

# LLOV: A Fast Static Data-Race Checker for OpenMP Programs

UTPAL BORA, SANTANU DAS, PANKAJ KUKREJA, SAURABH JOSHI, and RAMAKRISHNA UPADRASTA, IIT Hyderabad, India  
SANJAY RAJOPADHYE, Colorado State University, Fort Collins CO, USA

In the era of Exascale computing, writing efficient parallel programs is indispensable and at the same time, writing sound parallel programs is highly difficult. While parallel programming is easier with frameworks such as OpenMP, the possibility of data races in these programs still persists. In this paper, we propose a fast, lightweight, language agnostic, and static data race checker for OpenMP programs based on the LLVM compiler framework. We compare our tool with other state-of-the-art data race checkers on a variety of well-established benchmarks. We show that the precision, accuracy, and the F1 score of our tool is comparable to other checkers while being orders of magnitude faster. To the best of our knowledge, this work is the only tool among the state-of-the-art data race checkers that can verify a FORTRAN program to be data race free.

Additional Key Words and Phrases: OpenMP, Polyhedral Compilation, Program Verification, Data Race Detection, Static Analysis

## ACM Reference Format:

Utpal Bora, Santanu Das, Pankaj Kukreja, Saurabh Joshi, Ramakrishna Upadrista, and Sanjay Rajopadhye. 2019. LLOV: A Fast Static Data-Race Checker for OpenMP Programs. 1, 1 (December 2019), 22 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

The benefits of heterogeneous parallel programming in obtaining high performance from modern complex hardware architectures are indisputable among the scientific community. Although indispensable for its efficiency, parallel programming is prone to errors. This is crucial since programming errors could result in significant monetary losses or prove to be a risk factor where human safety systems are involved. Historical incidents such as Therac-25 accidents [39], the Ariane 5 flight 501 failure [43], EDS Child Support IT failure, and Knight Capital Group trading glitch [53] have been directly attributed to software errors and testify to the need for bug detection mechanisms. The detection of errors a priori, therefore, could significantly reduce this risk and make programs more robust and dependable.

In this work, we propose a solution to the problem of statically detecting errors (data races) in OpenMP parallel programs. We developed a data race detection tool based on LLVM/Clang/Flang [29, 32, 44] that is amenable to various languages, such as C/C++/Objective-C as well as FORTRAN. To the best of our knowledge, our work is the *first static OpenMP data race detection tool based on language independent intermediate representation of LLVM* (henceforth called LLVM-IR) [30].

**Contributions:** In this paper, we make the following contributions:

- Implementation of a *fast, static, and language agnostic* OpenMP data race checker in the LLVM framework based on its intermediate representation (LLVM-IR) using the polyhedral framework Polly [27]. Our tool can also certify that a program is *data race free* along with detecting race conditions in OpenMP programs.

---

Authors' addresses: Utpal Bora; Santanu Das; Pankaj Kukreja; Saurabh Joshi; Ramakrishna Upadrista, IIT Hyderabad, India; Sanjay Rajopadhye, Colorado State University, Fort Collins CO, USA.

---

2019. XXXX-XXXX/2019/12-ART \$15.00  
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

- We create a benchmark suite in FORTRAN, named DataRaceBench FORTRAN, which is a manifestation of DataRaceBench v1.2 [40] in FORTRAN. DataRaceBench v1.2 is a benchmark suite consisting of programs written in C/C++ using OpenMP kernels. We make DataRaceBench FORTRAN open source [37]. For all the tools that can analyze FORTRAN programs for data races, this benchmark suite can allow evaluation of such tools in a standardized fashion.
- We make a comparative study of well known data race checker tools on a standard set of OpenMP benchmarks. We evaluate these tools on various metrics such as precision, recall, accuracy, F1 score, Diagnostic Odds Ratio and running times. We show that LLOV performs quite well on these metrics while completely outperforming its competitors in terms of runtime.

The rest of the paper is divided as follows: we start with the motivation for our work in Section 2, and describe common data race conditions in Section 3. In Section 4, we discuss the verifier implementation details, our proposed algorithm, and list out the advantages of our approach over the existing dynamic tools. Section 5 discusses related work in OpenMP data race detection along with their differences from our approach. Our results and comparison with other verifiers are presented in Section 6 and finally, we conclude in Section 7.

## 2 BACKGROUND AND MOTIVATION

Multithreading support in hardware architectures has been very common in recent times with the number of cores per socket going upto 56 in Intel® Xeon® Platinum 9282 with 2 threads per core and upto 72 in Intel® Xeon Phi™ Processor 7290F (accelerator) with 4 threads per core. The Top500 [61] November 2018 list of supercomputers comprises systems with cores per socket ranging from 6 to 260. As these systems are simultaneous multithreading (SMT) systems, operating systems with support for SMT and/or Symmetric Multi Processing (SMP) can benefit from execution of a large number of threads in parallel. The memory (not cache) is shared among all the threads in a node with uniform memory access (UMA) or non-uniform memory access (NUMA), enabling shared memory multithreading.

In the past, the scientific community wrote parallel programs in C/C++ and FORTRAN using either language extensions or with APIs to run them across different nodes in a cluster or grid. With the advent of multi-core processors, the focus shifted to a shared memory programming model, e.g., the pthreads library/run-time system [33] coupled with a vanilla language like C/C++, or a parallel language like coarray-FORTRAN or HPF.

Later on, languages having structured parallelism such as Cilk [9], Julia [7], Chappel [12, 35], X10 [13], and others started gaining popularity in the community. However, the community has extensively adopted structured parallel programming frameworks, such as OpenMP [17, 49], MPI [26], OpenACC [51], and OpenCL [28] because of easy migration from legacy sequential code to its parallel variant. The availability of efficient runtimes and versatile support for various architectures played a major role in popularizing these frameworks. Amongst these, in this work we focus on the OpenMP parallel programming framework.

The OpenMP programming paradigm [17, 49] introduced structured parallelism in C/C++ and FORTRAN. It supports *Single Program Multiple Data* (SMPD) programming model with multiple threads, *Single Instruction Multiple Data* (SIMD) programming model within a single thread in CPUs with SIMD hardware extension, as well as SIMD among threads of a *thread block* in GPUs. OpenMP enables divide-and-conquer paradigm with tasks and nested parallelism, provides a data environment for shared memory consistency, supports mutual exclusion and atomicity, and synchronization amongst threads.

However, inappropriate usage of the OpenMP model might introduce bugs into an application. Multithreaded applications running on a shared memory SMP system can suffer from access anomaly due to data dependences, where two threads access the same memory location, and at least one of the accesses is a write operation.

*Definition 2.1 (Data Race).* An execution of a concurrent program is said to have a *data race* when two different threads access the same memory location, these accesses are not protected by a mutual exclusion mechanism (e.g., locks), the order of the two accesses is non-deterministic and one of these accesses is a write.

*Definition 2.2 (Team of threads).* A set of OpenMP threads comprising a master thread and an optional group of sibling threads that participate in the potential execution of an OpenMP parallel region is called a *team*. The master thread is assigned thread id zero.

Though compilers do ensure that OpenMP constructs conform to the syntactic and semantic specifications [50], *none of the mainstream compilers*, such as GCC [56], LLVM [32], ICC [69], and PGI [52], provide *built-in data race detection* support. There exist dynamic tools to detect race conditions, but they either take a very long time to report races, or might miss some race conditions. This is because these tools are dependent on the execution schedule of the threads and the program parameters. The primary goal of our work is to provide a built-in data race detector for OpenMP parallel programs in the LLVM toolchain, using static analysis technique as discussed in Section 4.

### 3 COMMON RACE CONDITIONS IN OPENMP PROGRAMS

In this section, we will walk through, with examples, different race conditions frequently encountered in OpenMP programs.

#### 3.1 Missing data sharing clauses

Listing 1 shows an OpenMP worksharing construct `omp parallel for` with a data race. This program computes the sum of the squares of all the elements in the matrix `u`.

By default, OpenMP data environment considers variables as *shared* among all the threads in a team. In this example, variables `temp`, `i`, and `j` are marked as *private*, indicating that each thread will have its own copy of these variables. However, the variable `sum` (Line 5) of Listing 1 is not listed by any of the data sharing clauses. Therefore, the variable `sum` will be *shared* among all the threads in the team. Thus, each thread will work on the same shared variable and update it simultaneously without any synchronization, leading to a data race.

Listing 2 presents a program in FORTRAN with a data race due to a missing *private* clause corresponding to the variable `tmp` (Line 3). Due to such intricacies, a programmer is prone to make mistakes and inadvertently introduce data races in the program.

Our aim is to develop techniques and a tool that understand the semantics of OpenMP pragmas and clauses with all their subtleties.

```

1 #pragma omp parallel for private (temp,i,j)
2   for (i = 0; i < len; i++)
3     for (j = 0; j < len; j++) {
4       temp = u[i][j];
5       sum = sum + temp * temp;
6     }

```

Listing 1. DRB021: OpenMP Worksharing construct with data race

```

1 !SOMP PARALLEL DO
2   do i = 0, len - 1
3     tmp = a(i) + i
4     a(i) = tmp
5   end do
6 !SOMP end PARALLEL do

```

Listing 2. DRBF028: FORTRAN code with data race because of missing *private* clause

### 3.2 Loop carried dependences

OpenMP programs may suffer from race conditions due to incorrect parallelization strategies. Such race conditions may occur because of parallelization of loops with loop carried dependences.

For example, the loop nest in Listing 3 is parallel in the outer dimension (Line 1), but it is the inner loop (Line 3) that is marked parallel, which has a loop carried dependence because the read of  $b[i][j-1]$  (Line 4) is dependent on write to  $b[i][j]$  (Line 4) in the previous iteration. A correct parallelization strategy for this example would be to mark the outer loop as parallel in place of the inner loop.

As another example, Listing 4 is a parallel implementation of Floyd-Warshall's shortest path algorithm with a *benign race condition* [4]. This is because all the iterations of the (parallel) inner  $j$  loop read  $A[i][k]$  including the one where  $j=k$ , which also writes into  $A[i][k]$ . Hence the iterations before  $j=k$  need the previous value of  $A[i][k]$  and subsequent ones need the new, updated value. However, if all the matrix entries are non-negative,  $A[i][k]$  is unchanged by the assignment.

```

1  for (i=0; i<n; i++) {
2  #pragma omp parallel for
3      for (j=1; j<m; j++) {
4          b[i][j]=b[i][j-1];
5      }
6  }
```

Listing 3. DRB038: Example with Loop Carried Dependence

```

1  for (k = 1; k <= n; k++)
2  #pragma omp parallel for
3  for (i = 1; i <= n; i++)
4      for (j = 1; j <= n; j++)
5          A[i][j] =
6          min(A[i][k] + A[k][j], A[i][j]);
```

Listing 4. Parallel Floyd-Warshall Algorithm with a benign race condition

### 3.3 SIMD races

OpenMP supports SIMD constructs for both CPUs and GPUs. In CPUs, SIMD is supported with vector processing units, where the consecutive iterations of a loop can be executed in a SIMD processing unit within a single core by a single thread. This is contrary to other loop constructs where iterations are shared among different threads.

In the example in Listing 5, the loop is marked as SIMD parallel loop by the `pragma omp simd` (Line 1). This signifies that consecutive iterations assigned to a single thread can be executed concurrently in SIMD units, called vector arithmetic logic units (ALU), within a single core rather than executing sequentially in a scalar ALU. However, because of the forward loop carried dependence between write to  $a[i+1]$  (Line 3) in one iteration and read of  $a[i]$  (Line 3) in the previous iteration, concurrent execution of the consecutive iterations in a vector ALU will produce inconsistent results. Dynamic data race detection tools *fail to* detect race conditions in such cases as the execution happens within a single thread.

```

1 #pragma omp simd
2 for (int i=0; i<len-1; i++){
3     a[i+1] = a[i] + b[i];
4 }
```

Listing 5. DRB024: Example with SIMD data race

```

1 #pragma omp parallel shared(b, error) {
2 #pragma omp for nowait
3     for(i = 0; i < len; i++)
4         a[i] = b + a[i]*5;
5 #pragma omp single
6     error = a[9] + 1;
7 }
```

Listing 6. DRB013: Example with data race due to improper synchronization

### 3.4 Synchronization issues

Improper synchronization between threads is a common cause of data race conditions in concurrent programs. OpenMP can have both explicit and implicit synchronizations associated with different constructs.

The constructs `parallel`, `for`, `workshare`, `sections`, and `single` have an implicit barrier at the end of the construct. This ensures that all threads in the team wait for others to proceed further. This enforcement can be overcome with the `nowait` clause where threads in a team, after completion of the construct, are no longer bound to wait for the other unfinished threads. However, improper use of the `nowait` clause can result in data races as shown in Listing 6. In this example, a thread executing the `parallel for` (Line 3) will not wait for the other threads in the team because of the `nowait` clause (Line 2). Threads that have finished executing the `for` loop are free to continue and execute the `single` construct (Line 5). Since there is a data dependence between write to `a[i]` (Line 4) and read of `a[9]` (Line 6), it might result in a data race.

Such cases are extremely difficult to reproduce as they are dependent on the order of execution of the threads. This particular order of execution may not manifest during runtime, therefore, making it hard to detect such cases for dynamic analysis tools. Static analysis techniques have an advantage over the dynamic techniques in detecting race conditions for such cases.

### 3.5 Control flow dependent on number of threads

Control flow dependent on number of threads available at runtime might introduce race conditions in a parallel program. In the example in Listing 7, race condition will arise only when thread IDs of two or more threads in the team are multiples of 2.

```

1 #pragma omp parallel
2   if (omp_get_thread_num() % 2 == 0) {
3     Flag = true;
4   }

```

Listing 7. Control flow dependent on number of threads

## 4 IMPLEMENTATION DETAILS

We now describe the architecture, the implementation and the algorithm of our tool. LLOV is built on top of LLVM-IR and can analyze OpenMP programs written in C/C++, FORTRAN, and any programming language that has a stable LLVM frontend. LLVM-IR can be generated from C/C++ programs using the Clang [29] frontend, and from FORTRAN programs using the Flang [44] frontend. The architecture of LLOV is shown in Figure 1.

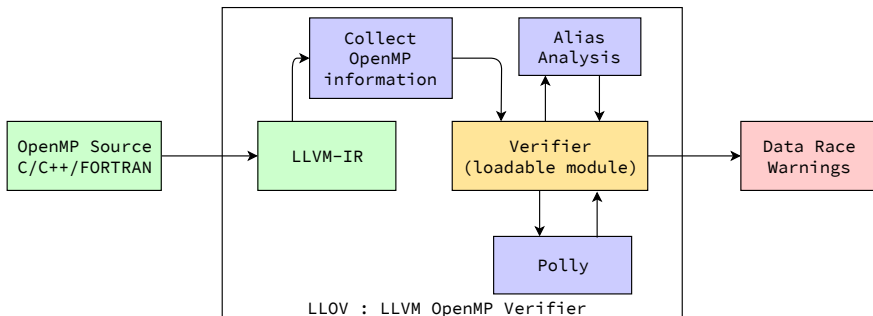


Fig. 1. Flow Diagram of LLVM OpenMP Verifier (LLOV)

## 4.1 LLOV architecture

LLOV has two phases: analysis and verification.

**4.1.1 Analysis phase.** In the first phase, we analyze the LLVM-IR to collect various OpenMP pragmas and additional information required for race detection. This analysis is necessary because OpenMP constructs are lowered to the IR by compiler frontends such as Clang [29] and Flang [44]. LLVM-IR [30] is sequential and does not have support for parallel constructs. Hence, parallel constructs in high level languages are represented in the IR as function calls. OpenMP pragmas are translated as function calls to the APIs of the OpenMP runtime library *libomp*.

As part of this analysis, for each OpenMP construct, we collect information such as memory locations, access types, storage modifiers, synchronization details, scheduling type, etc. An *in-memory representation* of the OpenMP directive contains a directive type, a schedule (if present), variable names and types, and a list of child directives for nested OpenMP constructs. An illustration of our representation is shown in Listing 8 for the example in Listing 6 (Section 3). The grammar for this representation is presented in BNF form in Listing 9.

```
Directive: OMP_Parallel
Variables:
  Private:  %.omp.ub = alloca i32, align 4
  Private:  %.omp.lb = alloca i32, align 4
  Shared:   i32* %i
  Firstprivate: i32* %len
  Firstprivate: i64 %vla
  Shared:   i32* %a
  Shared:   i32* %b
  Shared:   i32* %error
  Private:  %.omp.stride = alloca i32, align 4
  Private:  %.omp.is_last = alloca i32, align 4
Child Directives:
1: Directive: OMP_Workshare_Loop
   Schedule type : Static Schedule (chunksize=1)
2: Directive: OMP_Workshare_single
3: Directive: OMP_Barrier
```

Listing 8. In-memory representation of a directive

```
<Directive> ::= <Dtype> [ Sched ]
              { <Var> } { <Directive> }
<Dtype> ::= parallel | for | simd
          | workshare | single
          | master | critical
<Var> ::= <Vtype> val
<Vtype> ::= private | firstprivate
          | shared | lastprivate
          | reduction | threadprivate
<Sched> ::= [ <modifier> ]
           [ ordered ] <Stype>
           <chunk>
<modifier> ::= monotonic
             | nonmonotonic
<Stype> ::= static | dynamic
          | guided | auto | runtime
<chunk> ::= positive-int-const
```

Listing 9. BNF of in-memory representation of directives

**4.1.2 VERIFICATION PHASE.** LLOV checks for data races only in regions of a program marked parallel by one of the structured parallelism constructs of OpenMP listed in Table 1. A crucial property of OpenMP constructs is that the specification [50] allows only structured blocks within a pragma. A structured block must not contain arbitrary jumps into or out of it. In other words, a structured block closely resembles a Single Entry Single Exit (SESE) region [1] used in loop analyses. The polyhedral framework [27] can precisely model such an SESE region for powerful analyses and complex transformations.

The polyhedral framework is based on *exact dependence analysis* [24], by which the exact dependence information can be expressed as (piecewise, pseudo-) affine functions. The framework is more powerful than any approximations returned by conventional dependence testing algorithms. Polly [27, 45] is the polyhedral compilation engine in LLVM framework. Polly relies on the Integer Set Library (ISL) [64] to perform exact dependence analysis, and also provides transformations like loop tiling, loop fusion, and outer loop vectorization. Dependences are modelled as ISL relations and transformations are performed on integer sets using ISL operations. The input to Polly is serial code that is marked as Static Control Part (SCoP) and the output is tiled or parallel code enabling vectorization.

Table 1. Comparison of OpenMP pragma handling by OpenMP aware tools. (Y for Yes, N for No)

OpenMP Pragma	LLOV	POLYOMP	DRACO	SWORD	OpenMP Pragma	LLOV	POLYOMP	DRACO	SWORD
#pragma omp parallel	Y	Y	Y	Y	#pragma omp critical	Y	N	N	Y
#pragma omp for	Y	Y	Y	Y	#pragma omp parallel sections	N	N	N	Y
#pragma omp parallel for	Y	Y	Y	Y	#pragma omp sections	N	N	N	Y
#pragma omp atomic	Y	N	N	Y	#pragma omp declare reduction	N	N	N	N
#pragma omp threadprivate	Y	N	N	N	#pragma omp task	N	N	N	N
#pragma omp master	Y	N	N	Y	#pragma omp taskgroup	N	N	N	N
#pragma omp single	Y	N	N	Y	#pragma omp taskloop	N	N	N	N
#pragma omp simd	Y	N	Y	N	#pragma omp taskwait	N	N	N	N
#pragma omp parallel for simd	Y	N	Y	N	#pragma omp teams	N	N	N	N
#pragma omp distribute	Y	N	N	N	#pragma omp barrier	N	N	N	Y
#pragma omp ordered	Y	N	N	N	#pragma omp target map	N	N	N	N

LLOV is built using the Polly infrastructure, but does not use its transformation capabilities. Our primary goal is to perform analyses whereas Polly is designed for complex transformations and optimized code generation. The input to LLOV is *explicitly parallel code* that uses the structured parallelism of OpenMP. With an assumption that the input code has a serial schedule<sup>1</sup>, LLOV calculates its dependence information by using the dependence analyzer of Polly, and then analyzes the parallel constructs using the dependence information to check for data races. Consequently, LLOV can deterministically state whether a program contains data race or is race free.

Polly was designed as an automatic parallelization pass using polyhedral dependence analysis. Its analysis and transformation phases were closely coupled and the analyses were not directly usable from outside Polly, either by other analyses or optimization passes in LLVM. We modified and extended the analysis phase of Polly so that its internal data structures became accessible from LLVM. Other changes involved modifying the dependence analysis of Polly so that its Reduced Dependence Graph (RDG) could be computed *on-demand* for a section of code, thereby reducing its compilation time.<sup>2</sup>

## 4.2 Race detection algorithm

In the analysis phase, we gather information provided by OpenMP’s structured parallelism constructs. We model a section of code, marked as parallel by one of the OpenMP constructs listed in the Table 1, in the polyhedral framework as static affine control parts (SCoPs). This enhances performance as LLOV avoids analyzing unnecessary fragments of the program.

For each SCoP, we query the race detection Algorithm 1 to check for the presence of dependences. A race condition is flagged when the set of the memory accesses in the reduced dependence graph (RDG) and the set of the shared memory accesses within the SCoP is not disjoint. When dependences are absent, the SCoP is parallel and the corresponding program segment is guaranteed to be data race free. Hence, we can verify—fully statically—the *absence of data race* in a program.

<sup>1</sup>This assumption is trivial to prove because of the input C language semantics.

<sup>2</sup>The initial version of the implementation for exposing Polly’s dependence analysis information to LLVM was published in Polly (as part of Google Summer of Code 2016 project) “Polly as an analysis pass in LLVM” [11].

**Algorithm 1: Race Detection Algorithm**


---

```

Input:  $L$ 
Output: result
1 Function isRaceFree( $L$ ):
2    $SCoP = \text{ConstructSCoP}(L)$ ;
3    $RDG = \text{ComputeDependences}(SCoP)$ ;
4    $depth = \text{GetLoopDepth}(L)$ ;
5   if isParallel( $RDG, depth$ ) then
6     | result = "Program is race free.";
7   else
8     | result = "Data Race detected.";
9   return result
10 End Function

```

---

**Algorithm 2: Algorithm to check parallelism**


---

```

Input:  $RDG, dim$ 
Output: True/False
1 Function isParallel( $RDG, dim$ ):
2   if  $RDG$  is Empty then
3     | return True;
4   else
5     Flag = True;
6     while Dependence  $D$  in  $RDG$  do
7       |  $D' = \text{Project Out first } dim \text{ dimensions}$ 
8         | from  $D$ ;
9         | if  $D'$  is Empty then
10          | | continue;
11          | else
12            | | Flag = False;
13            | | break;
14          | return Flag;
15 End Function

```

---

### 4.3 Advantages over dynamic race detection tools

Our static race detection tool LLOV has several advantages over state-of-the-art dynamic race detection tools.

**4.3.1 Detects races in SIMD constructs.** Our tool, LLOV, can detect SIMD races within a single thread, such as those shown in Listing 5. Our algorithm can detect parallelism of the SIMD loop within the loop nest, but even when the loop is not parallel, there is a possibility of data races due to concurrent execution of consecutive iterations in the SIMD units within a single core. Dynamic race detection tools are based on different techniques such as vector clocks, happens-before relations, locksets, monitors, offset-span-labels, etc. and *fail to* detect such race conditions present within a single thread.

**4.3.2 Independent of runtime thread schedule.** Being a static analysis tool, LLOV has the added advantage that the race detection is not dependent on the order of the execution of the threads. This is a major drawback of the dynamic tools because they need to be run multiple times for each specific number of threads to detect races dependent on control flow.

**4.3.3 Independent of input size.** As our analysis can accept parametric variables and we solve the problem with parametric integer programming, there is no limitation on the input size. On the other hand, the dynamic data race detection tools need to be run multiple times for each program parameter to capture races. It is computationally not feasible to cover all the possible input sizes and hence, a dynamic tool can never be complete.

**4.3.4 Independent of number of threads.** Our analysis is not dependent on the number of threads available during runtime. However, all known dynamic tools have to be run multiple times with different numbers of threads to detect races. Hence, dynamic tools might miss out race conditions arising from control flow dependence on the number of threads at runtime.

### 4.4 Limitations of LLOV

LLOV is in active development and in the current version, we tried to cover the frequently used OpenMP v4.5 pragmas. In this version, LLOV does not provide support for the OpenMP constructs



for synchronization, device offloading, and tasking. Since our tool is based on the polyhedral framework, its application is limited to affine programs only. We will explore handling non-affine programs in the future.

## 5 RELATED WORK

There has been extensive work on data race detection in parallel programs; many static, dynamic, and hybrid analyses approaches have been proposed. Mellor-Crummey et. al. [46] proposed race detection in fork-join multithreaded programs using Offset-Span labelling of the nodes. ERASER [58] proposed lockset based approach for race detection. Most earlier tools have focused on pthread based programs, while recent works such as OMPVERIFY [4], ARCHER [2], ROMP [34], SWORD [3], DRACO [66], and POLYOMP [15] are for OpenMP based programs.

In the following subsections, we discuss the state-of-the-art tools and categorize them based on their analyses and their approaches.

### 5.1 Static Tools

There are multiple static data race detectors in the literature. Common techniques used for static analyses are lockset based approach or modeling race condition as a linear programming problem (integer-linear or parametric integer-linear), and appropriately using an ILP or SMT solver to check for its satisfiability. Here, we briefly cover the state-of-the-art static data race detection tools for OpenMP programs. We limit the discussion to OpenMP race detection tools only.

OMPVERIFY [4] is a polyhedral model based static data race detection tool that detects incorrectly specified `omp parallel for` constructs. OMPVERIFY computes the Polyhedral Reduced Dependency Graph (PRDG) to reason about possible violations of true dependences, write-write conflicts, and stale read problems in the OpenMP `parallel for` construct. Although our approach is inspired by OMPVERIFY, we have implemented a different algorithm to detect parallelism of a loop nest. And, while OMPVERIFY can handle only `omp parallel for`, LLOV can handle many more pragmas as listed in Table 3. Moreover, the prototype implementation of OMPVERIFY is in Eclipse CDT/CODAN framework using the AlphaZ [68] polyhedral framework whereas LLOV is based on the widely used LLVM compiler infrastructures. Finally, OMPVERIFY works at the AST level, whereas LLOV works on the language independent LLVM-IR level.

DRACO [66] is a static data race detection tool based on Polyhedral model and is built on the the ROSE compiler framework [57]. One significant advantage of LLOV over DRACO is that LLOV is based on LLVM-IR and is language agnostic, while DRACO is limited only to the C family of languages that can be compiled by the ROSE compiler.

POLYOMP [14] is a static data race detection tool based on polyhedral model. POLYOMP uses an extended polyhedral model to encode OpenMP loop nest information as constraints and uses the Z3 [18] solver to detect race conditions. In contrast, LLOV uses RDG (Reduced Dependence Graph) to determine parallelism in a region and infer race conditions based on the presence of data dependences.

Other static analysis tools like RELAY [65], LOCKSMITH [55] and RACERX [23] use ERASER's [58] lockset algorithm and can detect races in pthread based C/C++ programs. RACERD [8] is a static analysis tool for Java programs.

### 5.2 Dynamic Tools

Various dynamic race techniques have been proposed in the literature. The most well known among them are based on Locksets [58], Happens-before [38] relations, and Offset-Span labels [47].

ARCHER [2] uses both static and dynamic analyses for race detection. It uses happens-before relations [38, 59] which enforces multiple runs of the program to find races. ARCHER reduces the

analysis space of pthread based tool TSAN-LLVM by instrumenting only parallel sections of an OpenMP program. As ARCHER uses shadow memory to keep track of each memory access, memory requirement still remains the problem for memory bound programs.

SWORD [3] is a dynamic tool based on operational semantic rules and uses OpenMP tools framework OMPT [22]. SWORD uses locksets to implement the semantic rules by taking advantage of the events tracked by OMPT. SWORD logs runtime traces for each thread consisting of all the OpenMP events and memory accesses using OMPT APIs. In the second offline phase, it analyzes the traces for concurrent threads using Offset-Span labels and detects unsynchronized memory accesses in two concurrent threads. If no synchronization is used on a common memory access, data race condition is flagged. SWORD cannot detect races in OpenMP SIMD, tasks and target offloading constructs.

ROMP [34] is a dynamic data race detection tool for OpenMP programs. ROMP maintains access history for each memory access. An access history consists of the access event, access type, and any set of locks associated with the access. ROMP constructs a task graph for the implicit and explicit OpenMP tasks and analyzes concurrent events. If an access event is concurrent and the memory is not protected by mutual exclusion mechanisms, ROMP flags data race warnings. ROMP builds upon the offset-span-labels of OpenMP threads and constructs task graph to detect races.

HELGRIND [63] is a dynamic data race detection tool built in Valgrind framework [48] for C/C++ multithreaded programs. HELGRIND maintains *happens before* relations for each pair of memory accesses and forms a directed acyclic graph (DAG). If there is no path from one location to another in the happens-before graph, data race is flagged.

VALGRIND DRD [62] is another dynamic race detection tool in Valgrind. It can detect races in multithreaded C/C++ programs. It is based on happens-before relation similar to HELGRIND.

THREADSANITIZER [59] is a dynamic data race detection tool based on Valgrind for multi-threaded C, C++ programs. It employs a hybrid approach of happens-before annotations and maintains locksets for read and write operations on shared memory. It maintains a state machine as metadata called shadow memory. It reports a race condition when two threads access the same memory location and their corresponding locksets are not disjoint. Because of shadow memory requirement for each memory access, THREADSANITIZER's memory requirement grows linearly with the amount of memory shared among threads. Binary instrumentation also increases the runtimes by around 5x-30x [60].

TSAN-LLVM [60] is based on THREADSANITIZER [59]. TSAN-LLVM uses LLVM to instrument the binaries in place of Valgrind. TSAN-LLVM instrumented binaries incur less runtime overhead compared to THREADSANITIZER. However, it still has similar memory requirements and remains a bottleneck for larger programs.

There are other dynamic analysis tools like ERASER [58], FASTTRACK [25], CORD [36], and RACE-TRACK [67] that use lockset algorithm to detect races in parallel programs. IFRIT [21] is a dynamic race detection tool based on Interference Free Regions(IFR). Since they are not specific to OpenMP, we have not discussed them here.

*OpenMP aware tools:* Majority of the tools are either POSIX thread based or are specific to race detection in inherent parallelism of various programming languages, such as Java, C#, X10, and Chappel. The tools OMPVERIFY [4], ARCHER [2], POLYOMP [15], DRACO [66], SWORD [66], and ROMP [34] are the only ones that exploit the intricate details of structured parallelism in OpenMP.

*Polyhedral model based static analysis tools:* To the best of our knowledge, OMPVERIFY [4], DRACO [66], and POLYOMP [14] are the only tools based on polyhedral framework; meaning that at least theoretically, they are the only ones that use the exact dependence analysis of polyhedral compilation.

LLOV is different from all these tools as it works on the LLVM-IR and collects OpenMP pragmas lowered to library calls. This makes LLOV language independent. Also, the analysis phase of LLOV could be used for any other purpose, like generating task graphs from the LLVM-IR.

## 6 EXPERIMENTAL RESULTS

In this section we describe our experimental setup, provide details on our experiments and compare the results with other tools on a set of benchmarks.

### 6.1 Experimental Setup

Table 2. Race detection tools with the version numbers used for comparison

Tools	Source	Version / Commit
HELGRIND [63]	Valgrind	3.13.0
VALGRIND DRD [62]	Valgrind	3.13.0
TSAN-LLVM [59]	LLVM	6.0.1
ARCHER [2]	git master branch	fc17353
SWORD [3]	git master branch	7a08f3c
ROMP [34]	git master branch	6a0ad6d

We compare LLOV with the state-of-the-art data race detection tools as listed in Table 2. Some of the relevant tools were left out of our experimentation either because of their unavailability or due to their inability to handle OpenMP programs. We had issues setting up ROMP [34] even with the help of the authors of ROMP, hence we could not consider it for our comparison. OMPVERIFY [4] is a prototype implementation in the Eclipse CDT/CODAN framework and can detect races in *omp parallel for* constructs only. Neither binary nor source code for POLYOMP [14] is available in the open. Similarly, DRACO [66] is not available. The authors of DRACO informed us that DRACO is still in the development phase and will be made available only after it is released as a product. We did not consider INTEL INSPECTOR due to its proprietary nature and licensing issues.

For evaluation we chose DataRaceBench v1.2 [41, 42], DataRaceBench FORTRAN [37] (a FORTRAN implementation of DataRaceBench v1.2), and OmpSCR v2.0 [19] benchmark suits. All these benchmarks use OpenMP for parallelization and have known data races. The benchmarks cover OpenMP v4.5 pragmas comprehensively and contain race conditions due to common mistakes in OpenMP programming as listed earlier (Section 3).

DataRaceBench v1.2 [40], a *seeded* OpenMP benchmark with known data race conditions, consists of 116 microbenchmark kernels, out of which 59 kernels have data races, while the remaining 57 kernels do not have any data races. DataRaceBench FORTRAN has the same breakdown. OmpSCR v2.0 is a benchmark suite for high performance computing using OpenMP v3.0 APIs. The benchmark consists of C/C++ and FORTRAN kernels which demonstrate the usefulness and pitfalls of the parallel programming paradigm with both correct and incorrect parallelization strategies. The kernels range from parallelization of simple loops with dependences to more complex parallel implementations of algorithms such as Mandelbrot set generator, Molecular Dynamics simulation, Pi ( $\pi$ ) calculation, LU decomposition, Jacobi solver, fast Fourier transforms (FFT), and Quicksort.

**Performance Metrics Notations.** We define terminology used for performance metrics as follows:

- **True Positive (TP):** If the evaluation tool correctly detects a data race present in the kernel it is a True Positive test result. A higher number of true positives represents a better tool.

- **True Negative (TN):** If the benchmark does not contain a race and the tool declares it as race-free, then it is a true negative case. A higher number of true negatives represents a better tool.
- **False Positives (FP):** If the benchmark does not contain any race, but the tool reports a race condition, it is a false positive. False Positives should be as low as possible.
- **False Negatives (FN):** False Negative test result is obtained when the tool fails to detect a known race in the benchmark. These are the cases that are missed by the tool. A lower number of false negatives are desirable.

We consider the following statistical measures as performance metrics in our experiments.

- **Precision :** Precision is the measure of closeness of the outcomes of prediction. Thus, a higher value of precision represents that the tool will more often than not identify a race condition when it exists.

$$Precision = \frac{TP}{TP + FP}$$

- **Recall :** Recall gives the total number of cases detected out of the maximum data races present. A higher recall value means that there are less chances that a data race is missed by the tool. It is also called true positive rate (TPR).

$$Recall = \frac{TP}{TP + FN}$$

- **Accuracy :** Accuracy gives the chances of correct reports out of all the reports, as the name suggests. A higher value of accuracy is always desired and gives overall measure of the efficacy of the tool.

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN}$$

- **F1 Score :** The harmonic mean of precision and recall is called the F1 score. An F1 score of 1 can be achieved in the best case when both precision and recall are perfect. The worst case F1 score is 0 when either precision or recall is 0.

$$F1\ Score = 2 * \frac{Precision * Recall}{Precision + Recall}$$

- **Diagnostic odds ratio (DOR) :** It is the ratio of the positive likelihood ratio (LR+) to the negative likelihood ratio (LR-).

$$DOR = \frac{LR+}{LR-} \text{ where,}$$

$$\text{Positive Likelihood Ratio (LR+)} = \frac{TPR}{FPR},$$

$$\text{Negative Likelihood Ratio (LR-)} = \frac{FNR}{TNR},$$

$$\text{True Positive Rate (TPR)} = \frac{TP}{TP + FN},$$

$$\text{False Positive Rate (FPR)} = \frac{FP}{FP + TN},$$

$$\text{False Negative Rate (FNR)} = \frac{FN}{FN + TP} \text{ and}$$

$$\text{True Negative Rate (TNR)} = \frac{TN}{TN + FP}$$

**DOR** is the measure of the ratio of the odds of race detection being positive given that the test case has a data race, to the odds of race detection being positive given the test case does not have a race.

**System configuration.** We performed all our experiments on a system with two Intel Xeon E5-2697 v4 @ 2.30GHz processors, each having 18 cores and 2 threads per core, totalling 72 threads and 128GB of RAM. The system runs 64 bit Ubuntu 18.04.2 LTS server with Linux kernel version 4.15.0-48-generic. LLOV is currently based on the LLVM/Polly version release 7.0.1 and can be upgraded to the latest LLVM/Polly versions with minimal changes.

Similar to Liao et al. [40], our experiments use two parameters: (i) the number of OpenMP threads and (ii) the input size for variable length arrays. The number of threads that we considered for the experiments are {3, 36, 45, 72, 90, 180, 256}. For the 16 variable length kernels, we considered 6 different array sizes as follows: {32, 64, 128, 256, 512, 1024}. With each particular set of parameters,

Table 3. Maximum number of Races reported by different tools in DataRaceBench 1.2

Tools	Race: Yes		Race: No		Coverage/116
	TP	FN	TN	FP	
HELGRIND	56	3	2	55	116
VALGRIND DRD	56	3	26	31	116
TSAN-LLVM	57	2	2	55	116
ARCHER	56	3	2	55	116
SWORD	47	4	24	4	79
LLOV	45	3	28	9	85

we ran each of the 116 kernels 5 times. Both the number of threads and array sizes can be found in prior studies [40, 42] and we have used the same for uniformity. Since the dynamic tools depend on the execution order of the threads, multiple runs are required. The 16 kernels with variable length arrays were run 3360 (5 times x 16 kernels x 7 thread sizes x 6 array sizes) times in total. The remaining 100 kernels were run 3500 (5 times x 100 kernels x 7 thread sizes) times in total. For all experiments we used a timeout of 600 seconds for both compilation and execution.

## 6.2 Experimental Results

**6.2.1 DataRaceBench 1.2.** Table 3 provides comparison in terms of the number of races detected in DataRaceBench v1.2. Column 1 indicates the name of the tool. The column with titles “Race:Yes” and “Race:No” indicate if the benchmark had a race or not. The subcolumns “TP” and “FN” denote whether the tool was able to find the race or not when the benchmark had a race. Similarly, the subcolumns “FP” and “TN” denote if the tool erroneously reported a race or reported the absence of a race when the benchmark did not have a race. As dynamic analysis tools are run multiple times, if a tool reports a race in any of the runs, it is considered that the tool is classifying a benchmark as having a race.

LLOV could analyze 85 out of 116 (73%) kernels from the benchmark, and could detect 45 True Positives (TP) and 28 True Negatives (TN). As the analysis of our tool is conservative, it also produces 9 False Positives (FP). Moreover, it had 3 False Negatives (FN), which are mostly implementation issues that we believe can be fixed.

Due to the sound static analysis that LLOV implements, it could also prove that 28 of the kernels are *data race free*. LLOV is unique in this regard, other tools *are not able to make* such a claim. LLOV will report one of the following three cases: *data race detected* when the input program has a race, *data race free* when the input program is race free, and finally, *region not analyzed* when LLOV cannot analyze the input program. In addition, LLOV also reports if an OpenMP pragma is not supported (refer to Table 1). Due to these reasons, as of now LLOV does not provide complete coverage on DataRaceBench v1.2.

Due to compilation related issues, SWORD provides even lesser coverage on DataRaceBench v1.2.

Table 4 shows performance of the tool on various performance metrics defined earlier in the section. From Table 4 it appears that SWORD is the best whereas LLOV is a close second in terms of precision and accuracy. Since both SWORD and LLOV do not have complete coverage, a more appropriate strategy would be to compare all the tools on only those benchmarks which they are able to handle/cover.

Table 4. Precision, Recall and Accuracy of the tools on DataRaceBench 1.2

Tools	Precision	Recall	Accuracy	F1 Score	Diagnostic odds ratio
HELGRIND	0.50	0.95	0.50	0.66	0.68
VALGRIND DRD	0.64	0.95	0.71	0.77	15.66
TSAN-LLVM	0.51	<b>0.97</b>	0.51	0.67	1.04
ARCHER	0.50	0.95	0.50	0.66	0.68
SWORD	<b>0.92</b>	0.92	<b>0.90</b>	<b>0.92</b>	<b>70.50</b>
LLOV	0.83	0.94	0.86	0.88	46.67

Table 5. Maximum number of Races reported by different tools in common 66 kernels of DataRaceBench 1.2

Tools	Race: Yes		Race: No		Coverage/116
	TP	FN	TN	FP	
HELGRIND	46	1	2	17	66
VALGRIND DRD	46	1	13	6	66
TSAN-LLVM	46	1	2	17	66
ARCHER	46	1	2	17	66
SWORD	46	1	18	1	66
LLOV	44	3	16	3	66

Table 6. Precision, Recall and Accuracy of the tools on common 66 kernels of DataRaceBench 1.2

Tools	Precision	Recall	Accuracy	F1 Score	Diagnostic odds ratio
HELGRIND	0.73	<b>0.98</b>	0.73	0.84	5.41
VALGRIND DRD	0.88	<b>0.98</b>	0.89	0.93	99.67
TSAN-LLVM	0.73	<b>0.98</b>	0.73	0.84	5.41
ARCHER	0.73	<b>0.98</b>	0.73	0.84	5.41
SWORD	<b>0.98</b>	<b>0.98</b>	<b>0.97</b>	<b>0.98</b>	<b>828.00</b>
LLOV	0.94	0.94	0.91	0.94	78.22

Table 5 shows how tools classify benchmarks with respect to data race on 66 benchmarks that all the tools are able to handle/cover. Table 6 provides a comparison of the tools on 66 benchmarks on various performance metrics. It is indeed the case that on the benchmarks that SWORD is able to handle, it achieves the highest precision, accuracy, recall, F1 score and diagnostic odds ratio. LLOV is a close second with respect to precision, accuracy and F1 score. One must keep in mind that both SWORD and LLOV may gain advantage in terms of these metrics because of lesser coverage. A crucial point to note is that while SWORD crashes on many benchmarks, LLOV provides a graceful reporting and exit on benchmarks it is not able to cover, providing a better user experience.

Though LLOV does not come out on top on various metric such as coverage, precision etc., it completely outshines other tools in terms of runtime. Figure 2a shows the performance of the tool with respect to the run time. Since dynamic tools run benchmarks multiple times we report the average time taken for each benchmark, and the total time is the sum of these averages. In Figure 2a,  $y$ -axis represents total time taken in seconds by a tool on a logarithmic scale. Time taken

by LLOV to analyze all 116 kernels is a mere 43 seconds. On the other hand, other tools take orders of magnitude more time as compared to LLOV. The reason for SWORD performing the worst when all 116 benchmarks are considered is because the compilation process itself times out for several benchmarks. Timeout for every run is 600 seconds including the compile time. For LLOV the only time required is the time to compile as it does its analysis at compile time. In addition, as LLOV is able to detect the cases it can not analyze, the exits for such programs are graceful. The power of static analysis in LLOV is particularly evident in Figure 2 as the time taken remains constant irrespective of the number of threads.

The Polyhedral framework is known for large compile-time overheads because of the exponential nature of algorithms used in dependence analysis, scheduling and code-generation. The analysis time increases exponentially with the number of dimensions in the loop nest. However, very few programs in the real world have very large loop-depths.

In DataRaceBench v1.2, there are six tiled and parallel versions of PolyBench/C 3.2 kernels. The tiled version of matrix multiplication kernel DRB042-3mm-tile-no has 408 OpenMP parallel loops and contributes to around 60% of total time taken by LLOV for all the 116 kernels. Although such kernels are not very common in real world scenarios, as they are generated by polyhedral tools such as Pluto [10], verification of code generated by such automatic tools remains a challenge.

Figure 2b shows the runtime performance of the tools on 66 benchmarks that all the tools are able to cover. LLOV outperforms all the other tools on this subset of benchmarks by orders of magnitude. SWORD performance is much better on this subset and comes third with TSAN while ARCHER coming second. The performance of TSAN and ARCHER only marginally differ. ARCHER instruments the program and then invokes TSAN. Therefore, it is possible that for some subset of benchmarks, such an instrumentation may not help ARCHER achieve better performance than TSAN.

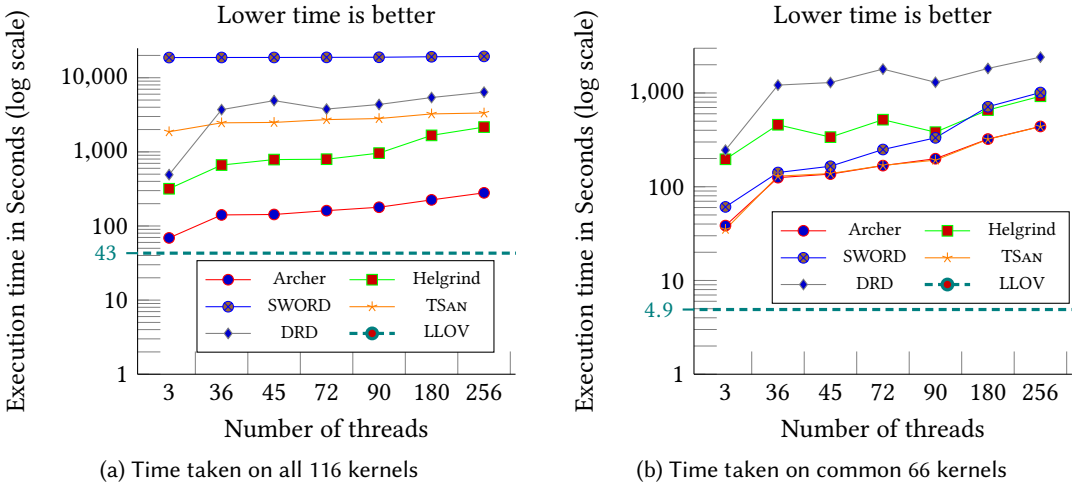


Fig. 2. DataRaceBench v1.2 total time taken on logarithmic scale

6.2.2 *DataRaceBench 1.2 FORTRAN*. Since LLOV is based on LLVM-IR, it is language independent. To demonstrate this we reimplemented DataRaceBench 1.2 in FORTRAN 95 (Github) rewriting 92 of the 116 DataRaceBench v1.2 kernels in FORTRAN. The other kernels, such as {41, 42, 43, 44, 55, 56},

are Polybench kernels that were tiled and/or parallelized by the POCC [54] polyhedral tool, and are not amenable to easy re-writing in FORTRAN manually.

The kernel in Listing 2 (Section 3) is from DataRaceBench 1.2 FORTRAN which has a data race because of the write to shared variable `tmp` (Line 3) and the read from `tmp` (Line 4). The race in this example can be avoided by explicitly stating that `tmp` is a private variable for each thread using the *private* clause. LLOV could detect such races because of missing data sharing clauses, such as *private*, *reduction*, *firstprivate*, and *lastprivate*.

To verify these FORTRAN kernels, we used LLVM FORTRAN frontend Flang [44] which is under active development and has officially been accepted in 2019 as a sub-project under the LLVM Compiler Infrastructure umbrella project. We used Flang version 7.1.0 (git commit hash a81f07070b) to generate LLVM-IR from FORTRAN source code and ran LLOV on the generated IR.

Table 7. Maximum number of Races reported by different tools in DataRaceBench FORTRAN

Tools	Race: Yes		Race: No		Coverage/92
	TP	FN	TN	FP	
HELGRIND	46	6	4	36	92
VALGRIND DRD	45	7	21	19	92
LLOV	34	6	19	5	64

Initial experiments show that our analysis is able to detect race conditions in OpenMP kernels of DataRaceBench FORTRAN. To the best of our knowledge, LLOV is the *only static tool* to be able to detect races in OpenMP programs written in FORTRAN. Table 7 shows that LLOV could detect 34 True Positives (TP) and also confirm that 19 kernels are data race free (TN). LLOV also produced 5 False Positives (FP) along with 6 False Negatives (FN).

Kernels {72, 78, 79, 94, 112} are not analyzed by LLOV because the current version of Flang does not support the corresponding OpenMP directives. This explains the difference in the numbers from DataRaceBench v1.2 (Table 3) and DataRaceBench FORTRAN (Table 7). Flang produced segmentation faults for the kernels {84, 85} having OpenMP *threadprivate* variables. As Flang is in active development, we believe that more OpenMP directives will be supported in the forthcoming releases. Another challenge we face is detecting polyhedral static control parts (SCoPs) in Flang generated IR using Polly [45]. The IR generated by Clang and Flang for functionally equivalent source code differs considerably. Hence, the native implementations of analyses and optimization passes in LLVM needs to be modified to have uniform results on LLVM-IR generated by Clang and Flang.

**6.2.3 OmpSCR v2.0.** We evaluated all the tools listed in Table 2 on OmpSCR v2.0 [19, 20] kernels. For the dynamic tools, we used default program parameters provided with argument *-test*. We updated OmpSCR v2.0 [19] to run with latest compilers and created scripts to test all the data race checkers. The updated version of OmpSCR can be found at [https://github.com/utpalbora/OmpSCR\\_v2.0.git](https://github.com/utpalbora/OmpSCR_v2.0.git).

We manually verified OmpSCR v2.0 benchmark suite. Table 8 divides the benchmarks in three categories: 1) Manually verified kernels with data races, 2) Manually verified race free kernels and 3) Unverified kernels. Out of 25 kernels, 11 remained unverified due to various complexities such as recursive calls using OpenMP pragmas. Every cell in the Table 8 denotes how many different regions are reported as containing a race by a tool. Every verified kernel having a race contains only one parallel region containing races. All the races reported by a tool belonging to a single parallel region is counted as only one. The reason for combining all the races in a region is because otherwise, the number of races reported for dynamic tools becomes quite high (e.g., several races



Table 8. Number of Races detected in OmpSCR v2.0 benchmark (CT is Compilation Timeout)

Kernel	LLOV	HELGRIND	DRD	TSAN	ARCHER	SWORD
Manually verified kernels with data races						
c_loopA.badSolution	1	1	1	1	1	1
c_loopA.solution2	1	1	1	1	1	0
c_loopA.solution3	1	1	1	1	1	0
c_loopB.badSolution1	1	1	1	1	1	1
c_loopB.badSolution2	1	1	1	1	1	1
c_loopB.pipelineSolution	1	1	1	1	1	0
c_md	1	2	2	2	1	CT
c_lu	1	1	1	1	1	0
Manually verified race free kernels						
c_loopA.solution1	0	2	1	2	1	0
c_mandel	0	1	0	1	1	0
c_pi	0	1	0	1	1	0
c_jacobi01	1	2	1	0	0	CT
c_jacobi02	1	1	1	0	0	CT
c_jacobi03	0	1	1	0	0	CT
Unverified kernels						
c_fft	1	1	1	1	1	CT
c_fft6	1	1	0	1	1	CT
c_qsort	0	1	1	1	1	CT
c_GraphSearch	0	0	0	0	0	0
cpp_qsomp1	0	0	0	0	0	0
cpp_qsomp2	0	0	0	0	0	0
cpp_qsomp3	0	0	0	0	0	0
cpp_qsomp4	0	0	0	0	0	0
cpp_qsomp5	0	0	0	0	0	0
cpp_qsomp6	0	0	0	0	0	0
cpp_qsomp7	0	0	0	0	0	0

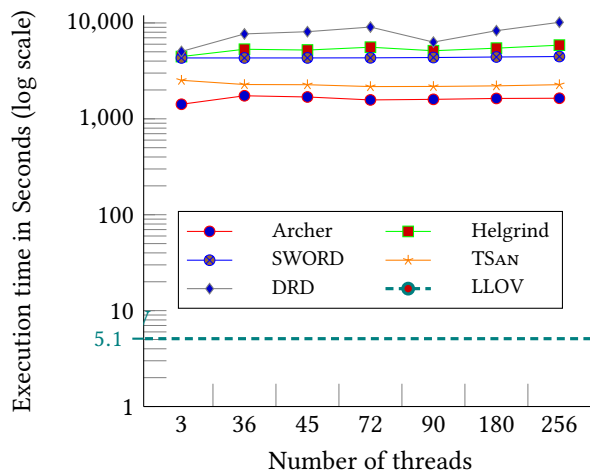


Fig. 3. OmpSCR v2.0 total execution time by different tools on logarithmic scale

Table 9. Comparison of different tools on OmpSCR v2.0

Tools	Race: Yes		Race: No		Coverage/14
	TP	FN	TN	FP	
HELGRIND	8	0	0	9	14
VALGRIND DRD	8	0	2	5	14
TSAN-LLVM	8	0	3	5	14
ARCHER	8	0	3	3	14
SWORD	3	4	3	0	10
LLOV	8	0	4	2	14

reported for a single array). Some of the programs contain multiple parallel regions. If a tool reports races in two distinct regions, then we count the tool as reporting 2 races. Since every verified kernel with race has only one parallel region having a race, races reported in other race-free regions are counted towards *false positives*.

As shown in Table 8, LLOV is able to detect true data races in Molecular Dynamics (`c_md`), `c_loopA` and `c_loopB` kernels. LLOV produced false positives for Jacobi kernels due to conservative alias analysis in LLVM. All the false positives flagged by LLOV are because of shared double pointer variables. In addition, SWORD ends up with compiler timeout (CT) for kernels such as Molecular Dynamics, Quicksort (`c_qsort`), FFT, and Jacobi. The time taken to detect races in all the kernels by the tools is shown in Figure 3. It can be seen from the figure that LLOV completes its analysis for the entire benchmark in just *5.1 seconds*, while the other state-of-the-art tools take orders of magnitude longer.

## 7 CONCLUSIONS AND FUTURE WORK

In this paper, we present LLOV, a language agnostic, OpenMP aware, static analysis based data race detection tool which is developed on top of the LLVM compiler framework. As LLOV operates at LLVM-IR level, it can support a multitude of programming languages supported by the LLVM infrastructure. We successfully demonstrate the language agnostic nature of LLOV by performing data race checks on a standard set of benchmarks written in C, C++, and FORTRAN.

Our experiments show that LLOV performs reasonably well in terms of precision and accuracy while being *most performant* with respect to other tools by a large margin. Though at present, LLOV supports only some of the pragmas offered by OpenMP, it gracefully exits on input programs that contain pragmas that it is unable to handle. We would like to further enrich LLOV by adding support for various OpenMP pragmas that are not supported at present. Many such pragmas offer an engineering challenge as the structural information is not available at LLVM-IR level, and such information has to be reconstructed from the IR.

Our tool is currently based on the polyhedral compilation framework, Polly. The use of approximate dependence analysis [31] readily available in LLVM may further increase the capability and scalability of LLOV. We also plan to extend the support for dynamic control flow and irregular accesses using the extended polyhedral framework [5, 6, 16].

## REFERENCES

- [1] Alfred V. Aho, Monica S. Lam, Ravi Sethi, and Jeffrey D. Ullman. 2006. *Compilers: Principles, Techniques, and Tools (2<sup>nd</sup> Edition)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- [2] Simone Atzeni, Ganesh Gopalakrishnan, Zvonimir Rakamaric, Dong H Ahn, Ignacio Laguna, Martin Schulz, Gregory L Lee, Joachim Protze, and Matthias S Müller. 2016. ARCHER: effectively spotting data races in large OpenMP applications. In *Parallel and Distributed Processing Symposium, 2016 IEEE International*. IEEE, IEEE, Chicago, IL, USA, 53–62.
- [3] Simone Atzeni, Ganesh Gopalakrishnan, Zvonimir Rakamaric, Ignacio Laguna, Gregory L Lee, and Dong H Ahn. 2018. Sword: A bounded memory-overhead detector of OpenMP data races in production runs. In *2018 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, IEEE, Vancouver, BC, Canada, 845–854.
- [4] Vamshi Basupalli, Tomofumi Yuki, Sanjay Rajopadhye, Antoine Morvan, Steven Derrien, Patrice Quinton, and David Wonnacott. 2011. ompVerify: polyhedral analysis for the OpenMP programmer. In *International Workshop on OpenMP*. Springer, Springer Berlin Heidelberg, Berlin, Heidelberg, 37–53.
- [5] Marouane Belaoucha, Denis Barthou, Adrien Eliche, et al. 2010. FADALib: an open source C++ library for fuzzy array dataflow analysis. *Procedia Computer Science* 1, 1 (2010), 2075–2084.
- [6] Mohamed-Walid Benabderrahmane, Louis-Noël Pouchet, Albert Cohen, and Cédric Bastoul. 2010. The Polyhedral Model is More Widely Applicable Than You Think. In *Proceedings of the 19th Joint European Conference on Theory and Practice of Software, International Conference on Compiler Construction (CC'10/ETAPS'10)*. Springer-Verlag, Berlin, Heidelberg, 283–303. [https://doi.org/10.1007/978-3-642-11970-5\\_16](https://doi.org/10.1007/978-3-642-11970-5_16)
- [7] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. 2017. Julia: A fresh approach to numerical computing. *SIAM review* 59, 1 (2017), 65–98. <https://doi.org/10.1137/141000671>
- [8] Sam Blackshear, Nikos Gorogiannis, Peter W. O&#39;Hearn, and Ilya Sergey. 2018. RacerD: Compositional Static Race Detection. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 144 (Oct. 2018), 28 pages. <https://doi.org/10.1145/3276514>
- [9] Robert D. Blumofe, Christopher F. Joerg, Bradley C. Kuszmaul, Charles E. Leiserson, Keith H. Randall, and Yuli Zhou. 1996. Cilk: An Efficient Multithreaded Runtime System. *J. Parallel and Distrib. Comput.* 37, 1 (1996), 55 – 69. <https://doi.org/10.1006/jpdc.1996.0107>
- [10] Uday Bondhugula, Albert Hartono, J. Ramanujam, and P. Sadayappan. 2008. A Practical Automatic Polyhedral Parallelizer and Locality Optimizer. In *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '08)*. ACM, New York, NY, USA, 101–113. <https://doi.org/10.1145/1375581.1375595>
- [11] Utpal Bora, Johannes Doefert, Tobias Grosser, and Ramakrishna Upadrasta. 2016. GSoc 2016: PolyhedralInfo - Polly as an analysis pass in LLVM. <https://llvmddevelopersmeetingbay2016.sched.com/event/8Z2Z/lightning-talks>. [Online; accessed 08-May-2019].
- [12] B.L. Chamberlain, D. Callahan, and H.P. Zima. 2007. Parallel Programmability and the Chapel Language. *Int. J. High Perform. Comput. Appl.* 21, 3 (Aug. 2007), 291–312. <https://doi.org/10.1177/1094342007078442>
- [13] Philippe Charles, Christian Grothoff, Vijay Saraswat, Christopher Donawa, Allan Kielstra, Kemal Ebcioglu, Christoph von Praun, and Vivek Sarkar. 2005. X10: An Object-oriented Approach to Non-uniform Cluster Computing. In *Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications (OOPSLA '05)*. ACM, New York, NY, USA, 519–538. <https://doi.org/10.1145/1094811.1094852>
- [14] Prasanth Chatarasi, Jun Shirako, Martin Kong, and Vivek Sarkar. 2016. An extended polyhedral model for SPMD programs and its use in static data race detection. In *International Workshop on Languages and Compilers for Parallel Computing*. Springer, Springer International Publishing, Cham, 106–120.
- [15] P. Chatarasi, J. Shirako, and V. Sarkar. 2015. Polyhedral Optimizations of Explicitly Parallel Programs. In *2015 International Conference on Parallel Architecture and Compilation (PACT)*. IEEE, San Francisco, CA, USA, 213–226. <https://doi.org/10.1109/PACT.2015.44>
- [16] Jean-François Collard, Denis Barthou, and Paul Feautrier. 1995. Fuzzy Array Dataflow Analysis. In *Proceedings of the Fifth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPOPP '95)*. ACM, New York, NY, USA, 92–101. <https://doi.org/10.1145/209936.209947>
- [17] Leonardo Dagum and Rameshm Enon. 1998. OpenMP: an industry standard API for shared-memory programming. *Computational Science & Engineering, IEEE* 5, 1 (1998), 46–55.
- [18] Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems*, C. R. Ramakrishnan and Jakob Rehof (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 337–340.
- [19] A.J. Dorta. 2004. OpenMP Source Code Repository. <https://sourceforge.net/projects/ompscr/files/OmpSCR/OmpSCR>. [Online; accessed 19-May-2019].
- [20] A. J. Dorta, C. Rodriguez, and F. de Sande. 2005. The OpenMP source code repository. In *13th Euromicro Conference on Parallel, Distributed and Network-Based Processing*. IEEE Computer Society, Washington, DC, USA, 244–250. <https://doi.org/10.1109/EMPDP.2005.41>

- [21] Laura Effinger-Dean, Brandon Lucia, Luis Ceze, Dan Grossman, and Hans-J. Boehm. 2012. IFRit: Interference-free Regions for Dynamic Data-race Detection. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications (OOPSLA '12)*. ACM, New York, NY, USA, 467–484. <https://doi.org/10.1145/2384616.2384650>
- [22] Alexandre E. Eichenberger, John Mellor-Crummey, Martin Schulz, Michael Wong, Nawal Copty, Robert Dietrich, Xu Liu, Eugene Loh, and Daniel Lorenz. 2013. OMPT: An OpenMP Tools Application Programming Interface for Performance Analysis. In *OpenMP in the Era of Low Power Devices and Accelerators*, Alistair P. Rendell, Barbara M. Chapman, and Matthias S. Müller (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 171–185.
- [23] Dawson Engler and Ken Ashcraft. 2003. RacerX: Effective, Static Detection of Race Conditions and Deadlocks. In *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles (SOSP '03)*. ACM, New York, NY, USA, 237–252. <https://doi.org/10.1145/945445.945468>
- [24] Paul Feautrier. 1991. Dataflow analysis of array and scalar references. *International Journal of Parallel Programming* 20, 1 (1991), 23–53. <https://doi.org/10.1007/BF01407931>
- [25] Cormac Flanagan and Stephen N. Freund. 2009. FastTrack: Efficient and Precise Dynamic Race Detection. In *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '09)*. ACM, New York, NY, USA, 121–133. <https://doi.org/10.1145/1542476.1542490>
- [26] MPI Forum. 2019. MPI-3.1 (May 2019). <https://www.mpi-forum.org/docs/>. [Online; accessed 08-August-2019].
- [27] Tobias Grosser, Armin Größlinger, and Christian Lengauer. 2012. Polly - Performing Polyhedral Optimizations on a Low-Level Intermediate Representation. *Parallel Processing Letters* 22, 4, Article 1 (2012), 27 pages. <https://doi.org/10.1142/S0129626412500107>
- [28] Khronos OpenCL Working Group. 2019. The OpenCL Specification, v2.2-11 (July 2019). <https://www.khronos.org/registry/OpenCL/>. [Online; accessed 08-August-2019].
- [29] LLVM Developer Group. 2019. Clang: A C language family frontend for LLVM. <https://clang.llvm.org/>. [Online; accessed 08-August-2019].
- [30] LLVM Developer Group. 2019. LLVM Language Reference Manual. <https://llvm.org/docs/LangRef.html>. [Online; accessed 08-August-2019].
- [31] LLVM Developer Group. 2019. Loop Access Info, Class Reference. [https://llvm.org/doxygen/classllvm\\_1\\_1LoopAccessInfo.html](https://llvm.org/doxygen/classllvm_1_1LoopAccessInfo.html). [Online; accessed 08-May-2019].
- [32] LLVM Developer Group. 2019. The LLVM Compiler Infrastructure. <http://llvm.org/>. [Online; accessed 08-May-2019].
- [33] POSIX Austin Joint Working Group. 2018. IEEE Standard for Information Technology–Portable Operating System Interface (POSIX(R)) Base Specifications, Issue 7. *IEEE Std 1003.1-2017 (Revision of IEEE Std 1003.1-2008)* (Jan 2018), 1–3951. <https://doi.org/10.1109/IEEEESTD.2018.8277153>
- [34] Yizi Gu and John Mellor-Crummey. 2018. Dynamic Data Race Detection for OpenMP Programs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '18)*. IEEE Press, Piscataway, NJ, USA, Article 61, 12 pages. <http://dl.acm.org/citation.cfm?id=3291656.3291738>
- [35] Cray Inc. 2019. Chapel language specification (version 0.9). <http://chapel.cray.com/papers.html>. [Online; accessed 08-May-2019].
- [36] Baris Kasikci, Cristian Zamfir, and George Candea. 2012. CoRD: A Collaborative Framework for Distributed Data Race Detection. In *Proceedings of the Eighth USENIX Conference on Hot Topics in System Dependability (HotDep'12)*. USENIX Association, Berkeley, CA, USA, 4–4. <http://dl.acm.org/citation.cfm?id=2387858.2387862>
- [37] Pankaj Kukreja, Himanshu Shukla, and Utpal Bora. 2019. DataRaceBench FORTRAN. [https://github.com/IITH-Compilers/drb\\_fortran](https://github.com/IITH-Compilers/drb_fortran). [Online; accessed 19-October-2019].
- [38] Leslie Lamport. 1978. Time, Clocks, and the Ordering of Events in a Distributed System. *Commun. ACM* 21, 7 (July 1978), 558–565. <https://doi.org/10.1145/359545.359563>
- [39] N. G. Leveson and C. S. Turner. 1993. An investigation of the Therac-25 accidents. *Computer* 26, 7 (July 1993), 18–41. <https://doi.org/10.1109/MC.1993.274940>
- [40] Chunhua Liao, Pei-Hung Lin, Joshua Asplund, Markus Schordan, and Ian Karlin. 2017. DataRaceBench: A Benchmark Suite for Systematic Evaluation of Data Race Detection Tools. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '17)*. ACM, New York, NY, USA, Article 11, 14 pages. <https://doi.org/10.1145/3126908.3126958>
- [41] Chunhua Liao, Pei-Hung Lin, Joshua Asplund, Markus Schordan, and Ian Karlin. 2018. DataRaceBench v1.2.0. <https://github.com/LLNL/dataracebench>. [Online; accessed 19-May-2019].
- [42] Chunhua Liao, Pei-Hung Lin, Markus Schordan, and Ian Karlin. 2018. A Semantics-Driven Approach to Improving DataRaceBench's OpenMP Standard Coverage. In *Evolving OpenMP for Evolving Architectures*, Bronis R. de Supinski, Pedro Valero-Lara, Xavier Martorell, Sergi Mateo Bellido, and Jesus Labarta (Eds.). Springer International Publishing, Cham, 189–202.
- [43] Jacques-Louis Lions et al. 1996. Flight 501 failure. (1996).

- [44] LLVM/Flang. 2019. Flang: A Fortran Compiler Targeting LLVM. <https://github.com/flang-compiler/flang/wiki>. [Online; accessed 08-May-2019].
- [45] LLVM/Polly. 2019. Polly: LLVM Framework for High-Level Loop and Data-Locality Optimizations. <https://polly.llvm.org/>. [Online; accessed 08-May-2019].
- [46] John Mellor-Crummey. 1991. On-the-fly Detection of Data Races for Programs with Nested Fork-join Parallelism. In *Proceedings of the 1991 ACM/IEEE Conference on Supercomputing (Supercomputing '91)*. ACM, New York, NY, USA, 24–33. <https://doi.org/10.1145/125826.125861>
- [47] John Mellor-Crummey. 1991. On-the-fly Detection of Data Races for Programs with Nested Fork-join Parallelism. In *Proceedings of the 1991 ACM/IEEE Conference on Supercomputing (Supercomputing '91)*. ACM, New York, NY, USA, 24–33. <https://doi.org/10.1145/125826.125861>
- [48] Nicholas Nethercote and Julian Seward. 2003. Valgrind: A program supervision framework. <http://valgrind.org/>. [Online; accessed 08-May-2019].
- [49] OpenMP Architecture Review Board. 1997. OpenMP Application Programming Interface. <https://www.openmp.org>. [Online; accessed 19-October-2019].
- [50] OpenMP Architecture Review Board. 2015. OpenMP Application Programming Interface Version 4.5. <https://www.openmp.org/wp-content/uploads/openmp-4.5.pdf>. [Online; accessed 19-May-2019].
- [51] OpenACC Organization. 2019. The OpenACC Application Programming Interface,v2.7 (November 2018). <https://www.openacc.org/specification>. [Online; accessed 08-August-2019].
- [52] PGI. 2018. PGI Accelerator Compilers with OpenACC Directives. <https://www.pgroup.com/resources/accel.htm>. [Online; accessed 08-August-2019].
- [53] Nathaniel Popper. 2012. Knight Capital says trading glitch cost it \$440 million. *New York Times* 2 (2012).
- [54] Louis-Noël Pouchet. 2018. PoCC, the POLYhedral Compiler Collection package. A full source-to-source polyhedral compiler. <https://sourceforge.net/projects/pocc/>. [Online; accessed 19-October-2019].
- [55] Polyvios Pratikakis, Jeffrey S Foster, and Michael Hicks. 2011. LOCKSMITH: Practical static race detection for C. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 33, 1 (2011), 3.
- [56] GNU Project. 2019. GCC, the GNU Compiler Collection. <https://gcc.gnu.org/>. [Online; accessed 08-August-2019].
- [57] Dan Quinlan and Chunhua Liao. 2011. The ROSE source-to-source compiler infrastructure. In *Cetus users and compiler infrastructure workshop, in conjunction with PACT*, Vol. 2011. Citeseer, 1.
- [58] Stefan Savage, Michael Burrows, Greg Nelson, Patrick Sobalvarro, and Thomas Anderson. 1997. Eraser: A dynamic data race detector for multithreaded programs. *ACM Transactions on Computer Systems (TOCS)* 15, 4 (1997), 391–411.
- [59] Konstantin Serebryany and Timur Iskhodzhanov. 2009. ThreadSanitizer: Data Race Detection in Practice. In *Proceedings of the Workshop on Binary Instrumentation and Applications (WBIA '09)*. ACM, New York, NY, USA, 62–71. <https://doi.org/10.1145/1791194.1791203>
- [60] Konstantin Serebryany, Alexander Potapenko, Timur Iskhodzhanov, and Dmitriy Vyukov. 2012. Dynamic Race Detection with LLVM Compiler. In *Proceedings of the Second International Conference on Runtime Verification (RV'11)*. Springer-Verlag, Berlin, Heidelberg, 110–114. [https://doi.org/10.1007/978-3-642-29860-8\\_9](https://doi.org/10.1007/978-3-642-29860-8_9)
- [61] TOP500.Org. 2019. Top 500 Supercomputer Sites. <http://www.top500.org/>. [Online; accessed 08-May-2019].
- [62] Valgrind-project. 2007. DRD: a thread error detector. <http://valgrind.org/docs/manual/drd-manual.html>. [Online; accessed 08-May-2019].
- [63] Valgrind-project. 2007. Helgrind: a thread error detector. <http://valgrind.org/docs/manual/hg-manual.html>. [Online; accessed 08-May-2019].
- [64] Sven Verdoolaege. 2010. Isl: An Integer Set Library for the Polyhedral Model. In *Proceedings of the Third International Congress Conference on Mathematical Software (ICMS'10)*. Springer-Verlag, Berlin, Heidelberg, Article 1250010, 4 pages. <http://dl.acm.org/citation.cfm?id=1888390.1888455>
- [65] Jan Wen Voung, Ranjit Jhala, and Sorin Lerner. 2007. RELAY: static race detection on millions of lines of code. In *Proceedings of the 6th joint meeting of the European software engineering conference and the ACM SIGSOFT symposium on The foundations of software engineering*. ACM, ACM, New York, NY, USA, 205–214.
- [66] Fangke Ye, Markus Schordan, Chunhua Liao, Pei-Hung Lin, Ian Karlin, and Vivek Sarkar. 2018. Using Polyhedral Analysis to Verify OpenMP Applications are Data Race Free. In *2018 IEEE/ACM 2nd International Workshop on Software Correctness for HPC Applications (Correctness)*. IEEE, IEEE, Dallas, TX, USA, 42–50. <https://doi.org/10.1109/Correctness.2018.00010>
- [67] Yuan Yu, Tom Rodeheffer, and Wei Chen. 2005. RaceTrack: Efficient Detection of Data Race Conditions via Adaptive Tracking. *SIGOPS Oper. Syst. Rev.* 39, 5 (Oct. 2005), 221–234. <https://doi.org/10.1145/1095809.1095832>
- [68] Tomofumi Yuki, Gautam Gupta, DaeGon Kim, Tanveer Pathan, and Sanjay Rajopadhye. 2012. Alphaz: A system for design space exploration in the polyhedral model. In *International Workshop on Languages and Compilers for Parallel Computing*. Springer, Springer Berlin Heidelberg, Berlin, Heidelberg, 17–31.

- [69] Intel Developer Zone. 2019. The Intel C++ Compiler. <https://software.intel.com/en-us/c-compilers>. [Online; accessed 08-August-2019].