

Final Paper  
for  
CSC 431



## Contents

<b>1 Overview</b>	<b>1</b>
1.1 Parsing	2
1.2 Static Semantics	2
1.3 Intermediate Representation	2
1.4 Optimizations	2
1.5 Code Generation and Register Allocation	3
1.6 Other	3
<b>2 Analysis</b>	<b>3</b>
2.1 OptimizationBenchmark	4
2.2 BenchMarkishTopics	4
2.3 Fibonacci	5
2.4 GeneralFunctAndOptimize	6
2.5 TicTac	6
2.6 bert	7
2.7 biggest	8
2.8 binaryConverter	8
2.9 creativeBenchMarkName	9
2.10 fact_sum	10
2.11 hailstone	10
2.12 hanoi_benchmark	11
2.13 killerBubbles	12
2.14 mile1	12
2.15 mixed	13
2.16 primes	14
2.17 programBreaker	14
2.18 stats	15
2.19 uncreativeBenchmark	16
2.20 wasteOfCycles	16

## 1 Overview

I chose initially to write my compiler in Java. But after learning more Scala to work on my project for [REDACTED], I chose to use Scala for the later parts (optimization, code generation, and register allocation).

## 1.1 Parsing

The compiler creates an Antlr `CharStream` from the input file. This is then passed to the generated lexer, and the resulting token stream is passed to the generated parser. The success or failure of the parse is represented with FunctionalJava's `Option` type, with a failed parse resulting in `None` and a successful parse resulting in a `Some` containing the parse tree. Since the parser handles printing parse errors, there is no need to store the parse error in the result structure.

## 1.2 Static Semantics

An interface for all program validations was defined and used to provide a uniform way to run them. This allows validations to be added or removed as desired, making the compiler more modular.

Return statement checking was implemented as a program validation. It ensures each function has a return-equivalent statement in it. A return-equivalent statement is defined as a block containing a return-equivalent statement, a conditional statement in which both branches contain a return-equivalent statement, a recursive invocation of the function, or an actual return statement. With those definitions, checking a function for a return-equivalent statement requires only a simple recursive function.

Typechecking was implemented as another program validation. An interface named `Typed` was created and all type-checkable AST elements implement it. This interface provides a single method, `type()`, which returns a FunctionalJava `Validation`<sup>1</sup> containing either an error (of type `ValidationException`) or the type of the element it was called on, represented as one of the provided types. For statements, `type()` will return `VoidType` if the types for any subexpressions are valid. Thus, to typecheck a function, we call the `type()` method on its block statement and check if the result is successful or not (i.e., if the result is a failure or an instance of `VoidType`).

## 1.3 Intermediate Representation

As the result of parsing and validation, an abstract syntax tree is produced, representing the program in a language-dependent fashion. The compiler converts this to a list of functions, each represented as a control flow graph, with LLVM IR instructions stored in each block of the graph. The LLVM instructions provide a layer of abstraction between the code and the hardware instruction set, while not being dependent on the language's structure. LLVM IR is also designed to be in Static Single Assignment form. This makes some program transformations easier and more efficient, as it stores the use-def chain as part of the IR[1]. Converting to a CFG reorients the structure of the program code around the control flow, rather than the syntax of the source language[2]. This also helps to abstract the program representation into a unified shape for all source languages.

In the compiler presented here, `Derive4J` was used to generate classes and a Visitor pattern implementation in Java for the instructions in the IR. The rest of the classes for the CFG are just plain Java classes. The CFG basic block class stores links to its child blocks to create the graph, but it is not necessary to traverse these links to visit all blocks, as blocks are also stored in a list attached to the function they are from. This makes printing them out much simpler.

`Derive4J` worked reasonably well, although its generated pattern matching was less flexible than Scala's. I would probably choose to use Scala case classes if I were to reimplement this part of the code.

## 1.4 Optimizations

Two optimizations were implemented for this compiler: constant propagation and useless code elimination.

To implement constant propagation, the Sparse Conditional Constant Propagation algorithm was used. This algorithm tracks the blocks that are able to be executed, given what is known about the constant values in the program. For example, if a conditional branch has a constant value for a condition, then the algorithm will ignore the block that the untaken branch would have jumped to (unless, of course, it is used by another branch statement). Within the blocks that *are* executed, each instruction is abstractly executed in order

---

<sup>1</sup>A `Validation` is an `Either` designed for failure-checking. Thus, it provides some convenience methods for accumulating failures across a list of `Validations`.

to determine its result value. This is effectively interpretation of the instructions, but with the ability to handle cases where the actual input values of the instructions are unknown or unknowable.

Useless code elimination is very straightforward: instructions with unused results are eliminated, except for those with possible side-effects (i.e., calls).

## 1.5 Code Generation and Register Allocation

The SSA code generated by the compiler up to this point contains phi nodes which provide the logic to enable separate execution paths to join into one. Phi nodes do not represent an actual instruction; rather, a single register can be written to by any possible execution path, and the result will always be present for the next instruction. In order to convert phi nodes to plain instructions, it is necessary to substitute a register that will hold the result value from whichever branch executes. A new virtual register is generated to hold the phi value between the parent basic blocks and the block containing the phi node. New `mov` instructions are then added to the parent blocks immediately above the branch blocks ending them. Finally, the phi node is replaced with a `mov` instruction copying the value from the new temporary register to the target register of the phi instruction.

It was noticed that the resulting code sometimes would load a function argument into a register, only to move it to `r0` for the call. In order to avoid this, a minor optimization was added to the register allocator to attempt to reduce all such needless `mov` operations. When doing the graph coloring, the allocator is passed a mapping from registers to registers they are moved to. This is then checked for each register, and if a destination register is not in use, it will be used instead of choosing a new one. This allows the `mov` operation to then be eliminated in a later minor optimization.

This unneeded `mov` elimination optimization occurs as the code of the program is being written into the output file, by ignoring unneeded `movs`.

One noteworthy thing about the implementation of the stack-based translation is that it makes no attempt to be clever about the loads and stores involved: it potentially will spill a register containing a stack address to the stack, rather than throwing it away and recalculating it when needed.

## 1.6 Other

In ██████ this quarter, my team and I built ██████ We chose to use Scala to do so, as ██████ had some experience with it. As I became more familiar with Scala, I found I rather liked it, so I transitioned this project over to it. Thus, the early parts of the project are written in Java, typically using the Functional Java and Derive4J libraries, and the later parts are written in Scala.

Also, for amusement, I chose to use a more functional style in many places. I believe this is my first large project in such a style, so there are a lot of icky parts, as I don't really have a good feel for how to implement things cleanly yet.

The translation to ARM assembly is mostly working; I hope that the main remaining bug is a failure to add loop bodies as their own children. As far as I can tell, this leads to the register allocator deciding it can overwrite the registers that are needed by the loop during the first loop run, resulting in complete nonsense.

## 2 Analysis

These results were run on a Intel Core i7-2600K CPU running in frequency autoscaling mode, with approximately 11 Gb of free memory. Some other programs were running, but system load is believed to have been below 1. The OS is Ubuntu 16.04 LTS. LLVM 3.8 was targeted, as it was available on the testing machine. Gradle 3.4.1 running on Java 1.8 was used to build the code (and it is believed the same JVM was used to run it). A Ruby script was used to gather the reported data. It uses the `benchmark` library from the Ruby standard library to collect the timing data. The times reported are wall clock times. The graphs below are based on ten runs of the program generated by each configuration. The instruction counts are reported by the compiler in a file named `compileStats` which is overwritten after each compilation.

As previously noted in class, the LLVM results are uneven at best. Unsurprisingly, Clang generally beats everything else when all the optimizations are enabled. But the optimizations, and even the use of stack-versus register-based code, do not consistently improve performance. This is likely due to Clang's inability

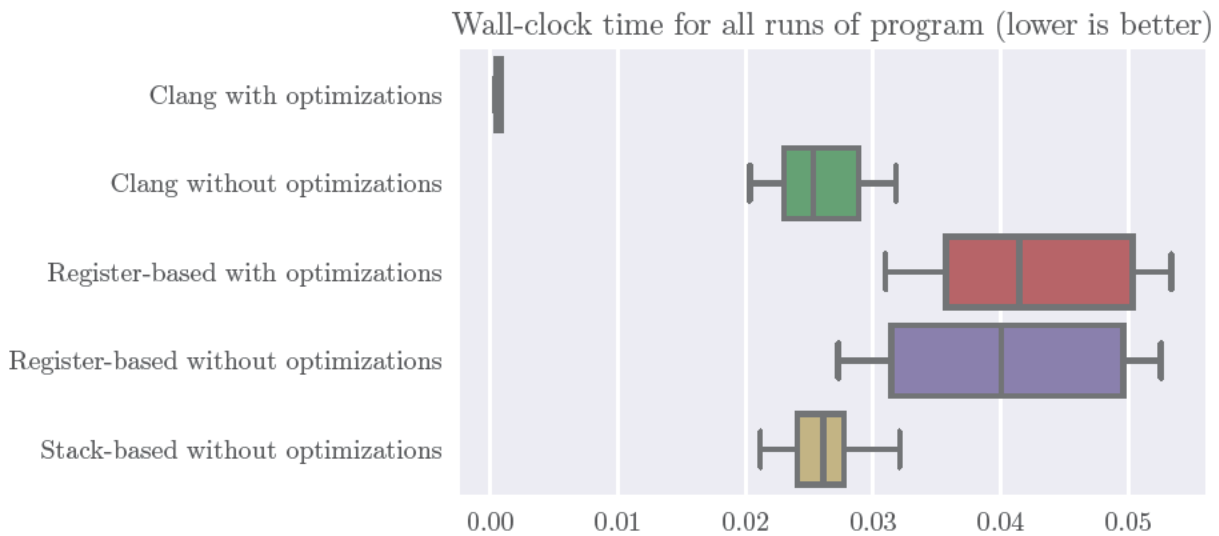
to use smart register allocation without enabling optimizations, and thus the LLVM IR ends up using the stack anyway.

It does appear that in general the register-based code is significantly shorter, with an average reduction in length of 46.53%

## 2.1 OptimizationBenchmark

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$2.621 \times 10^{-2}$	$3.376 \times 10^{-3}$
Register-based without optimizations	10	$4.020 \times 10^{-2}$	$9.788 \times 10^{-3}$
Register-based with optimizations	10	$4.248 \times 10^{-2}$	$8.491 \times 10^{-3}$
Clang without optimizations	10	$2.584 \times 10^{-2}$	$4.086 \times 10^{-3}$
Clang with optimizations	10	$6.089 \times 10^{-4}$	$1.944 \times 10^{-4}$

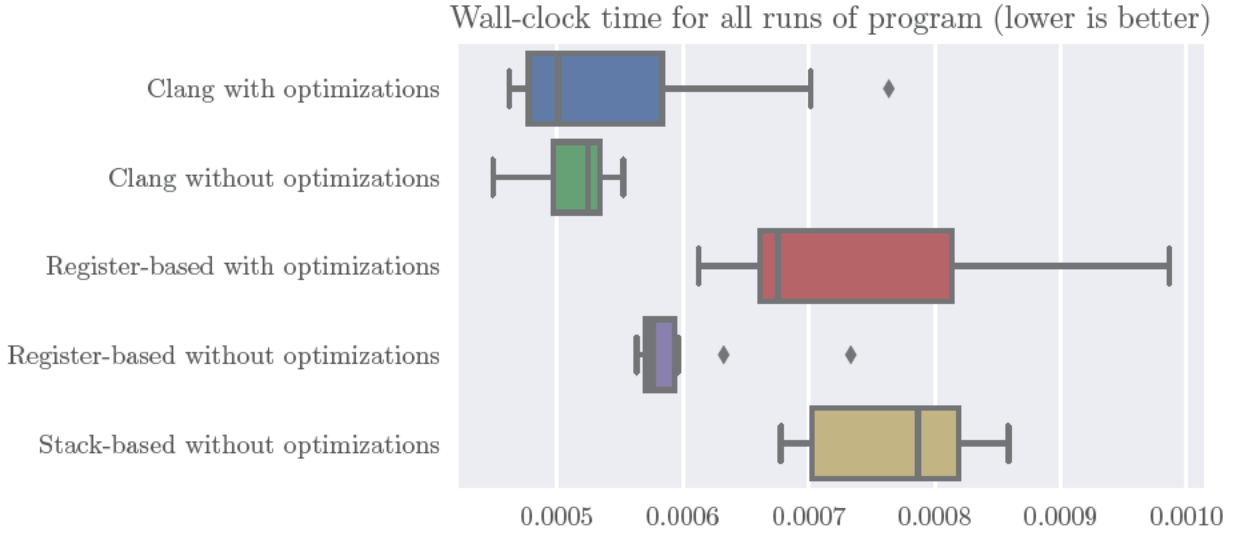
Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	974	
Register-based without optimizations	369	62.11
Register-based with optimizations	137	62.87



## 2.2 BenchMarkishTopics

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$7.687 \times 10^{-4}$	$6.722 \times 10^{-5}$
Register-based without optimizations	10	$5.973 \times 10^{-4}$	$5.188 \times 10^{-5}$
Register-based with optimizations	10	$7.390 \times 10^{-4}$	$1.246 \times 10^{-4}$
Clang without optimizations	10	$5.118 \times 10^{-4}$	$3.574 \times 10^{-5}$
Clang with optimizations	10	$5.488 \times 10^{-4}$	$1.059 \times 10^{-4}$


Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	113	
Register-based without optimizations	65	42.48
Register-based with optimizations	65	0.



### 2.3 Fibonacci

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	1.365	$2.098 \times 10^{-2}$
Register-based without optimizations	10	1.358	$2.267 \times 10^{-2}$
Register-based with optimizations	10	1.343	$3.103 \times 10^{-2}$
Clang without optimizations	10	1.444	$2.678 \times 10^{-2}$
Clang with optimizations	10	$8.591 \times 10^{-1}$	$2.740 \times 10^{-2}$


Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	26	
Register-based without optimizations	17	34.62
Register-based with optimizations	17	0.



## 2.4 GeneralFuncAndOptimize

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	1.265	$3.027 \times 10^{-2}$
Register-based without optimizations	10	1.200	$3.027 \times 10^{-2}$
Register-based with optimizations	10	1.129	$1.931 \times 10^{-2}$
Clang without optimizations	10	$7.036 \times 10^{-1}$	$2.975 \times 10^{-2}$
Clang with optimizations	10	$1.936 \times 10^{-1}$	$2.161 \times 10^{-2}$

Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	141	
Register-based without optimizations	81	42.55
Register-based with optimizations	67	17.28




## 2.5 TicTac

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$7.322 \times 10^{-4}$	$1.017 \times 10^{-4}$
Register-based without optimizations	10	$7.606 \times 10^{-4}$	$6.114 \times 10^{-5}$
Register-based with optimizations	10	$8.551 \times 10^{-4}$	$7.489 \times 10^{-5}$
Clang without optimizations	10	$5.138 \times 10^{-4}$	$6.619 \times 10^{-5}$
Clang with optimizations	10	$7.559 \times 10^{-4}$	$6.792 \times 10^{-5}$

Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	469	
Register-based without optimizations	342	27.08
Register-based with optimizations	337	1.462

Wall-clock time for all runs of program (lower is better)

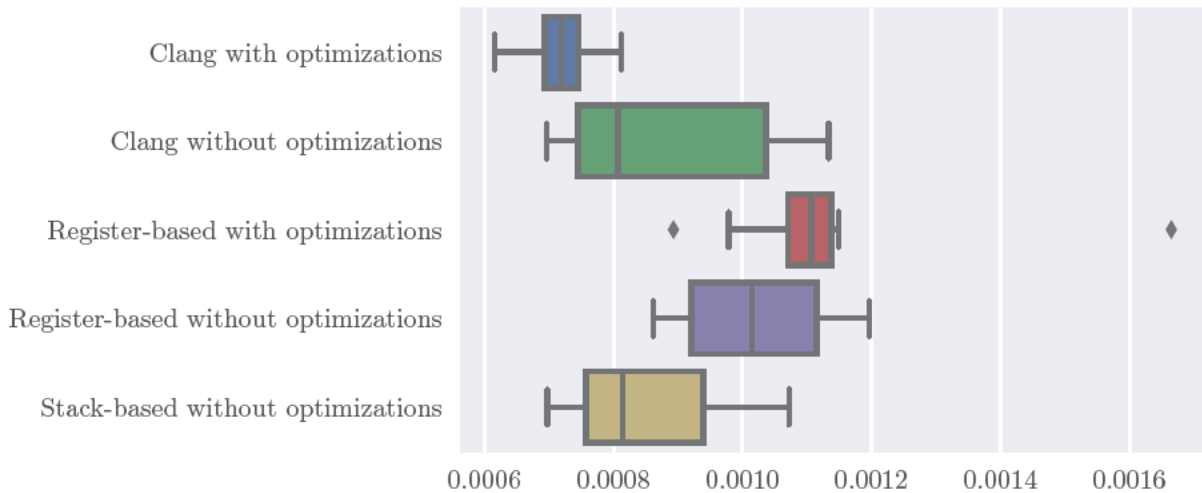


## 2.6 bert

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$8.452 \times 10^{-4}$	$1.225 \times 10^{-4}$
Register-based without optimizations	10	$1.026 \times 10^{-3}$	$1.226 \times 10^{-4}$
Register-based with optimizations	10	$1.132 \times 10^{-3}$	$2.037 \times 10^{-4}$
Clang without optimizations	10	$8.867 \times 10^{-4}$	$1.712 \times 10^{-4}$
Clang with optimizations	10	$7.214 \times 10^{-4}$	$5.677 \times 10^{-5}$

Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	617	
Register-based without optimizations	359	41.82
Register-based with optimizations	334	6.964


Wall-clock time for all runs of program (lower is better)



## 2.7 biggest

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$1.397 \times 10^{-3}$	$1.257 \times 10^{-4}$
Register-based without optimizations	10	$6.346 \times 10^{-4}$	$4.718 \times 10^{-5}$
Register-based with optimizations	10	$6.443 \times 10^{-4}$	$5.013 \times 10^{-5}$
Clang without optimizations	10	$9.310 \times 10^{-4}$	$1.876 \times 10^{-4}$
Clang with optimizations	10	$6.578 \times 10^{-4}$	$6.883 \times 10^{-5}$

Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	79	
Register-based without optimizations	42	46.84
Register-based with optimizations	41	2.381

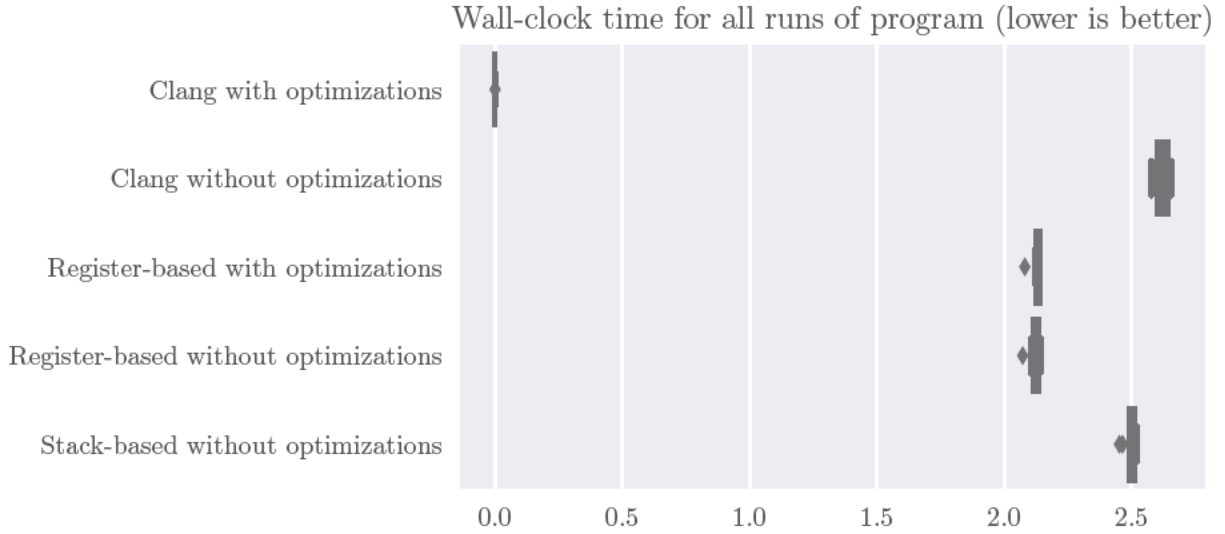


## 2.8 binaryConverter

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	2.499	$2.252 \times 10^{-2}$
Register-based without optimizations	10	2.122	$2.088 \times 10^{-2}$
Register-based with optimizations	10	2.127	$1.743 \times 10^{-2}$
Clang without optimizations	10	2.623	$2.636 \times 10^{-2}$
Clang with optimizations	10	$7.595 \times 10^{-4}$	$5.312 \times 10^{-5}$

Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	113	
Register-based without optimizations	43	61.95
Register-based with optimizations	43	0.






## 2.9 creativeBenchMarkName

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	2.738	$1.161 \times 10^{-1}$
Register-based without optimizations	10	2.503	$3.726 \times 10^{-2}$
Register-based with optimizations	10	2.442	$3.451 \times 10^{-2}$
Clang without optimizations	10	2.873	$3.374 \times 10^{-2}$
Clang with optimizations	10	$6.279 \times 10^{-4}$	$7.064 \times 10^{-5}$

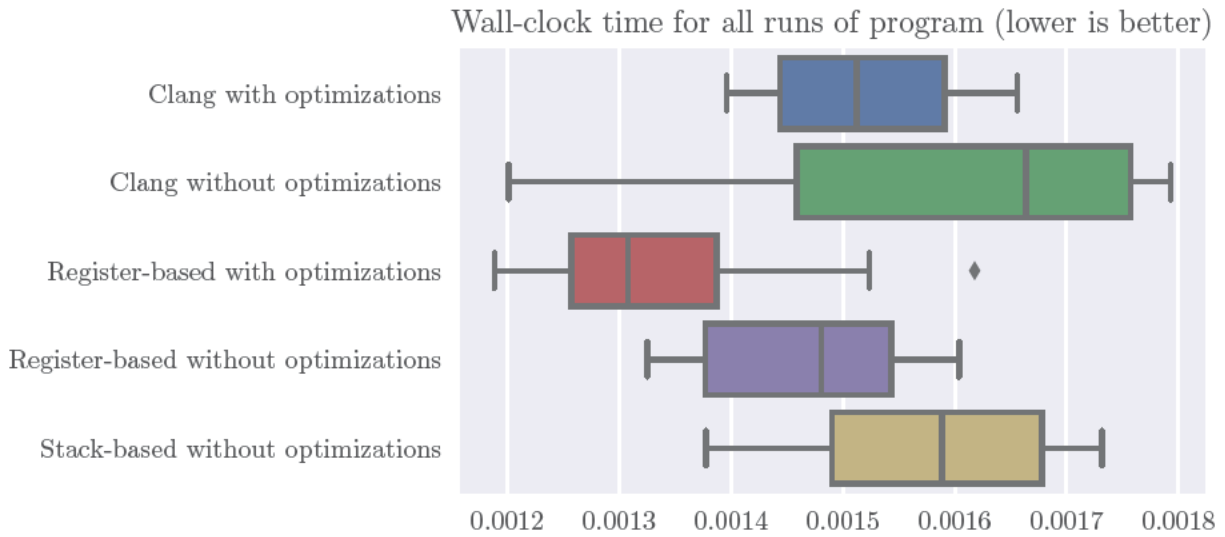
Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	230	
Register-based without optimizations	116	49.57
Register-based with optimizations	113	2.586



## 2.10 fact\_sum

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$1.576 \times 10^{-3}$	$1.260 \times 10^{-4}$
Register-based without optimizations	10	$1.463 \times 10^{-3}$	$9.963 \times 10^{-5}$
Register-based with optimizations	10	$1.346 \times 10^{-3}$	$1.340 \times 10^{-4}$
Clang without optimizations	10	$1.603 \times 10^{-3}$	$1.999 \times 10^{-4}$
Clang with optimizations	10	$1.514 \times 10^{-3}$	$9.361 \times 10^{-5}$


Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	66	
Register-based without optimizations	34	48.48
Register-based with optimizations	30	11.76



## 2.11 hailstone

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$7.530 \times 10^{-4}$	$5.787 \times 10^{-5}$
Register-based without optimizations	10	$7.821 \times 10^{-4}$	$6.118 \times 10^{-5}$
Register-based with optimizations	10	$5.299 \times 10^{-4}$	$6.103 \times 10^{-5}$
Clang without optimizations	10	$6.275 \times 10^{-4}$	$1.512 \times 10^{-4}$
Clang with optimizations	10	$4.154 \times 10^{-4}$	$4.345 \times 10^{-5}$


Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	45	
Register-based without optimizations	24	46.67
Register-based with optimizations	23	4.167



## 2.12 hanoi\_benchmark

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	3.840	$1.918 \times 10^{-2}$
Register-based without optimizations	10	3.708	$9.307 \times 10^{-2}$
Register-based with optimizations	10	3.851	$6.220 \times 10^{-2}$
Clang without optimizations	10	4.100	$1.789 \times 10^{-2}$
Clang with optimizations	10	1.593	$3.334 \times 10^{-2}$

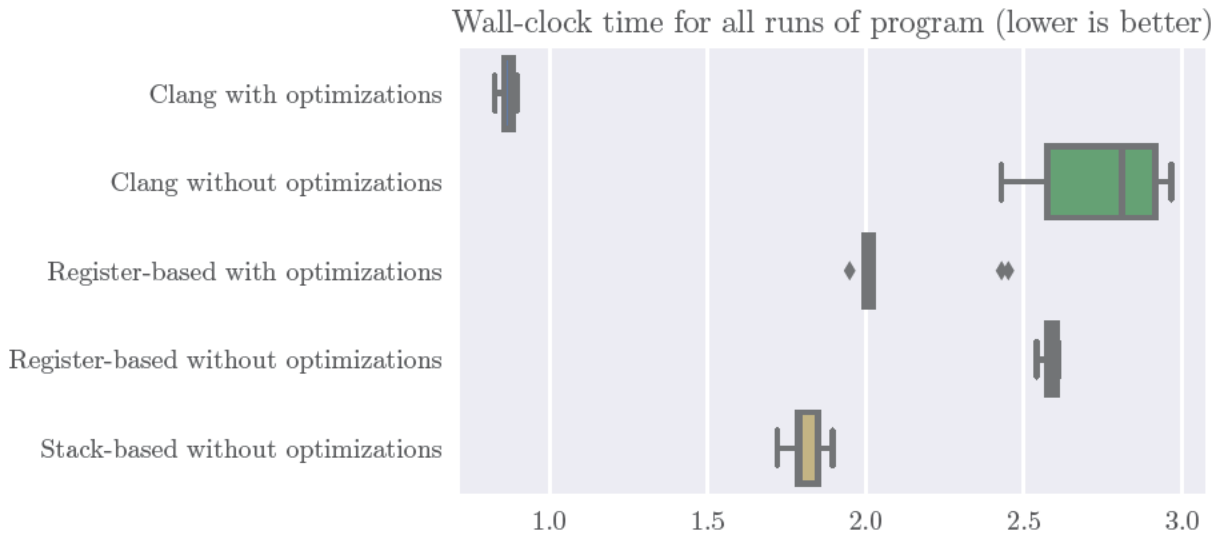
Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	194	
Register-based without optimizations	126	35.05
Register-based with optimizations	126	0.



### 2.13 killerBubbles

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	1.809	$5.489 \times 10^{-2}$
Register-based without optimizations	10	2.584	$2.479 \times 10^{-2}$
Register-based with optimizations	10	2.088	$1.868 \times 10^{-1}$
Clang without optimizations	10	2.747	$1.941 \times 10^{-1}$
Clang with optimizations	10	$8.693 \times 10^{-1}$	$2.463 \times 10^{-2}$

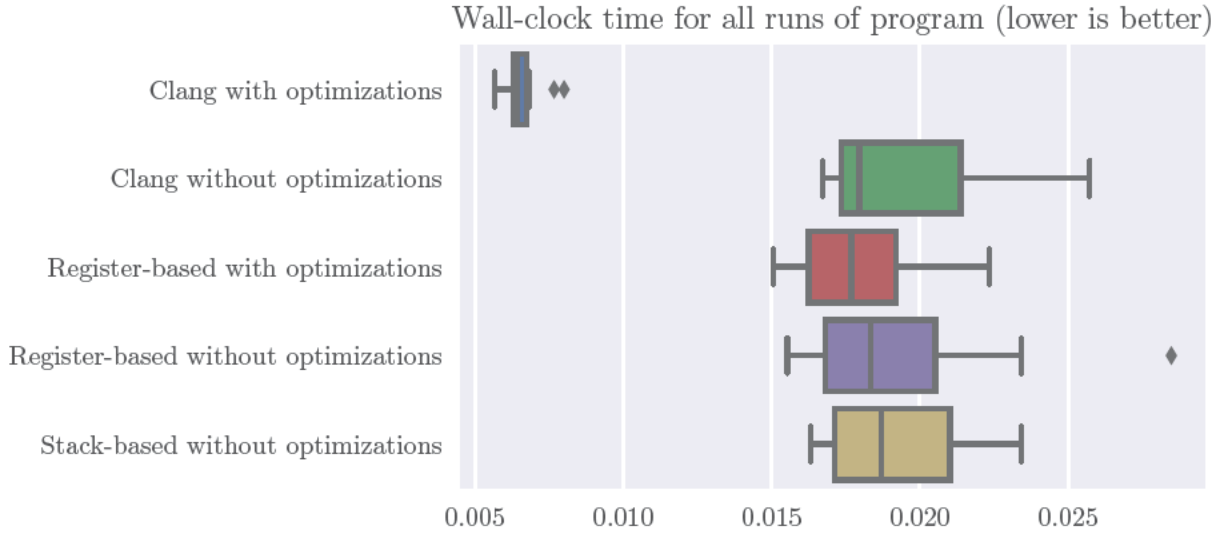
Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	185	
Register-based without optimizations	96	48.11
Register-based with optimizations	93	3.125



### 2.14 mile1

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$1.925 \times 10^{-2}$	$2.508 \times 10^{-3}$
Register-based without optimizations	10	$1.944 \times 10^{-2}$	$3.949 \times 10^{-3}$
Register-based with optimizations	10	$1.811 \times 10^{-2}$	$2.408 \times 10^{-3}$
Clang without optimizations	10	$1.935 \times 10^{-2}$	$3.076 \times 10^{-3}$
Clang with optimizations	10	$6.601 \times 10^{-3}$	$7.293 \times 10^{-4}$

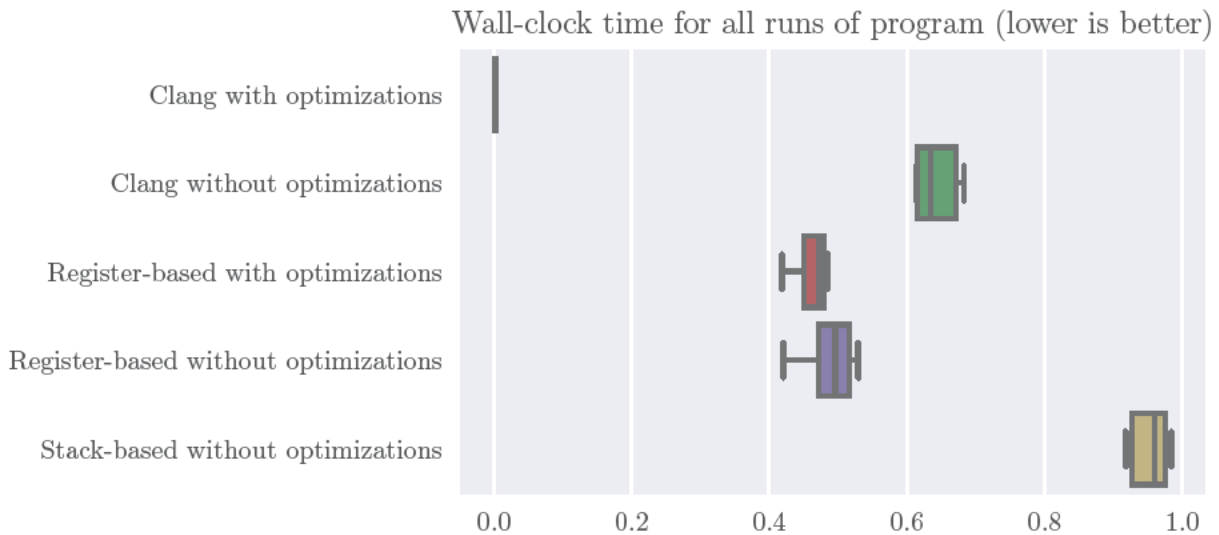
Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	70	
Register-based without optimizations	33	52.86
Register-based with optimizations	31	6.061



## 2.15 mixed

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$9.518 \times 10^{-1}$	$2.654 \times 10^{-2}$
Register-based without optimizations	10	$4.903 \times 10^{-1}$	$3.368 \times 10^{-2}$
Register-based with optimizations	10	$4.607 \times 10^{-1}$	$2.470 \times 10^{-2}$
Clang without optimizations	10	$6.425 \times 10^{-1}$	$2.972 \times 10^{-2}$
Clang with optimizations	10	$2.013 \times 10^{-3}$	$2.634 \times 10^{-4}$

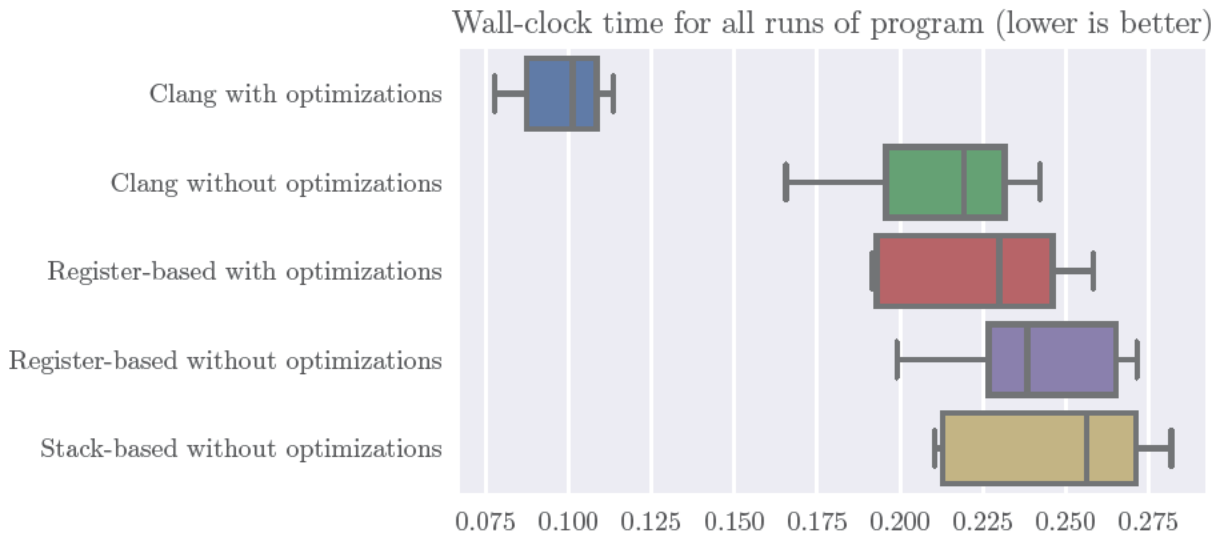
Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	201	
Register-based without optimizations	116	42.29
Register-based with optimizations	96	17.24



## 2.16 primes

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$2.458 \times 10^{-1}$	$3.041 \times 10^{-2}$
Register-based without optimizations	10	$2.416 \times 10^{-1}$	$2.442 \times 10^{-2}$
Register-based with optimizations	10	$2.228 \times 10^{-1}$	$2.762 \times 10^{-2}$
Clang without optimizations	10	$2.127 \times 10^{-1}$	$2.498 \times 10^{-2}$
Clang with optimizations	10	$9.811 \times 10^{-2}$	$1.301 \times 10^{-2}$


Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	87	
Register-based without optimizations	38	56.32
Register-based with optimizations	38	0.



## 2.17 programBreaker

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$2.777 \times 10^{-3}$	$1.405 \times 10^{-4}$
Register-based without optimizations	10	$2.819 \times 10^{-3}$	$2.282 \times 10^{-4}$
Register-based with optimizations	10	$2.760 \times 10^{-3}$	$2.806 \times 10^{-4}$
Clang without optimizations	10	$2.447 \times 10^{-3}$	$4.070 \times 10^{-4}$
Clang with optimizations	10	$1.507 \times 10^{-3}$	$2.898 \times 10^{-4}$

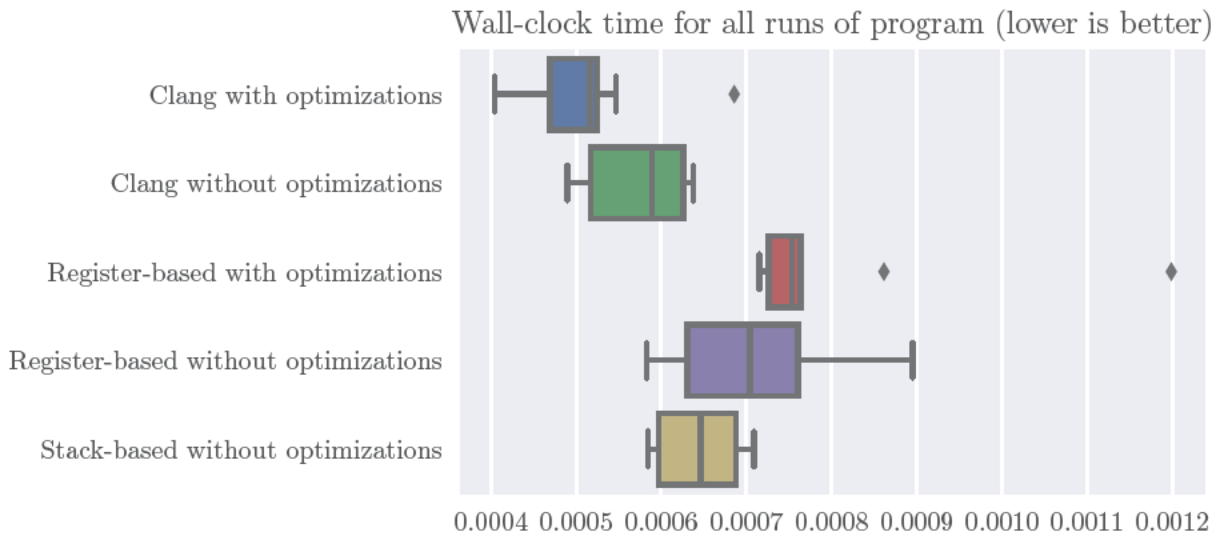
Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	69	
Register-based without optimizations	33	52.17
Register-based with optimizations	31	6.061



## 2.18 stats

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$6.438 \times 10^{-4}$	$4.921 \times 10^{-5}$
Register-based without optimizations	10	$7.148 \times 10^{-4}$	$1.017 \times 10^{-4}$
Register-based with optimizations	10	$7.987 \times 10^{-4}$	$1.468 \times 10^{-4}$
Clang without optimizations	10	$5.725 \times 10^{-4}$	$5.830 \times 10^{-5}$
Clang with optimizations	10	$5.117 \times 10^{-4}$	$7.439 \times 10^{-5}$

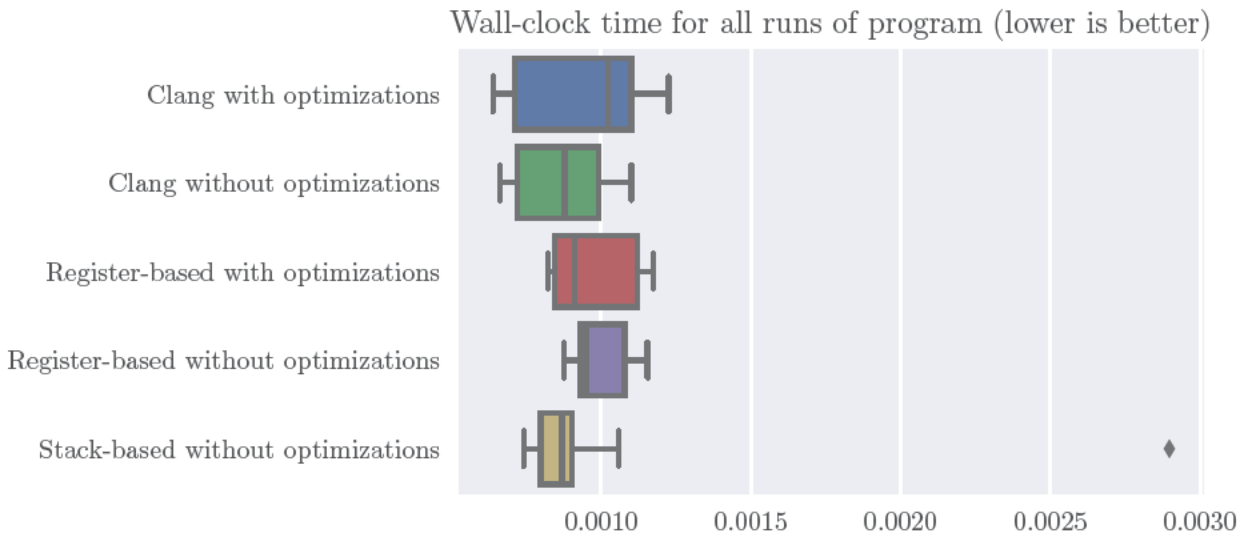
Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	232	
Register-based without optimizations	101	56.47
Register-based with optimizations	101	0.



## 2.19 uncreativeBenchmark

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$1.059 \times 10^{-3}$	$6.522 \times 10^{-4}$
Register-based without optimizations	10	$9.949 \times 10^{-4}$	$9.713 \times 10^{-5}$
Register-based with optimizations	10	$9.727 \times 10^{-4}$	$1.467 \times 10^{-4}$
Clang without optimizations	10	$8.684 \times 10^{-4}$	$1.588 \times 10^{-4}$
Clang with optimizations	10	$9.309 \times 10^{-4}$	$2.235 \times 10^{-4}$

Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	298	
Register-based without optimizations	206	30.87
Register-based with optimizations	206	0.

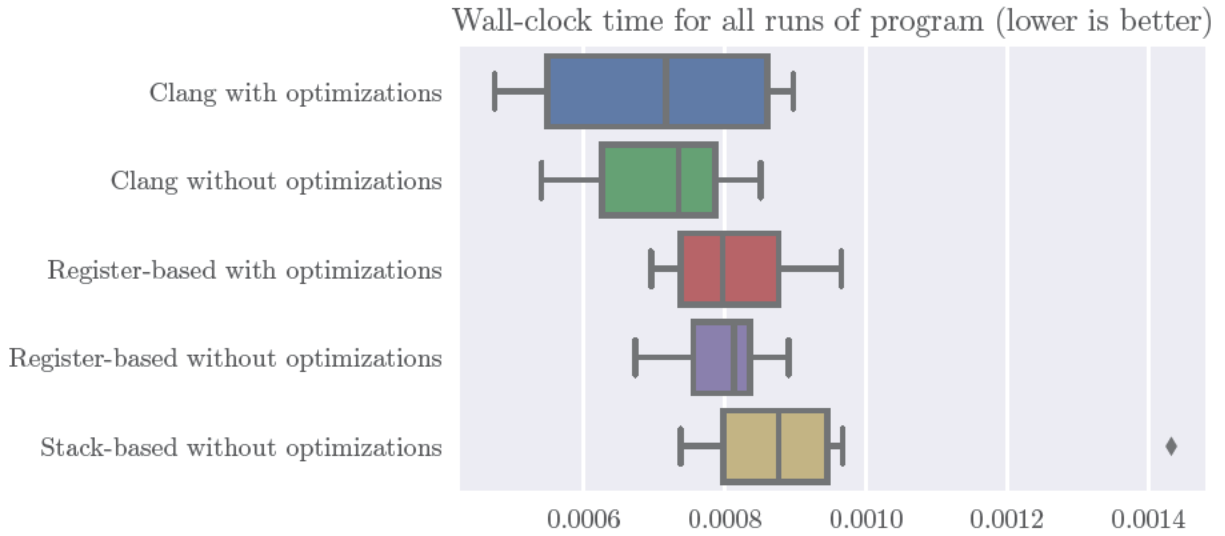


## 2.20 wasteOfCycles

Type	Runs	Average Time (s)	Sample Std. Dev. (s)
Stack-based without optimizations	10	$9.152 \times 10^{-4}$	$1.993 \times 10^{-4}$
Register-based without optimizations	10	$7.955 \times 10^{-4}$	$6.660 \times 10^{-5}$
Register-based with optimizations	10	$8.087 \times 10^{-4}$	$9.029 \times 10^{-5}$
Clang without optimizations	10	$7.115 \times 10^{-4}$	$1.044 \times 10^{-4}$
Clang with optimizations	10	$7.020 \times 10^{-4}$	$1.742 \times 10^{-4}$

Type	Instruction Count	Percent improvement over previous
Stack-based without optimizations	42	
Register-based without optimizations	20	52.38
Register-based with optimizations	20	0.





## References

- [1] Matthias Braun et al. “Simple and Efficient Construction of Static Single Assignment Form.” In: (). Ed. by Ranjit Jhala and Koen De Bosschere. URL: <http://dblp.uni-trier.de/db/conf/cc/cc2013.html#BraunBHLMZ13>; [http://dx.doi.org/10.1007/978-3-642-37051-9\\_6](http://dx.doi.org/10.1007/978-3-642-37051-9_6).
- [2] Keith D. Cooper and Linda Torczon. *Engineering a Compiler*. 2nd ed. Morgan Kaufmann, Jan. 12, 2004, p. 800. ISBN: 1-55860-699-8.