. If not, it checks no memory invalidation happens if I would be hoisted outside the loop. Finally, if I is a call instruction, it checks (i) if I can modify any memory location, (ii) if the only memory accessed are via arguments to the call, (iii) and if any sub-loop can modify the same memory accessed via arguments by the call I.

Algorithm 2 shows the NOELLE’s implementation that relies on the high-level PDG abstraction. It checks if I is currently under analysis (i.e., in the stack `s`). If not, it checks instruction that I depends on whether it is outside the loop or a loop invariant. Notice that this algorithm is smaller, simpler, and more precise than Algorithm 1 (Figure 4).

**Algorithm 1:** *isInvariant\_llvm(Instruction I, Loop L, Dominator DT, AliasAnalysis AA)*


---

```

Result: Return true if instruction I is an invariant in loop L
/* Simplified logic of LLVM implementation */
for operand in I.getOperands() do
  | if operand is defined in L then return False;
end
if isa<LoadInst>(I) then
  | for Instruction j in L do
  | | if getModRef(j, I) != NoMod then return False;
  | end
end
if isa<StoreInst>(I) then
  | for memory use MU in L do
  | | // Conservatively ensures no memory
  | | // use precedes this store
  | | if not DT.dominates(I, MU) then return False;
  | end
  | // Ensures no memory def/use would be
  | // invalidated by hoisting the store
  | M ← AA.getNearestDominatingMemoryAccess(I);
  | if M is in L then return False;
end
if call ← dyn_cast<CallInst>(I) then
  | if AA.getModRefBehavior(call) != NoMod then return False;
  | S ← AA.onlyMemoryAccessesAreArguments(call);
  | if not S then return False;
  | for Argument A of call do
  | | for sL in L->getSubLoops() do
  | | | for sl in sL do
  | | | | if AA.getModRef(A, sl) != NoMod then return False;
  | | | end
  | | end
  | end
end
return True;

```

---

**Algorithm 2:** *isInvariant\_noelle(Instruction I, Loop L, PDG dg, Stack s)*


---

```

Result: Return true if instruction I is an invariant in loop L
/* Implementation using high level abstraction PDG instead of
low level abstractions alias analysis and dominators */
if I in s then return False;
s.push(I);
for PDG dependence j to I do
  | if j is in L then
  | | inv ← isInvariant_noelle(j, L, dg, s);
  | | if not inv then return False;
  | end
end
s.pop();
return True;

```

---

### 3 Transformations Built Upon NOELLE

This section describes the code transformations built upon NOELLE. Table 3 summarizes them and their LoC.

Each transformation relies on several of NOELLE’s abstractions. Table 4 shows the abstractions used by them. It is important to notice that every abstraction is used by more than one custom tool suggesting their wide applicability.

**HELIX** parallelizes a loop by distributing its iterations between cores [23, 24, 42]. Each iteration is sliced into several sequential and parallel segments. Different instances of the same static sequential segment run sequentially between the cores while everything else can overlap.

HELIX uses PRO, FR, and L of NOELLE to identify the most profitable loops to parallelize. HELIX uses the PDG and ENV to identify and organize the live-in and live-out of each chosen loop. LB and T abstractions are then used to generate the parallel version of a loop.

HELIX uses aSCCDAG, INV, IV, and the RD abstractions to identify the SCCs that need to be executed sequentially. HELIX uses DFE to translate SCCs into sequential segments. SCD is then used to reduce the size of each sequential segment as well as to schedule them within the body of each parallelized loop. Moreover, HELIX uses IVS to perform chunking of loop iterations. Finally, HELIX uses AR to implement helper thread optimization [23].

**DSWP** parallelizes a loop by distributing its SCCs between cores [43]. Instances of a given SCC are executed by the same core to create a unidirectional communication between cores. DSWP uses NOELLE’s abstractions, similarly to how HELIX does while leveraging DSWP-specific knowledge to select the loops to parallelize and to parallelize them.

**CARAT** is co-designed with the underlying operating system to replace virtual memory. This compiler injects code to guard IR memory instructions that cannot be proved at compile time to be valid [46].

CARAT relies on the PDG, the aSCCDAG, and INV to identify the memory instructions that need guarding. Then, it uses DFE and PRO to avoid redundant guards of the same memory location. CARAT also uses L, LB, and IV to merge guards. Finally, SCD is used to place the guards in the code.

**Compiler-Based Timing** is co-designed with the underlying operating system to inject calls to OS routines [31] into a program. This compiler uses DFE and PRO to implement its specialized data flow analyses. It also uses L, FR, and LB to handle potentially-infinite loops. Finally, it uses CG to improve the accuracy of its time analyses.

**PRVJeeves** selects the pseudo-random value generators (PRVG) for a randomized program (e.g., Monte Carlo simulations) [38]. To do so, it uses the PDG, CG, and DFE to identify the allocations and uses of the PRVGs. Then, PRVJeeves uses PRO to prune the design space (e.g., PRVGs not used frequently are left unmodified). Moreover, it uses L, LB, INV, and IV to identify the uses of a vector of PRVGs. Finally, PRVJeeves uses SCD to place the uses of a PRVG in the code.

**DOALL** parallelizes a loop that has no loop-carried data dependencies by distributing its iterations among cores [34]. DOALL’s implementation uses NOELLE’s abstractions similarly to the other parallelizing compilers (DSWP and HELIX), the difference being the loop selection process and parts of the parallelized code generation. Yet, the loop selection process and parts of the parallelized-code generation are naturally different from the other parallelization techniques.

**Table 3.** Custom tools built upon NOELLE

Custom tool	Description	LLVM	LLVM + NOELLE	Percent reduction
Time Squeezer (TIME)	Compiler to optimize compare instructions for timing speculative architectures	510	92	82.0%
Compiler-based timing (COOS)	Compiler to inject calls to Operating System routines to replace hardware interrupts	1641	495	69.8%
Loop Invariant Code Motion (LICM)	Hoist loop invariants outside their loop	2317	170	92.7%
DOALL	Parallelizing compiler that applies the DOALL code parallelization technique	5512	321	94.2%
Dead Function Elimination (DEAD)	Reduce the number of functions without increasing the total number of IR instructions	7512	61	99.2%
DSWP	Parallelizing compiler that applies the DSWP code parallelization technique	8525	775	90.9%
HELIX	Parallelizing compiler that applies the HELIX code parallelization technique	15453	958	93.8%
PRVJeeves (PRVJ)	Compiler to select the Pseudo Random Value Generators for the program given as input	17863	456	97.4%
CARAT	Inject memory guards to potentially incorrect memory instructions	21899	595	97.3%
Perspective (PERS)	Parallelizing compiler that minimizes speculation and privatization costs	33998	22706	33.2%

**Loop Invariant Code Motion** hoists loop invariants outside their loop. It uses FR to hoist loop invariants from innermost loops to outermost ones. Then, it uses INV to identify instructions that could be hoisted. Finally, it uses LB to perform the hoist transformation.

**Time-Squeezer** generates code optimized for timing speculative micro-architectures [28, 29]. To this end, the compiler needs to decide when to swap the compare operands (and modify its uses), how to change the schedule of instructions, and where to inject instructions that modify the clock period of the underlying architecture. This custom tool uses DFE, L, and FR to decide where to inject clock-changing instructions. It then uses SCD to optimize the instruction sequence of a code region that uses the same clock period previously chosen. Finally, it uses ISL and PDG to analyze the compare instructions and their dependences.

## 4 Evaluation

This section presents evaluation results for NOELLE and the custom tools built upon NOELLE. Before presenting the results, we first describe our evaluation platform and our evaluation methodology. Our results show that each NOELLE’s abstraction can be used by several and significantly different custom tools. Results suggest that NOELLE’s implementation of a few abstractions that exist in LLVM is more precise than their LLVM counterparts. Finally, results suggest that we can build a custom tool in a few lines of code that is powerful enough to improve the performance or reduce the binary size compared to the mainline, widely adopted compilers like clang.

### 4.1 Experimental Setup

We have evaluated NOELLE and ten custom tools on the platform described next and by following the measurement methodology described here.

**Platform.** Our evaluation was done on a Dell PowerEdge R730 server with one Intel Xeon E5-2695 v3 Haswell processor running at 2.3GHz. The processor has 12 cores with 2-way hyperthreading, 35MB of last-level cache, and has a peak power consumption of 120W. The cores are supported by 256GB of main memory in 16 dual rank RDIMMs at 2133MHz. Turbo Boost was disabled, and no CPU frequency governors

**Table 4.** Each NOELLE’s abstraction is used by several custom tools.

Custom tool	NOELLE’s abstractions used																		
	PDG	asCCDAG	CG	ENV	T	D/E	PRO	SCD	L	LB	IV	IVS	INV	FR	ISL	RD	AR	LS	
HELIX	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓			✓	✓	✓
DSWP	✓	✓		✓	✓		✓	✓	✓	✓	✓	✓	✓	✓				✓	✓
CARAT	✓	✓				✓	✓	✓	✓	✓	✓		✓						✓
COOS			✓			✓	✓	✓	✓	✓				✓					✓
PRVJ	✓		✓			✓	✓	✓	✓	✓	✓		✓						✓
DOALL	✓	✓		✓	✓		✓		✓	✓	✓	✓	✓	✓			✓	✓	✓
LICM									✓	✓				✓	✓				✓
TIME	✓					✓		✓	✓	✓				✓	✓				✓
DEAD			✓													✓			
PERS	✓	✓																	

were used (i.e., all cores ran at a maximum frequency). The OS used is Red Hat Enterprise Linux Server 8 on kernel 4.18. NOELLE was built on top of LLVM 9 [37].


**Statistics and convergence.** Each data point we show in our evaluation is an average of multiple runs. We ran the relevant configurations as many times as necessary to achieve a tight confidence interval (95% of the measurements are within 5% of the mean).

### 4.2 Building Upon NOELLE Reduces Source Code

NOELLE simplifies the implementation of code analyses and transformations. Table 3 compares the implementations of 10 transformations when built upon NOELLE and when implemented only using LLVM abstractions. The reduction in LoC is significant, reducing the maintainability cost of these custom tools.

NOELLE abstractions are general enough to be useful by many and highly heterogeneous custom tools. Table 4 shows that each abstraction is used by several custom tools. For example, the loop builder (LB) is used by eight custom tools out of 10. Moreover, it is important to notice the heterogeneity of these custom tools that use (for example) LB: parallelizing transformations, loop invariant code motion (LICM), compare instruction optimization and code generation for timing speculative micro-architecture (TIME), memory guard injector and optimization (CARAT), PRVG selector (PRVJ), and scheduler of OS routines within applications (COOS).





**Figure 3.** While LLVM is capable of proving the non-existence of most dependences, NOELLE disproves more by relying on state-of-the-art alias analysis techniques (SCAF [16])

### 4.3 NOELLE Abstractions

Next, we compare the subset of NOELLE’s abstractions that are also available in LLVM. These abstractions are loop invariants, loop induction variables, and dependences.

Figure 3 shows that NOELLE’s implementation of dependences within the PDG abstraction is more accurate than LLVM’s abstraction. LLVM is capable of proving a significant fraction of potential memory dependences non-existing. NOELLE further improves these results dramatically by leveraging state-of-the-art alias analyses [16, 35, 48].

Figure 4 compares the number of loop invariants identified by both LLVM and NOELLE. NOELLE identifies significantly more loop invariants than LLVM because the invariant abstraction of NOELLE is built using the PDG abstraction. This makes the invariant detection algorithm within NOELLE (Algorithm 2) smaller, more elegant, and more powerful compared to the LLVM one (Algorithm 1).

Finally, we computed the number of loop induction variables that govern a loop using both LLVM and NOELLE. We did so for the three benchmark suites for a total of 41 benchmarks. LLVM identifies only a few loop induction variables (11 total) among all loops for the 41 benchmarks. The reason is that LLVM’s induction variable analysis expects the input IR to have loops in the do-while shape. However, most loops in the 41 benchmarks have a while shape, and changing them into a do-while shape would reduce the applicability of all the implemented parallelization techniques. Instead, NOELLE identifies many loop induction variables (385 total) independently of the shape of the loop being analyzed.

### 4.4 Parallelizing Transformations Upon NOELLE

Next, we describe the parallelizing code transformations built upon NOELLE (HELIX, DSWP, DOALL) that do not rely on speculative techniques. This allows us to compare few-hundred lines of code implementations built upon NOELLE with the parallelizing transformations implemented by `icc` (Intel) and `gcc` (GNU) compilers.

Figure 5 shows the speedups we obtained in PARSEC and MiBench benchmark suites. The few missing benchmarks have failed to compile with the unmodified `clang` compiler, and therefore we cannot use them to test NOELLE-based tools. Figure 5 shows that the NOELLE-based small custom tools already extract more parallelism compared to what `gcc`

and `icc` extract. Furthermore, we analyzed the few benchmarks that NOELLE-based parallelizing tools could not extract significant performance benefits (e.g., `crc`). We found this is due to the lack of support for memory object cloning. This is arguably an abstraction that should exist in the parallelization techniques rather than within NOELLE as the latter is not specialized for parallelization purposes.

We also run these five parallelizing tools on 14 SPEC CPU2017 benchmarks (the only missing benchmark is `gcc`, which did not compile with `clang`). Speedups were obtained only by NOELLE-based parallelizing tools and are within 1% and 5% for these 14 benchmarks demonstrating the robustness of NOELLE abstractions. Speculative techniques are likely to be required to unlock further speedups on these benchmarks. We argue that speculative techniques should be implemented outside NOELLE as they are specific to the parallelization goal.

Finally, we have ported a state-of-the-art parallelizing compiler (Perspective [15]) together with the authors. We modified the original codebase to use the PDG and the aSCC DAG abstractions. This new version has preserved the performance shown in the authors’ original paper.


### 4.5 Reducing Binary Size with NOELLE

Binary size is an important optimization goal for both embedded systems and servers [18]. The compiler `clang` offers an optimization level specialized for this goal (`-Oz`). Dead-FunctionElimination further reduces the binary size by 6.3% on average among the 41 benchmarks considered.


## 5 Related Work

**Providing High-level Abstractions** Researchers have explored bringing high-level abstractions to compilers in many different ways. A few compilers that support automatic parallelization, including Polaris [20], a parallelizing compiler for Fortran programs, Cetus [3], a C compiler focusing on multicore, ROSE [10], a compiler for building custom compilation tools, operate on high-level abstractions and perform source-to-source translation, and thus miss opportunities presented only in low-level IRs including more fine-grained operations and more canonical code patterns.

Many domain-specific projects add new abstractions similar to NOELLE. SeaHorn [33] provides new abstractions for developing new verification techniques. Polly [8,



**Figure 4.** NOELLE detects significantly more invariants than LLVM even if the former relies on a simpler and shorter algorithm powered by higher-level abstraction (Algorithm 2) compared to LLVM (Algorithm 1).



**Figure 5.** Both gcc and icc did not obtain additional performance benefits from their parallelization techniques. Instead, NOELLE-based parallelizing tools generate additional benefits compared to their baseline, clang.

32], PLUTO [21], HALIDE [5, 45], Tiramisu [14, 19], and APOLLO [22] provide abstractions to suit polyhedral optimizations, which target loops characterized by regular control and data flows. TensorFlow [12], a widely used machine learning framework, uses high-level graph representations that allow graph optimizations more discoverable [13]. These projects focus on specific domains and their abstractions are not easily reusable for problems outside their domains.

Few domain-independent compilers combine low-level IR with high-level abstractions like NOELLE. SUIF compiler [11] provides low-level IR as well as higher-level constructs including loops, conditional statements, and array accesses [51]. The IMPACT compiler [25], which provides hierarchical IRs to enable optimizations at different levels. Unfortunately, they are not maintained anymore.

The LLVM community also has a Loop Optimization Working Group [39] that recently has started working on a few abstractions included in NOELLE, such as dependence graph. We plan to share NOELLE code with them. We also see value in maintaining NOELLE as a separate project that focuses mainly on performance rather than making a balance between performance, code size, and compilation time.

**LLVM Projects** As we have built NOELLE on top of LLVM, we want to know how NOELLE would do. To do this, we have exhaustively reviewed all 544 papers published in PLDI, CGO, and CC during the past five years (2016-2020). Out of these papers, 87 papers explicitly mention that they are built on top of LLVM by either implementing new passes, modifying the LLVM internals, or creating a new front-end/back-end based on LLVM IR. Out of these 87 papers,

- 26 (29.9%) use abstractions similar to the ones provided by NOELLE. Thus, they can potentially be re-implemented on top of NOELLE with significantly fewer lines of code and/or with better performance. We have implemented CARAT [46] and PRVJeeves [38] in NOELLE and presented the benefit in 3. Other examples include Spinal Node [36], which uses PDG as well as data flow analysis; Valence [52], which relies on call graph analysis; Clairvoyance [50], which relies on loop-carried dependence analysis.
- 10 (11.5%) provide new abstractions or implement analyses or transformations that fulfill NOELLE abstractions. We have already integrated SVF [47] and SCAF [16] within NOELLE. We plan to evaluate other examples [27, 40, 41, 44] in the future.
- 25 (28.7%) are doing tasks orthogonal to NOELLE’s abstractions. Nevertheless, they do not conflict with NOELLE because both implementations do not modify LLVM internals. Due to NOELLE’s modular and demand-driven design, future work can use NOELLE even if only a subset of abstractions are of interest.
- 26 (29.9%) paper modify LLVM internals or use alternative front-end/back-end. These projects need to be analyzed case by case for the possibility of integration with NOELLE.

In conclusion, 41.4% of the projects are highly likely to benefit from or contribute to NOELLE’s abstractions; 28.7% have the potential for future collaboration; 29.9% need investigation before integration.

## 6 Conclusion

Code analyses and transformations need to go beyond peephole and ILP optimizations for modern architectures. Their implementation requires high-level abstractions that are currently lacking in LLVM. This paper introduces NOELLE, an open-source compilation layer built upon LLVM that provides the required abstractions. NOELLE has been tested with ten highly diverse and complex tools that are built upon it. All these tools gain benefits compared to unmodified LLVM while dramatically reducing their LoC.

## References

- [1] [n. d.]. Automatic Parallelization. <https://www.intel.com/content/www/us/en/develop/documentation/cpp-compiler-developer-guide-and-reference/top/optimization-and-programming-guide/automatic-parallelization.html>.
- [2] [n. d.]. AutoParInGCC - GCC Wiki. <https://gcc.gnu.org/wiki/AutoParInGCC>.
- [3] [n. d.]. The Cetus Project. <https://engineering.purdue.edu/Cetus/>.
- [4] [n. d.]. GCC, the GNU Compiler Collection - GNU Project - Free Software Foundation (FSF). <https://gcc.gnu.org/>.
- [5] [n. d.]. Halide. <https://halide-lang.org/>.
- [6] [n. d.]. The LLVM Compiler Infrastructure Project. <http://llvm.org/>.
- [7] [n. d.]. MSVC Auto-parallelization. <https://docs.microsoft.com/en-us/cpp/parallel/auto-parallelization-and-auto-vectorization?view=vs-2019>.
- [8] [n. d.]. Polly - Polyhedral Optimizations for LLVM. <https://polly.llvm.org/>.
- [9] [n. d.]. Portable Hardware Locality (hwloc). <https://www.open-mpi.org/projects/hwloc>.
- [10] [n. d.]. Rose Compiler – Program Analysis and Transformation. <http://rosecompiler.org/>.
- [11] [n. d.]. The SUIF Compiler - SUIF 2. <https://suif.stanford.edu/suif/suif2/>.
- [12] [n. d.]. TensorFlow. <https://www.tensorflow.org/>.
- [13] [n. d.]. TensorFlow Graph Optimization with Grappler | TensorFlow Core. [https://www.tensorflow.org/guide/graph\\_optimization](https://www.tensorflow.org/guide/graph_optimization).
- [14] [n. d.]. Tiramisu Compiler. <http://tiramisu-compiler.org/>.
- [15] Sotiris Apostolakis, Ziyang Xu, Greg Chan, Simone Campanoni, and David I. August. 2020. Perspective: A Sensible Approach to Speculative Automatic Parallelization. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '20)*. Association for Computing Machinery, New York, NY, USA, 351–367. <https://doi.org/10.1145/3373376.3378458>
- [16] Sotiris Apostolakis, Ziyang Xu, Zujun Tan, Greg Chan, Simone Campanoni, and David I. August. 2020. SCAF: A Speculation-Aware Collaborative Dependence Analysis Framework. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2020)*. Association for Computing Machinery, New York, NY, USA, 638–654. <https://doi.org/10.1145/3385412.3386028>
- [17] Andrew W Appel. 2008. *Modern compiler implementation in Java*. Cambridge university press.
- [18] Grant Ayers, Nayana Prasad Nagendra, David I. August, Hyoun Kyu Cho, Svilen Kanev, Christos Kozyrakis, Trivikram Krishnamurthy, Heiner Litz, Tipp Moseley, and Parthasarathy Ranganathan. 2019. AsmDB: Understanding and Mitigating Front-End Stalls in Warehouse-Scale Computers. In *Proceedings of the 46th International Symposium on Computer Architecture (ISCA '19)*. Association for Computing Machinery, New York, NY, USA, 462–473. <https://doi.org/10.1145/3307650.3322234>
- [19] Riyadh Baghdadi, Jessica Ray, Malek Ben Romdhane, Emanuele Del Sozzo, Abdurrahman Akkas, Yunming Zhang, Patricia Suriana, Shoaib Kamil, and Saman P. Amarasinghe. 2019. Tiramisu: A Polyhedral Compiler for Expressing Fast and Portable Code. In *IEEE/ACM International Symposium on Code Generation and Optimization, CGO 2019, Washington, DC, USA, February 16-20, 2019*, Mahmut Taylan Kandemir, Alexandra Jimborean, and Tipp Moseley (Eds.). IEEE, 193–205. <https://doi.org/10.1109/CGO.2019.8661197>
- [20] Bill Blume, Rudolf Eigenmann, Keith Faigin, and John Grout. [n. d.]. *Polaris: The Next Generation in Parallelizing Compilers*. Technical Report. 18 pages.
- [21] Uday Bondhugula and Jagannathan Ramanujam. 2007. Pluto: A practical and fully automatic polyhedral parallelizer and locality optimizer. (2007).
- [22] Juan Manuel Martínez Caamaño, Aravind Sukumaran-Rajam, Artiom Baloian, Manuel Selva, and Philippe Clauss. 2017. APOLLO: Automatic Speculative POLyhedral Loop Optimizer. In *IMPACT 2017 - 7th International Workshop on Polyhedral Compilation Techniques*. 8.
- [23] Simone Campanoni, Timothy Jones, Glenn Holloway, Vijay Janapa Reddi, Gu-Yeon Wei, and David Brooks. 2012. HELIX: Automatic Parallelization of Irregular Programs for Chip Multiprocessing. In *Proceedings of the Tenth International Symposium on Code Generation and Optimization (CGO '12)*. ACM, New York, NY, USA, 84–93. <https://doi.org/10.1145/2259016.2259028>
- [24] Simone Campanoni, Timothy Jones, Glenn Holloway, Gu. Y. Wei, and David Brooks. 2012. The HELIX project: Overview and directions. In *DAC Design Automation Conference 2012*. 277–282. <https://doi.org/10.1145/2228360.2228412>
- [25] Pohua P Chong, Scott A Mohike, and Nancy J Warier. [n. d.]. IMPACT: An Architectural Framework for Multiple-Instruction-Issue Processors. ([n. d.]), 10.
- [26] Enrico A. Deiana, Vincent St-Amour, Peter A. Dinda, Nikos Hardavellas, and Simone Campanoni. 2018. Unconventional Parallelization of Nondeterministic Applications. In *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '18)*. ACM, New York, NY, USA, 432–447. <https://doi.org/10.1145/3173162.3173181>
- [27] Johannes Doerfert, Tobias Grosser, and Sebastian Hack. 2017. Optimistic Loop Optimization. In *Proceedings of the 2017 International Symposium on Code Generation and Optimization, CGO 2017, Austin, TX, USA, February 4-8, 2017*, Vijay Janapa Reddi, Aaron Smith, and Lingjia Tang (Eds.). ACM, 292–304.
- [28] Yuanbo Fan, Simone Campanoni, and Russ Joseph. 2019. Time squeezing for tiny devices. In *Proceedings of the 46th International Symposium on Computer Architecture, ISCA 2019, Phoenix, AZ, USA, June 22-26, 2019*. 657–670. <https://doi.org/10.1145/3307650.3322268>
- [29] Yuanbo Fan, Tianyu Jia, Jie Gu, Simone Campanoni, and Russ Joseph. 2018. Compiler-guided Instruction-level Clock Scheduling for Timing Speculative Processors. In *Proceedings of the 55th Annual Design Automation Conference (DAC '18)*. ACM, New York, NY, USA, Article 40, 6 pages. <https://doi.org/10.1145/3195970.3196013>
- [30] Jeanne Ferrante, Karl J Ottenstein, and Joe D Warren. 1987. The program dependence graph and its use in optimization. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 9, 3 (1987), 319–349.
- [31] Souradip Ghosh, Michael Cuevas, Simone Campanoni, and Peter Dinda. 2020. Compiler-based Timing For Extremely Fine-grain Preemptive Parallelism. In *Super Computing conference (SC)*.
- [32] Tobias Grosser, Hongbin Zheng, Raghesh Aloor, Andreas Simbürger, Armin Gröbinger, and Louis-Noël Pouchet. 2011. Polly-Polyhedral optimization in LLVM. In *Proceedings of the First International Workshop on Polyhedral Compilation Techniques (IMPACT)*, Vol. 2011. 1.
- [33] Arie Gurfinkel, Temesghen Kahsai, Anvesh Komuravelli, and Jorge A. Navas. 2015. The SeaHorn Verification Framework. In *Computer Aided Verification*, Daniel Kroening and Corina S. Păsăreanu (Eds.). Vol. 9206. Springer International Publishing, Cham, 343–361. [https://doi.org/10.1007/978-3-319-21690-4\\_20](https://doi.org/10.1007/978-3-319-21690-4_20)
- [34] Ali R Hurson, Joford T Lim, Krishna M Kavi, and Ben Lee. 1997. Parallelization of doall and doacross loops—a survey. In *Advances in*

- computers. Vol. 45. Elsevier, 53–103.
- [35] Nick P. Johnson, Jordan Fix, Stephen R. Beard, Taewook Oh, Thomas B. Jablin, and David I. August. 2017. A Collaborative Dependence Analysis Framework. In *Proceedings of the 2017 International Symposium on Code Generation and Optimization (CGO '17)*. IEEE Press, Piscataway, NJ, USA, 148–159. <http://dl.acm.org/citation.cfm?id=3049832.3049849>
- [36] Bongjun Kim, Seonyeong Heo, Gyeongmin Lee, Seungbin Song, Jong Kim, and Hanjun Kim. 2019. Spinal Code: Automatic Code Extraction for near-User Computation in Fogs. In *Proceedings of the 28th International Conference on Compiler Construction - CC 2019*. ACM Press, Washington, DC, USA, 87–98. <https://doi.org/10.1145/3302516.3307356>
- [37] Chris Latner and Vikram Adve. 2004. LLVM: A compilation framework for lifelong program analysis & transformation. In *Proceedings of the international symposium on Code generation and optimization: feedback-directed and runtime optimization*. IEEE Computer Society, 75.
- [38] Michael Leonard and Simone Campanoni. 2020. Introducing the Pseudorandom Value Generator Selection in the Compilation Toolchain. In *Proceedings of the 18th ACM/IEEE International Symposium on Code Generation and Optimization (CGO 2020)*. Association for Computing Machinery, New York, NY, USA, 256–267. <https://doi.org/10.1145/3368826.3377906>
- [39] llvm. [n. d.]. The Loop Optimization Working Group. <https://llvm.org/devmtg/2019-10/talk-abstracts.html#pan2>.
- [40] Maroua Maalej, Vitor Paisante, Pedro Ramos, Laure Gonnord, and Fernando Magno Quintão Pereira. 2017. Pointer Disambiguation via Strict Inequalities. In *Proceedings of the 2017 International Symposium on Code Generation and Optimization, CGO 2017, Austin, TX, USA, February 4-8, 2017*, Vijay Janapa Reddi, Aaron Smith, and Lingjia Tang (Eds.). ACM, 134–147.
- [41] Stanislav Manilov, Christos Vasiladiotis, and Björn Franke. 2018. Generalized Profile-Guided Iterator Recognition. In *Proceedings of the 27th International Conference on Compiler Construction - CC 2018*. ACM Press, Vienna, Austria, 185–195. <https://doi.org/10.1145/3178372.3179511>
- [42] Niall Murphy, Timothy Jones, Robert Mullins, and Simone Campanoni. 2016. Performance Implications of Transient Loop-carried Data Dependencies in Automatically Parallelized Loops. In *Proceedings of the 25th International Conference on Compiler Construction (CC 2016)*. ACM, New York, NY, USA, 23–33. <https://doi.org/10.1145/2892208.2892214>
- [43] G. Ottoni, R. Rangan, A. Stoler, and D. I. August. 2005. Automatic thread extraction with decoupled software pipelining. In *38th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO'05)*, 12 pp.–118. <https://doi.org/10.1109/MICRO.2005.13>
- [44] Ankush Phulia, Vaibhav Bhagee, and Sorav Bansal. 2020. OOEla: Order-of-Evaluation Based Alias Analysis for Compiler Optimization. In *Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2020, London, UK, June 15-20, 2020*, Alastair F. Donaldson and Emina Torlak (Eds.). ACM, 839–853. <https://doi.org/10.1145/3385412.3385962>
- [45] Jonathan Ragan-Kelley, Connelly Barnes, Andrew Adams, Sylvain Paris, Frédo Durand, and Saman Amarasinghe. 2013. Halide: A Language and Compiler for Optimizing Parallelism, Locality, and Recomputation in Image Processing Pipelines. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '13)*. Association for Computing Machinery, New York, NY, USA, 519–530. <https://doi.org/10.1145/2491956.2462176>
- [46] Brian Suchy, Simone Campanoni, Nikos Hardavellas, and Peter Dinda. 2020. CARAT: A Case for Virtual Memory through Compiler- and Runtime-Based Address Translation. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2020)*. Association for Computing Machinery, New York, NY, USA, 329–345. <https://doi.org/10.1145/3385412.3385987>
- [47] Yulei Sui and Jingling Xue. 2016. SVF: Interprocedural Static Value-Flow Analysis in LLVM. In *Proceedings of the 25th International Conference on Compiler Construction, CC 2016, Barcelona, Spain, March 12-18, 2016*, Ayal Zaks and Manuel V. Hermenegildo (Eds.). ACM, 265–266. <https://doi.org/10.1145/2892208.2892235>
- [48] Yulei Sui and Jingling Xue. 2016. SVF: interprocedural static value-flow analysis in LLVM. In *Proceedings of the 25th international conference on compiler construction*. 265–266.
- [49] Robert Tarjan. 1972. Depth-first search and linear graph algorithms. *SIAM journal on computing* 1, 2 (1972), 146–160.
- [50] Kim-Anh Tran, Trevor E. Carlson, Konstantinos Koukos, Magnus Sjölander, Vasileios Spiliopoulos, Stefanos Kaxiras, and Alexandra Jimborean. 2017. Clairvoyance: Look-Ahead Compile-Time Scheduling. In *Proceedings of the 2017 International Symposium on Code Generation and Optimization, CGO 2017, Austin, TX, USA, February 4-8, 2017*, Vijay Janapa Reddi, Aaron Smith, and Lingjia Tang (Eds.). ACM, 171–184.
- [51] Robert P Wilson, Robert S French, Christopher S Wilson, Saman P Amarasinghe, Jennifer M Anderson, Steve W K Tjiang, Shih-Wei Liao, Chau-Wen Tseng, Mary W Hall, Monica S Lam, and John L Hennessy. [n. d.]. *The SUIF Compiler System: A Parallelizing and Optimizing Research Compiler*. Technical Report. 7 pages.
- [52] Tong Zhou, Michael R. Jantz, Prasad A. Kulkarni, Kshitij A. Doshi, and Vivek Sarkar. 2019. Valence: Variable Length Calling Context Encoding. In *Proceedings of the 28th International Conference on Compiler Construction - CC 2019*. ACM Press, Washington, DC, USA, 147–158. <https://doi.org/10.1145/3302516.3307351>