

Writing Tucan

Writing an Optimizing Compiler in Rust

Thorsten Ball

November 2024

thorstenball.com

2016

THORSTEN BALL

WRITING AN
INTERPRETER
IN GO

2018

THORSTEN BALL

WRITING A

COMPILER

IN GO

2020

Tucan

- Optimizing compiler in Rust
- 18k lines of code, 0 third-party dependencies
- IR is a Control-Flow Graph in SSA form
- Optimizations
 - Dead Code Elimination
 - Sparse Conditional Constant Propagation
 - Dominator-based Value Numbering
 - Useless control-flow elimination
- Liveness analysis to compute live sets
- Linear scan register allocator
- x86-64 code generation
- Immix heap with GC

```
tucan master
tucan
├── .direnv
├── assets
├── samples
├── target
├── tucan_runtime
│   └── src
│       ├── heap.rs
│       ├── large_object_space.rs
│       ├── lib.rs
│       ├── object.rs
│       ├── raw.rs
│       └── space.rs
├── Cargo.toml
├── tucanc
│   ├── html
│   └── src
│       ├── ast
│       ├── codegen
│       ├── ir
│       ├── lexer
│       ├── middle
│       ├── optim
│       ├── parser
│       ├── regalloc
│       ├── assembly_runner.rs
│       ├── ast.rs
│       ├── builtins.rs
│       ├── check_ir.rs
│       ├── codegen.rs
│       └── compiler.rs
│       ├── debug.rs
│       ├── dumper.rs
│       ├── ir.rs
│       ├── lexer.rs
│       ├── lib.rs
│       ├── main.rs
│       ├── middle.rs
│       ├── optim.rs
│       ├── parser.rs
│       ├── regalloc.rs
│       ├── runtime_calls.rs
│       ├── runtime_tests.rs
│       ├── test_helpers.rs
│       └── test_macros.rs
├── tests
├── Cargo.lock
├── Cargo.toml
└── ...

tucanc/src/compiler.rs
pub struct Compiler
119     String { name: &'a str, code: &'a str },
120 }
121
122  ⚡ pub struct Compiler<'a> {      ⚡ Thorsten Ball, 2 years ago
123     use_regalloc: bool,
124     use_optim: bool,
125
126     dump_funk: Option<String>,
127     dumper: Option<Dumper>,
128
129     debug: bool,
130
131     input: Option<CompilerInput<'a>>,
132 }
133
134 impl Default for Compiler<'_> {
135     fn default() -> Self {
136         let dump_funk: Option<String> = match std::env::var(key: "TUCAN_DUMP") {
137             Ok(funk: String) if !funk.is_empty() => Some(funk),
138             _ => None,
139         };
140
141         Compiler {
142             use_regalloc: true,
143             use_optim: true,
144             debug: false,
145             input: None,
146             dump_funk,
147             dumper: None,
148         }
149     }
150 }
151
152 impl<'a> Compiler<'a> {
153     pub fn new() -> Self {
154         Self::default()
155     }
156
157     pub fn for_file(filename: &'a str) -> Self {
158         Compiler {
159             input: Some(CompilerInput::File { filename }),
160             ..Self::default()
161         }
162     }
163
164     pub fn for_string(name: &'a str, code: &'a str) -> Self {
165         Compiler {
166             input: Some(CompilerInput::String { name, code }),
167             ..Self::default()
168         }
169     }
170
171     pub fn regalloc(&mut self, use_regalloc: bool) -> &mut Self {
172         self.use_regalloc = use_regalloc;
173         self
174     }
175
176     pub fn optimize(&mut self, use_optim: bool) -> &mut Self {
177         self.use_optim = use_optim;
178         self
179     }
180
181     pub fn dump(&mut self, funk_name: &'a str) -> &mut Self {
```

No:

- Language design
- Fancy type systems
- Novel implementation strategies
- Building the perfect system

Yes:

- Figuring out how an optimizing compiler works

```
funk my_function(a: int, b: int) → int {
    let c = a + b + 18 - 5 + 3 + 2;
    let d = 18 - 5 + 3;
    let e = a + b + 30;
    let f = 100 - a - b + 3 + 2;
    let g = c + d + e; // dead code
    let h = a + b + c + e + f;

    let k = 50;

    let result = 0;
    for (let i = 0; i < 10; i = i + 1) {
        result = result + 5;
        if (i < a) {
            if (k ≥ d) {
                double_print(i);
            } else {
                print_num(i);
            }
        } else {
            result = result + c;
        }
        for (let j = 5; j > 0; j = j - 1) {
            result = result - j;
        }
        print_num(result);
        print_string("-");
    }
    return result;
}
```

Lessons Learned

The bird's name is Toucan

(But, hey: in German it's "Tukan")



Lexing & Parsing

Dora

Dominik Inführ's JIT compiler

The screenshot shows the GitHub repository page for 'dora' by 'dinfuehr'. The repository is public and has 11 watchers. The current branch is 'main', with 5 other branches and 0 tags. The repository contains 11 folders, each with a commit history and a date. The most recent commit is 'frontend: Switch to index traits' by 'dinfuehr' on 'yesterday'.

Folder	Commit	Time
.github/workflows	github: Enable release build tests on Windows again	5 months ago
.vscode	cannon: support enum values	4 years ago
bench	frontend: Require named arguments for classes with nam...	2 weeks ago
dora-asm	x64: Add more AVX instructions to boots	3 months ago
dora-bytecode	frontend: Allow unnamed access to classes and structs	2 weeks ago
dora-frontend	frontend: Switch to index traits	yesterday
dora-language-server	ls: Update dependencies	2 days ago
dora-parser	frontend: Start implementing pattern with named fields	last week
dora-runtime	frontend: Switch to index traits	yesterday
dora-sema-fuzzer	driver: Rename cmd.rs to flags.rs	2 weeks ago
dora	frontend: Start report warnings when errors aren't present.	2 weeks ago

Lesson:

Copy & Pasting & Deleting

... can kickstart projects

```
15 pub struct Parser<'a> {
16     lexer: Lexer,
17     token: Token,
18     peek_token: Token,
19     ast: &'a mut Ast,
20
21     id_counter: RefCell<usize>,
22     last_end: Option<u32>,
23 }
24
25 type ExprResult = Result<Box<Expr>, ParseErrorAndPos>;
26 type StmtResult = Result<Stmt, ParseErrorAndPos>;
27 type FunkResult = Result<Funk, ParseErrorAndPos>;
28
29 impl<'a> Parser<'a> {
30     pub fn new(reader: Reader, ast: &'a mut Ast) -> Parser<'a> {
31         let token: Token = Token::new(tok: TokenKind::Eof, pos: Position::r
32         let peek_token: Token = Token::new(tok: TokenKind::Eof, pos: Positi
33         let lexer: Lexer = Lexer::new(reader);
34
35         Parser {
36             lexer,
37             token,
38             peek_token,
39             ast,
40             id_counter: RefCell::new(1),
41             last_end: Some(0),
42         }
43     }
44
45     fn init(&mut self) -> Result<(), ParseErrorAndPos> {
46         self.advance_token()?;
47         self.advance_token()?;
48
49         Ok(())
50     }
51
52     fn read_token(&mut self) -> Result<Token, ParseErrorAndPos> {
53         self.last_end = if self.token.span.is_valid() {
54             Some(self.token.span.end())
55         } else {
56             None
57         };
58         let peek_token: Token = self.lexer.read_token()?;
59         let token: Token = mem::replace(dest: &mut self.peek_token, src: pe
60         Ok(mem::replace(dest: &mut self.token, src: token))
61     }
62 }
```

SSA

Static Single-Assignment

Original

```
a := b + c
b := c + 1
d := b + c
a := a + 1
e := a + b
```

SSA

```
a1 := b1 + c1
b2 := c1 + 1
d1 := b2 + c1
a2 := a1 + 1
e1 := a2 + b2
```

Original

```
if B then
  a := b
else
  a := c
end
... a ...
```

SSA

```
if B then
  a1 := b
else
  a2 := c
End
a3 :=  $\Phi(a_1, a_2)$ 
... a3 ...
```


All the cool kids have it

- Rust
- LLVM
- Go
- LuaJIT
- PyPy
- WebKit

Should be easy, right?

Lesson:

SSA ain't SSA

~/drive/notes/Tucan - Constructing SSA for Tucan from CFG.md

```
1 # Constructing SSA for Tucan from CFG
2
3 ## Current status
4
5 Okay, let's recap.
6
7 Right now we a program in the form of a control-flow graph (CFG) with each
8 instruction consisting of a target, and an instruction kind.
9
10 What we want is to turn this program into SSA form. SSA stands for Static Single
11 Assignment, another way to represent programs.
12
13 Why convert the program to SSA form? Because once a program is in SSA form it's
14 easier to run certain optimization passes over the program.
15
16 > If LLVM had not used the SSA form, we would need to run a separate data flow
17 > analysis to compute the use-def chains, which are mandatory for classical
18 > optimizations such as constant propagation and common subexpression elimination.
19
20 Example: https://hub.packtpub.com/introducing-llvm-intermediate-representation/
21
22 ## Resources
23
24 <!-- Links up here because it's easier to add to the bottom of a section without having to move the links all the time -->
25
26 [ssabook]: http://ssabook.gforge.inria.fr/latest/book.pdf
27 [simple]: https://link.springer.com/content/pdf/10.1007/3-540-46423-9\_8.pdf
28 [simpleandefficient]: https://pp.info.uni-karlsruhe.de/uploads/publikationen/braun13cc.pdf
29 [denkerpresentation]: https://marcusdenker.de/talks/08CC/08IntroSSA.pdf
30 [coirbuilder]: https://github.com/rsms/co/blob/master/src/ir/builder.ts#L1218
31 [firmgraphs]: https://pp.ipd.kit.edu/firm/GraphSnippets.html
32 [trufflegraphs]: https://chriseaton.com/truffleruby/basic-truffle-graphs/
33 [graalgraphs]: https://chriseaton.com/truffleruby/basic-graal-graphs/
34 [cliffclicktalk]: https://www.youtube.com/watch?v=98lt45Aj8mo
35 [gocompilerinternals]: https://eli.thegreenplace.net/2019/go-compiler-internals-adding-a-new-statement-to-go-part-2/
36 [llvmpresentation]: https://llvm.org/devmtg/2017-06/1-Davis-Chisnall-LLVM-2017.pdf
37 [simplefastdom]: https://www.cs.rice.edu/~keith/EMBED/dom.pdf
38 [waddledominator]: https://github.com/efritz/waddle/blob/master/src/main/scala/waddle/dominators/DominatorUtil.scala#L92
39 [dominancefrontierslecture]: http://pages.cs.wisc.edu/~fischer/cs701.f05/lectures/Lecture22.pdf
40
41 ### Books
42
43 Apparently Appel's book and the Cooper book contain algorithms for SSA.
44
45 ### Papers
46
47 There's the [Static Single Assignment Book][ssabook].
48
49 Then there's the paper [Simple Generation of Static Single-Assignment Form][simple]:
50
51 > Abstract. The static single-assignment (SSA) form of a program pro-
52 > vides data flow information in a form which makes some compiler opti-
53 > mizations easy to perform. In this paper we present a new, simple method
54 > for converting to SSA form, which produces correct solutions for nonre-
55 > ducible control-flow graphs, and produces minimal solutions for reducible
56 > ones. Our timing results show that, despite its simplicity, our algorithm
57 > is competitive with more established techniques.
58
59 And the simple and efficient version: [Simple and Efficient Construction of Static Single Assignment Form][simpleandefficient]
60
61 > Abstract. We present a simple SSA construction algorithm, which al-
62 > lows direct translation from an abstract syntax tree or bytecode into an
63 > SSA-based intermediate representation. The algorithm requires no prior
64 > analysis and ensures that even during construction the intermediate rep-
65 > resentation is in SSA form. This allows the application of SSA-based op-
66 > timizations during construction. After completion, the intermediate rep-
67 > resentation is in minimal and pruned SSA form. In spite of its simplicity,
68 > the runtime of our algorithm is on par with Cytron et al.'s algorithm.
69
70 This presentation also has a good overview of an SSA algorithm: [Markus Denker Presentation][denkerpresentation]
71
72 ### Code
73
74 #### Go
75
```



Thorsten Ball  @thorstenball · Dec 31, 2020



Thinking about SSA is melting my brain.



Slava Egorov
@mraleph



Just think about it as a data flow graph rather than anything else

3:54 PM · Dec 31, 2020



Rasmus Andersson ✓

@rsms



You could simply use the memory address of a node to identify it, rather than a name or managed numeric identifier (since a value never changes.)

4:08 AM · Jan 1, 2021

Lesson:

You have to read the code

72 **### Code**

73

74 **#### Go**

75

76 The ****Go source code**** also has conversion from AST to SSA and this is the end of the file:
<https://sourcegraph.com/github.com/golang/go@master/-/blob/src/cmd/compile/internal/go#L1106-1113>

77

78 Here is the part of the code that handles assignment:

<https://sourcegraph.com/github.com/golang/go@95ce805d14642a8e8e40fe1f8f50b9b5b/src/cmd/compile/internal/gc/ssa.go#L1254>

79

80 This is the most important file: ``src/cmd/compile/internal/gc/ssa.go``

81

82 **#### Co**

83

84 @rsms' IR builder for his ****Co**** language creates a CFG in SSA form in this file:
[IRBuilder](#) [[coirbuilder](#)] Pure gold!

85

86 **#### BinaryAnalysisPlatform**

87

88 Pass to transform to pruned SSA:

<https://sourcegraph.com/github.com/BinaryAnalysisPlatform/bap/-/blob/lib/bap/ssa.ml>

89

And more information here:

<https://sourcegraph.com/github.com/BinaryAnalysisPlatform/bap/-/blob/lib/bap/>

90

91 Linked to by Rijnard: <https://twitter.com/rvtond/status/1344008280171970560>

92

93 **#### irhydra**

94

95 @mraleph's irhydra:

<https://github.com/mraleph/irhydra/blob/master/saga/lib/src/flow/ssa.dart>

96

97 [[This comment](#)] [<https://twitter.com/mraleph/status/1343907821495181313?s=20>] b

98 interesting:

99

> Things really depend on how your IR looks like. For example in Dart VM each

> instruction is an allocated object so its identity (address) is its SSA name.

> the question of computing SSA form is just the question of placing phi's and

> computing their arguments

100

101 The ``ssa.dart`` file and the definitions in [``node.dart``]

(<https://github.com/mraleph/irhydra/blob/master/saga/lib/src/flow/node.dart>)

102 interesting.

103

104

105

Lesson:
You have to
read the
slides

08IntroSSA.ppt 26 / 52 100%

SSA

Dominance and SSA C

© Marcus Denker

The image shows a presentation slide titled "08IntroSSA.ppt" at slide 26 of 52. The slide content is partially obscured by a dark overlay on the left, which contains a vertical list of slide numbers from 21 to 26. Slide 26 is highlighted in blue. The main slide content on the right shows the title "SSA" and "Dominance and SSA C". Below the title is a control flow graph with nodes 1 through 13. Node 1 is the root. Node 2 is a child of 1. Node 3 is a child of 2 and has a self-loop. Node 4 is a child of 3. Node 5 is a child of 1 and is the root of a gray-shaded region containing nodes 6, 7, and 8. Node 6 and 7 are children of 5, and node 8 is a child of both 6 and 7. Node 13 is a child of both 4 and 8. The graph is labeled "© Marcus Denker".

SSA-based Compiler Design

1 / 412 | 100%

1

Lots of authors


Static Single Assignment Book

Friday 8th June, 2018
16:58

1

2

3



Lesson:

Papers aren't meant to teach

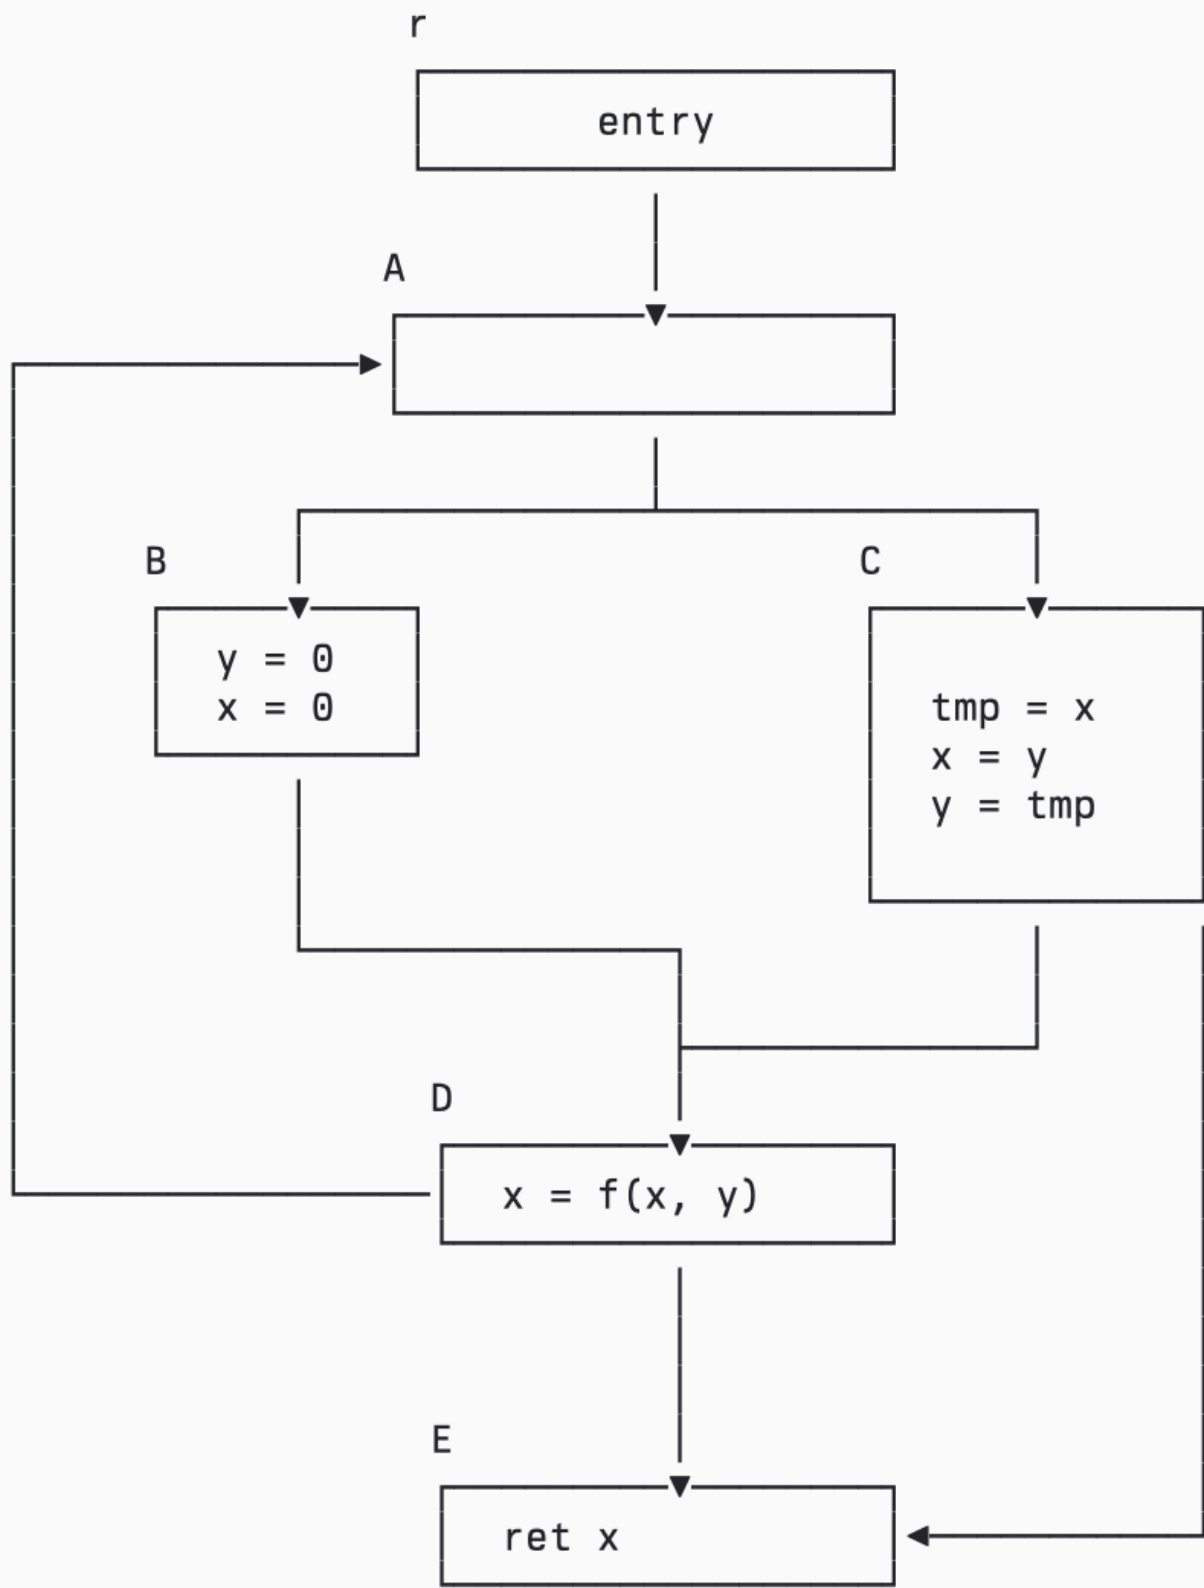
Textbooks aren't meant to teach

They are meant to share knowledge

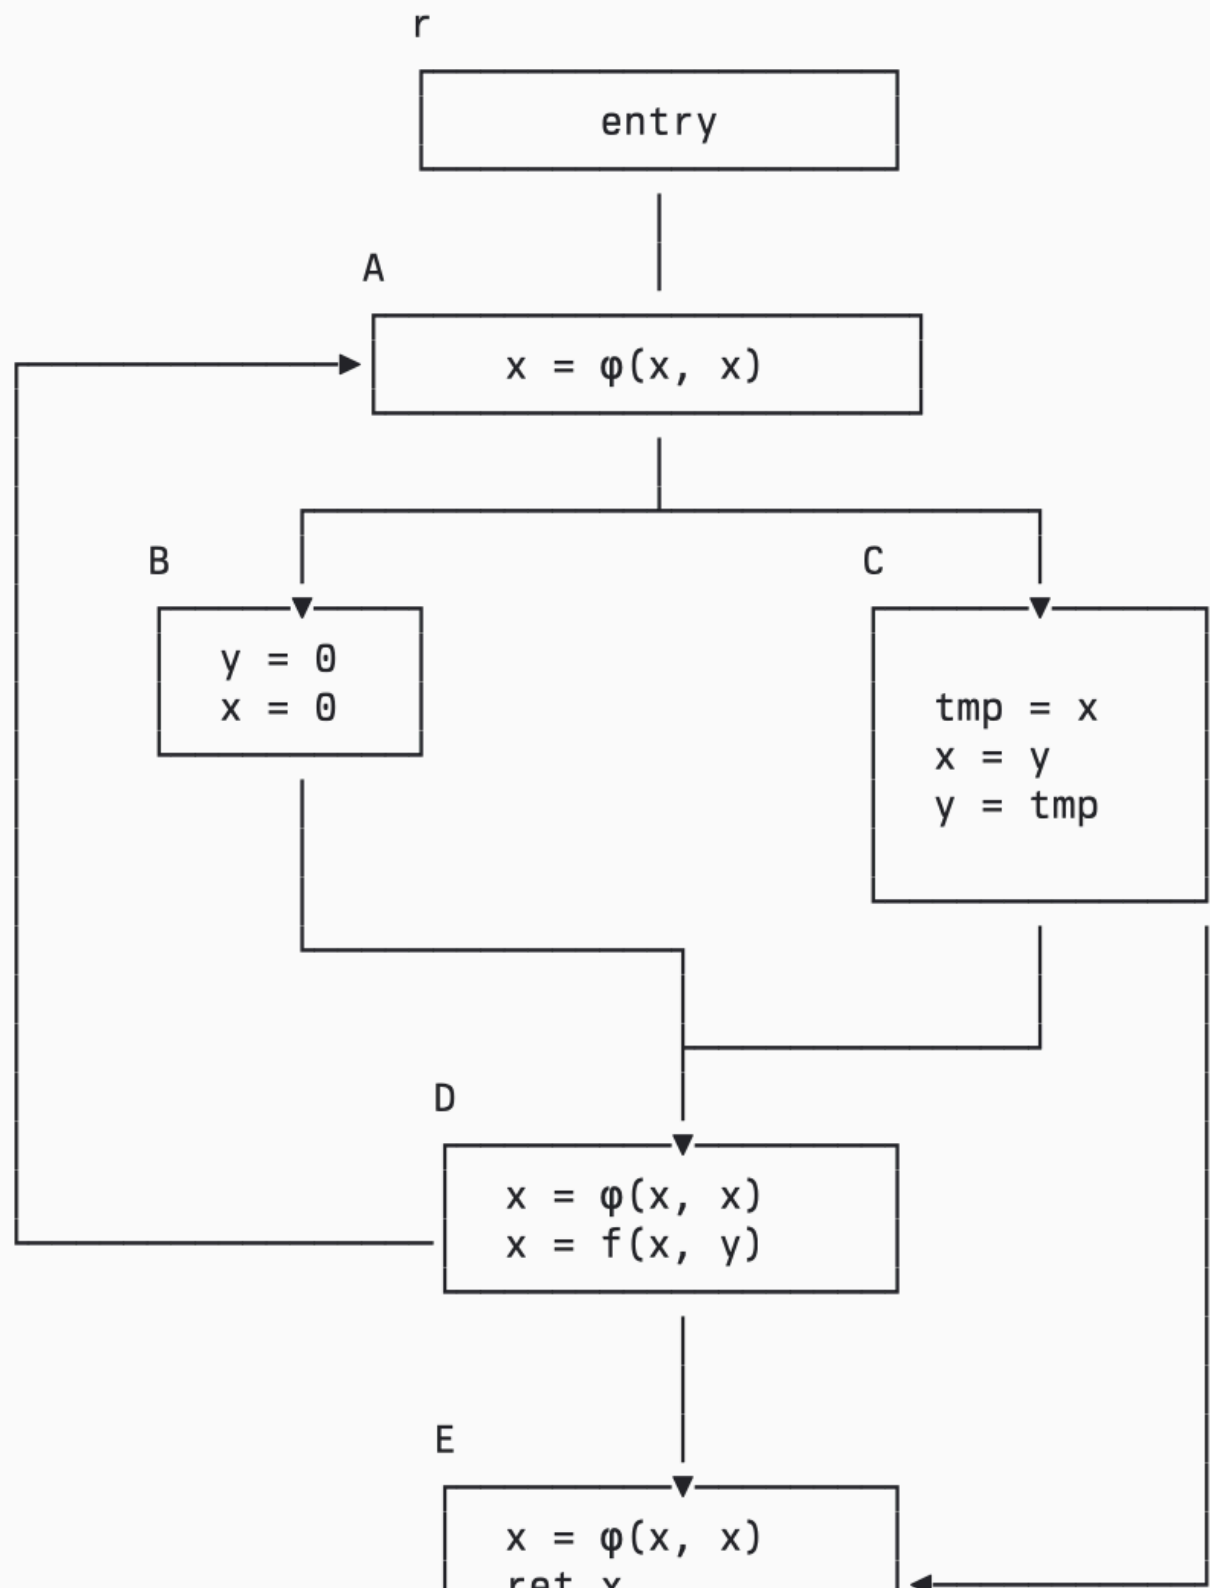
Lesson:

ASCII graphs are awesome

ϕ -function insertion



Before



After


```

34 // Phase 1 of SSA construction:  $\phi$ -function insertion
35 //
36 // Based on "3.1.2  $\phi$ -function insertion" in "Standard Construction and
37 // Destruction Algorithms" by J. Singer and F. Rastello in the SSA book.
38 pub fn insert_phi_functions(funk: &mut BlockFunk) {
39     let mut phis: HashMap<BlockId, Vec<Instruction>> = HashMap::new();
40     let assigns: HashMap<Place, HashSet<Var>> = funk.build_assigns();
41
42     for (place: Place, defs: HashSet<Var>) in assigns {
43         // Skipping the vars that are only written-to in a single block.
44         if defs.len() == 1 {
45             continue;
46         }
47
48         let mut done: HashSet<BlockId> = HashSet::new();
49         let mut work_list: Vec<&Block> = Vec::new();
50
51         let new_var: impl Fn(&BlockId) -> Var = |id: &BlockId| Var(*id, place);
52
53         let mut def_blocks: HashSet<BlockId> = HashSet::new();
54         for def_var: &Var in defs.iter() {
55             work_list.push(&funk[def_var.0]);
56             def_blocks.insert(def_var.0);
57         }
58
59         while let Some(block: &Block) = work_list.pop() {
60             for df_block_id: BlockId in block.dom_frontier.iter().cloned() {

```

Code Generation

Relatively straightforward

- Emit x86 ASM
- Assemble with GCC

```
// movq
pub fn emit_movq_const(&mut self, constant: &ir::Constant, target: RegOrOffset) -> EmitResult {
    match constant {
        ir::Constant::Int(i: &i64) => writeasm!(self, "\tmovq ${i}, {target}"),
        ir::Constant::Bool(true) => writeasm!(self, "\tmovq ${{}}, {{}}, &1, target),",
        ir::Constant::Bool(false) => writeasm!(self, "\tmovq ${{}}, {{}}, &0, target),",
        ir::Constant::String(s: &String) => unlowered_string!(s),
    }
}

pub fn emit_movq(&mut self, origin: RegOrOffset, target: RegOrOffset) -> EmitResult {
    writeasm!(self, "\tmovq {origin}, {target}")
}

// subq/addq
pub fn emit_subq_ir(&mut self, integer: &u64, reg: Reg) -> EmitResult {
    writeasm!(self, "\tsubq ${integer}, {reg}")
}

pub fn emit_addq_ir(&mut self, integer: &u64, reg: Reg) -> EmitResult {
    writeasm!(self, "\taddq ${integer}, {reg}")
}

pub fn emit_addq(&mut self, origin: RegOrOffset, reg: impl Into<RegOrOffset>) -> EmitResult {
    writeasm!(self, "\taddq {{}}, {{}}, origin, reg.into())
}

pub fn emit_subq(&mut self, origin: RegOrOffset, reg: impl Into<RegOrOffset>) -> EmitResult {
    writeasm!(self, "\tsubq {{}}, {{}}, origin, reg.into())
}

// cmp, jump, etc.
pub fn emit_cmpq_to_reg(&mut self, origin: RegOrOffset, reg: Reg) -> EmitResult {
    writeasm!(self, "\tcmpq {origin}, {reg}")
}

pub fn emit_cmpq_const(&mut self, constant: &ir::Constant, target: RegOrOffset) -> EmitResult {
    match constant {
        ir::Constant::String(s: &String) => unlowered_string!(s),
        ir::Constant::Bool(true) => writeasm!(self, "\tcmpq ${{}}, {{}}, &1, target),",
        ir::Constant::Bool(false) => writeasm!(self, "\tcmpq ${{}}, {{}}, &0, target),",
        ir::Constant::Int(i: &i64) => writeasm!(self, "\tcmpq ${i}, {target}"),
    }
}

pub fn emit_jump(&mut self, label: &str) -> EmitResult {
    writeasm!(self, "\tjmp {label}")
}
```

Lesson:
Debuggability is precious

(A lesson one might have to learn multiple times)

Register Allocation

Or: My Darkest Hour

It's hard

```
1 # Register allocation for Tucan
2
3 ## Resources
4
5 [linearscan]: https://link.springer.com/content/pdf/10.1007%
6 [craneliftregalloc]: https://github.com/bytecodealliance/was
7 [craneliftcorrectnessregalloc]: https://cfallin.org/blog/202
8 [craneliftcodegen]: https://blog.benj.me/2021/02/17/cranelif
9 [mikepallreverselinearscan]: http://lua-users.org/lists/lua-
10 [reverselinearscanblogpost]: http://brrt-to-the-future.blog
11 [wimmerlinearscan]: http://citeseerx.ist.psu.edu/viewdoc/dow
12 [optimizedinterval]: https://www.usenix.org/legacy/events/ve
13
14 ### Books
15
16 - Engineering a Compiler
17 - SSA Book (Seems like there's a newer version of the book h
18
19 ### Papers
20
21 - [Linear Scan Register Allocation in the Context of SSA For
22 - [Linear Scan Register Allocation on SSA form][wimmerlinear
23 - [Optimized Interval Splitting in a Linear Scan Register Al
24
25 ### Misc
26
27 - Cranelift's README on SSA-based regalloc: [craneliftregall
28 - Mike Pall on LuaJIT's "reverse linear scan" algorithm: [mi
29 - Blog post from one of the author's of MoarVM on "reverse l
30
31 ### Slides
32
33 - https://ethz.ch/content/dam/ethz/special-interest/infk/ins
```



Thorsten Ball  @thorstenball · Oct 31, 2021



Compilers folks: the computation of the LiveOut sets per block shown in the Linear-Scan-Regalloc-in-SSA paper (link.springer.com/content/pdf/10...) seems much simpler than the iterative fixed-point algorithms I've seen.

Is that only because SSA makes it simpler or am I missing something?

ranges $[1,3]$, $[3,7]$ are merged into a single range $[1,7]$.
The algorithm BUILDINTERVALS() traverses the control flow graph in an arbitrary order, finds out which values are live at the end of every block, and computes the ranges for these values as described above.

```
BuildIntervals()
  for each block b do
    live ← {}
    for each successor s of b do
      live ← live ∪ s.live
      for each φ-function phi in s do
        live ← live - {phi} ∪ {phi.opd(b)}
    for each instruction i in live do ADDRANGE(i,
b, b.last.n+1)
  for all instructions i in b in reverse order
  do
    live ← live - {i}
    for each operand opd of i do
      if opd ∉ live then
        live ← live ∪ {opd}
        ADDRANGE(opd, b, i.n)
```

Fig. 12 shows a sample program in source code and in intermediate representation with a ϕ -function for the value d and corresponding move instructions in the predecessor blocks. Fig. 13 shows the live intervals that are computed for this



Slava Egorov
@mrleph



Notice that it uses something called **b.live** (does not compute that!) - that's a live-in set for block b. In other words it expects that a liveness analysis has been run prior to constructing the intervals

Day after, 03 Nov 21, 6:30am.

The [[paper version of the algorithm in the SSA book](#)][[computingliveness](#)] has a clear list of requirements for the liveness analysis:

- > Since our algorithm exploits advanced program properties some prerequisites
- > have to be met by the input program and the compiler framework:
- > * The CFG of the input program is available.
- > * The program has to be in strict SSA form.
- > * A loop-nesting forest of the CFG is available.

But they also seem to present *another* algorithm, one that doesn't need the loop-nesting forest:

- > "Liveness Sets using Path Exploration"
- > For SSA programs, another approach is possible that follows the classical definition of liveness:

- > a variable is live at a program point p , if p belongs to a path of the CFG
- > leading from a definition of that variable to one of its uses without passing
- > through another definition of the same variable.

- > Therefore, the live-range of a variable can be computed using a backward traversal starting
- > on its uses and stopping when reaching its (unique) definition. For comparison,
- > we designed optimized implementations of this path-exploration principle (see
- > Section 5), for both SSA and non-SSA programs, and compared the efficiency
- > of the resulting algorithms with our novel non-iterative data-flow algorithm.

And that algorithm doesn't require a loop-nesting forest to have been built:

```
```javascript
// Compute liveness sets by exploring paths from variable uses.
function Compute_LiveSets_SSA_ByUse(CFG) {
 for each basic block B in CFG do // Consider all blocks successively
 for each $v \in \text{PhiUses}(B)$ do // Used in the ϕ of a successor block
 LiveOut(B) = LiveOut(B) \cup {v}
 Up_and_Mark(B, v)
 for each v used in B (ϕ excluded) do // Traverse B to find all uses
 Up_and_Mark(B, v)
}
```

# Liveness Analysis

```
// compute_loop_nesting_forest builds a loop-nesting forest (finding loops in a function and
// nesting them under each other).
//
// It's based on
//
// - Paul Havlak - Nesting of Reducible and Irreducible loops
//
// with the help of:
//
// - G. Ramalingam - Identifying Loops In Almost Linear Time
// - R. Endre Tarjan - Testing Flow Graph Reducibility
//
// Big difference to Havlak's algorithm is that we don't use the DFS' ranking to determine whether
// a block is an ancestor, but use dominators instead.
pub fn compute_loop_nesting_forests(funk: &BlockFunk) → LoopForest {
 let (order, numbers) = compute_dfs_numbers(funk);

 let mut non_back_preds: Vec<HashSet<BlockId>> = Vec::with_capacity(order.len());
 let mut back_preds: Vec<HashSet<BlockId>> = Vec::with_capacity(order.len());

 let mut forest = LoopForest::new();
 let mut union_find = UnionFind::new(order.clone());

 let dominates = |a, b| funk[b].doms.contains(a):
```



# regalloc: Just a 150 line function

```
fn linear_scan(context: &mut Context, funk: &mut BlockFunk, mut free: HashSet<codegen::Reg> {
51 let all_available_registers = free.clone();
52 let weights = calc_place_weights(context, funk);
53
54 // We initialize memory locations, which already assigns parameters to memory locations, making
55 // them "fixed" (in the terms of the paper) intervals
56 let mut mem = Locations::new(context, funk);
57
58 // We then use this HashSet here to determine whether an interval is fixed
59 let fixed = mem
60 .stack
61 .keys()
62 .chain(mem.registers.keys())
63 .copied()
64 .collect::<HashSet<_>>();
65
66 // Unhandled are all the intervals in increasing order of their start points.
67 let mut unhandled = build_worklist(context, funk, &mut mem);
68
69 let mut active: HashSet<PlaceInterval> = HashSet::new();
70 let mut inactive: HashSet<PlaceInterval> = HashSet::new();
71 let mut handled: HashSet<PlaceInterval> = HashSet::new();
72
73 // Pick and remove the first interval from unhandled
74 while let Some(cur) = unhandled.pop_front() {
75 // Check for active intervals that expired
76 expire_active_intervals(
77 &mut active,
78 &mut handled,
79 &mut inactive,
80 &mut free,
81 &mem,
82 &cur,
83);
84
85 // Check for inactive intervals that expired or become reactivated
86 check_inactive_intervals(
87 &mut active,
88 &mut handled,
89 &mut inactive,
90 &mut free,
91 &mem,
92 &cur,
93);
94
95 // Collect available registers, by removing the registers assigned to the "inactive" and
96 // "unhandled" intervals (the latter only if they have been assigned a register).
97 let mut available_regs = free.clone();
98 for reg in inactive
99 .iter()
100 .chain(unhandled.iter().filter(|pi| fixed.contains(&pi.place)))
101 .filter(|pi| pi.intervals.overlap(cur.intervals))
102 .filter_map(|pi| mem.get_reg(&pi.place))
103 {
104 available_regs.remove(®);
105 }
106
107 // Check if there is a register available to assign the current interval to.
108 if available_regs.is_empty() {
109 // If not... we check whether the weight of our current interval is higher than
110 // the weight of some of the intervals that have been assigned a register.
111 //
112 // If that's the case, then we bump the lower-weight interval from the registers, take
113 // the register, and spill the lower-weight interval to the stack.
114
115 let mut reg_weights = all_available_registers
116 .iter()
117 .cloned()
118 .map(|rl (r, 0)|
119 .collect::<HashMap<codegen::Reg, i32>>());
120
121 for int in active
122 .iter()
123 .chain(inactive.iter())
124 .chain(unhandled.iter().filter(|pi| fixed.contains(&pi.place)))
125 .filter(|int| int.intervals.overlap(cur.intervals))
126 {
127 if let Some(reg) = mem.get_reg(&int.place) {
128 let w = if fixed.contains(&int.place) {
129 FIXED_PLACE_WEIGHT
130 } else {
131 weights[&int.place]
132 };
133 reg_weights.entry(reg).and_modify(|e| *e += w);
134 }
135 }
136
137 // Get weight of current interval
138 let cur_weight = match weights.get(&context.place_rep(funk.id, &cur.place)) {
139 Some(w) => w,
140 None => panic!(
141 "no weight for place {} (rep: {})",
142 cur.place,
143 context.place_rep(funk.id, &cur.place)
144),
145 };
146
147 // Find a register candidate with the lowest weight
148 match reg_weights.iter().min_by_key(|e| e.1) {
149 // Check whether our interval has higher or lower weight.
150 Some((reg_candidate, reg_weight)) if cur_weight > reg_weight => {
151 // If it's higher, we spill the assigned intervals to the stack:
152 //
153 // 1. Assign memory locations to the intervals occupied by reg_candidate.
154 // 2. Move all active or inactive intervals to which reg_candidate was assigned
155 // to handled.
156 // 3. Assign memory locations to them.
157 for set in [&mut active, &mut inactive] {
158 set.retain(|pi| {
159 if let Some(true) = mem
160 .get_reg(&pi.place)
161 .map(|int_reg| int_reg == *reg_candidate)
162 {
163 handled.insert(*pi);
164 mem.registers.remove(&pi.place);
165 mem.add_to_stack(pi.place);
166 false
167 } else {
168 true
169 }
170 });
171 }
172 mem.assign(cur.place, (*reg_candidate).into());
173 active.insert(cur);
174 }
175 - => {
176 // If it's lower, it's the current interval that needs to be spilled to the
177 // stack.
178 // But only if it's not a fixed location:
179 if !fixed.contains(&cur.place) {
180 mem.add_to_stack(cur.place);
181 }
182 handled.insert(cur);
183 }
184 }
185
186 } else {
187 let used_reg = if fixed.contains(&cur.place) {
188 mem.get_reg(&cur.place).unwrap()
189 } else {
190 let reg = available_regs.drain().next().unwrap();
191 mem.assign(cur.place, reg.into());
192 };
193
194 free.remove(&used_reg);
195 active.insert(PlaceInterval {
196 place: cur.place,
197 intervals: cur.intervals,
198 });
199 }
200
201 }
202
203 context.memory_locations.insert(funk.id, mem);
204 }
```

**Lesson:**

I can do hard things?

# Bonus Lesson





Thorsten Ball

to hanspeter.moessenboeck@jku.at

Nov 25, 2021 (3 years ago)

Sent Reply to all Actions

Hallo!

Ich baue privat grade einen Compiler in Rust und finde Ihr Paper "Linear Scan Register Allocation in the Context of SSA Form and Register Constraints" sehr gut und nützlich.

Ich habe aber eine Frage zu Section 4:





Fig. 12. Sample program in source code and in intermediate representation

This should be  
... = b + d  
right?



Thorsten Ball

to hanspeter.moessenboeck@jku.at

Nov 25, 2021 (3 years ago)

Sent Reply to all Actions

Hallo!

Ich baue privat grade einen Compiler in Rust und finde Ihr Paper "Linear Scan Register Allocation in the Context of SSA Form and Register Constraints" sehr gut und nützlich.

Ich habe aber eine Frage zu Section 4:

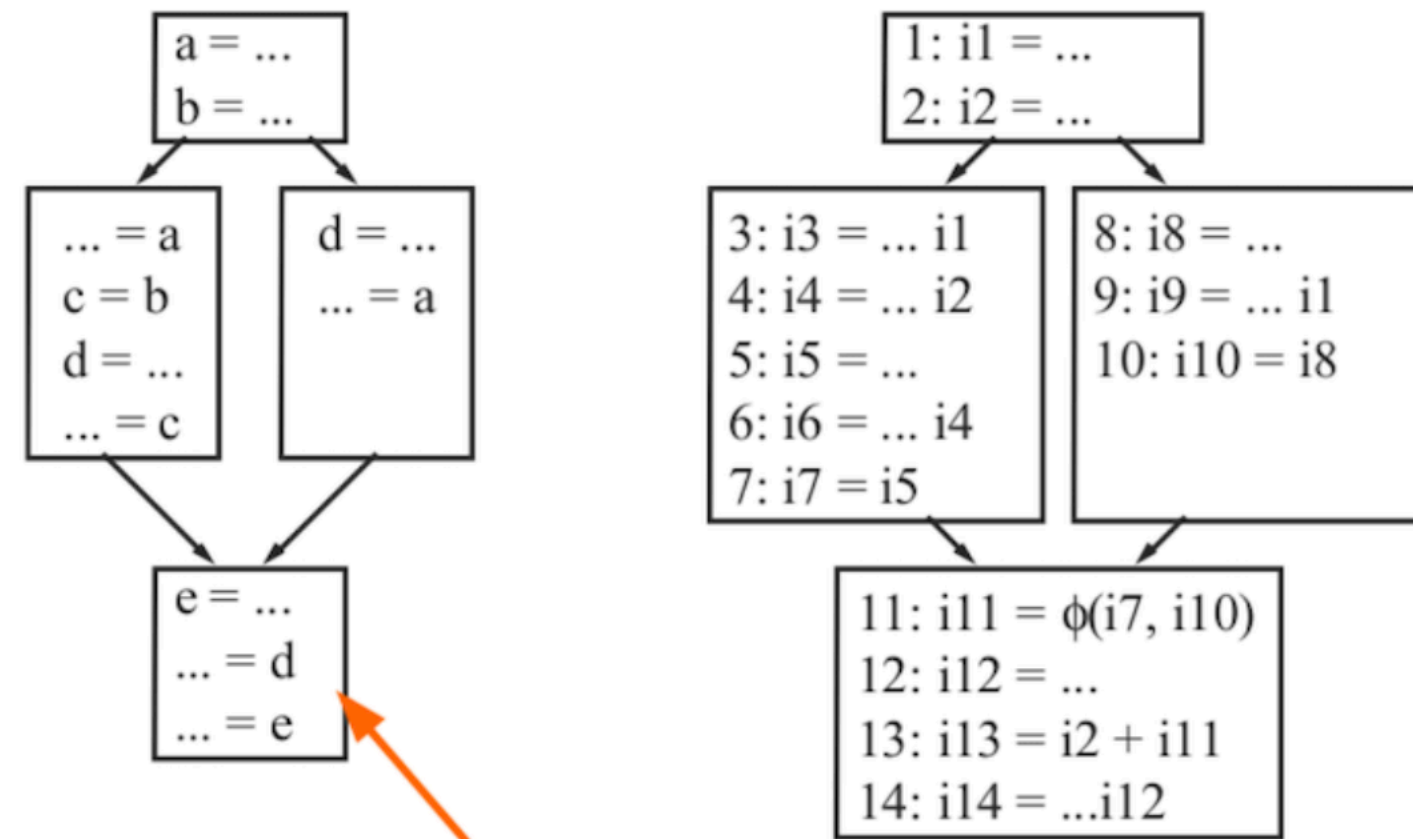


Fig. 12. Sample program in source code and in intermediate representation

This should be  
... = b + d  
right?



Peter Mössenböck

to Thorsten Ball

Nov 25, 2021 (3 years ago)

Archive Reply Actions

Lieber Herr Ball,

Sie haben natürlich recht, und ich wundere mich, dass das noch niemandem aufgefallen ist.

Es muss entweder im linken Diagram heißen

... = b + d

oder im rechten Diagramm

i13 = i11

Da aber in Fig.13 das Live-Intervall von b bis Instruktion 13 geht, ist die erste Interpretation die richtige.

Vielen Dank für den Hinweis.

Beste Grüße

Hanspeter Mössenböck



**Lesson:**

Even experts make mistakes

# Optimizations

# Optimizations

```
// Sparse Conditional Constant Propagation.
//
// This is based on "10.7.1 Combining Optimizations" in Cooper/Torczon - Engineering A Compiler,
// which is based on the "Sparse Simple Constant Propagation (SSCP)" algorithm presented earlier in
// Section 9.3.6
```

```
// "In SSCP [and SCCP], the algorithm initializes the value associated with each SSA name to
// \top , which indicates that the algorithm has no knowledge of the SSA name's
// value.
//
// If the algorithm subsequently discovers that SSA name x has the known constant value Ci, it
// models that knowledge by assigning Value(x) the semi-lattice element Ci.
//
// If it discovers that x has a changing value, it models that fact with the value \perp ."
```

```
#[derive(PartialEq, Eq, Clone, Debug)]
```

```
enum Value {
 Top, // \top
 Const(Constant),
 Bot, // \perp
}
```

```
impl Value {
 fn is_bot(&self) -> bool {
 matches!(self, Value::Bot)
 }
}
```

```
// Rules for Meet:
```

```
// $\top \wedge x = x \quad \forall x$
// $\perp \wedge x = \perp \quad \forall x$
// $C_i \wedge C_j = C_i \quad \text{if } C_i = C_j$
// $C_i \wedge C_j = \perp \quad \text{if } C_i \neq C_j$
```

```
fn meet(&self, other: &Self) -> Self {
 match (self, other) {
 (Value::Top, other: &Value) | (other: &Value, Value::Top) => other.clone(),
 (Value::Bot, _) | (_, Value::Bot) => Value::Bot,
 (Value::Const(i: &Constant), Value::Const(j: &Constant)) => {
 if i == j {
 Value::Const(i.clone())
 } else {
 Value::Bot
 }
 }
 }
}
```

```
#[test]
fn handling_phi_functions_diamond_shape() {
 // This is taken from p8 of "Briggs, Cooper, Simpson - Value Numbering"
```

```
// BEFORE:
```

```
b0
a_0 = param(0)
b_0 = param(1)
c_0 = param(2)
d_0 = param(3)
e_0 = param(4)
f_0 = param(5)
```

```
u_0 = a_0 + b_0
v_0 = c_0 + d_0
w_0 = e_0 + f_0
```

```
b1
x_0 = c_0 + d_0
y_0 = c_0 + d_0
```

```
b2
u_1 = a_0 + b_0
x_1 = e_0 + f_0
y_1 = e_0 + f_0
```

```
b3
u_2 = $\phi(u_0, u_1)$
x_2 = $\phi(x_0, x_1)$
y_2 = $\phi(y_0, y_1)$
z_0 = u_2 + y_2
u_3 = a_0 + b_0
```

```
// AFTER:
```

```
b0
a_0 = 11
b_0 = 22
c_0 = 33
d_0 = 44
e_0 = 55
f_0 = 66
```

```
u_0 = a_0 + b_0
v_0 = c_0 + d_0
w_0 = e_0 + f_0
```

```
b1
b2
```

```
b3
x_2 = $\phi(v_0, w_0)$ <-- IMPORTANT
z_0 = u_0 + x_2 <-- IMPORTANT
```

```
versioned_registers![
```

```
[a_0],
[b_0],
[c_0],
[d_0],
[e_0],
[f_0],
[u_0, u_1, u_2, u_3],
[v_0],
[w_0],
[x_0, x_1, x_2],
[y_0, y_1, y_2],
[z_0]
```

```
];
```

```
let b0_ins: Vec<Instruction> = vec![
```

```
mov!(a_0, param!(0)),
mov!(b_0, param!(1)),
mov!(c_0, param!(2)),
mov!(d_0, param!(3)),
mov!(e_0, param!(4)),
mov!(f_0, param!(5)),
```

**Lesson:**

Rust !  Graph Manipulation



# Uh, oh

~/drive/notes/Tucan - Building an optimizing compiler.md

```
1 # Tucan - Building an optimizing compiler
2
3 - [Compiler Resources](./Compiler\ Resources.md)
4 - [Tucan - Intermediate Representation (IR)](./Tucan\ -\ Intermediate\ Representation\ IR.md)
5 - [Tucan - Constructing SSA for Tucan from CFG](./Tucan\ -\ Constructing\ SSA\ for\ Tucan\ from\ CFG.md)
6 - [Tucan - Codegen for Tucan](./Tucan\ -\ Codegen\ for\ Tucan.md)
7 - [Tucan - Last expression statement is returned](./Tucan\ -\ Last\ expression\ statement\ is\ returned.md)
8 - [Tucan - Liveness Analysis](./Tucan\ -\ Liveness\ Analysis.md)
9 - [Tucan - Register allocation for Tucan](./Tucan\ -\ Register\ allocation\ for\ Tucan.md)
10 - [Tucan - Optimizations](./Tucan\ -\ Optimizations.md)
11 - [Tucan - Fixing regalloc liveness bug](./Tucan\ -\ Fixing\ regalloc\ liveness\ bug.md)
12 - [Tucan - Fixing phi operand bug](./Tucan\ -\ Fixing\ phi\ operand\ bug.md)
13 - [Tucan - Fixing another bug in register allocator](./Tucan\ -\ Fixing\ another\ bug\ in\ register\ allocator.md)
14 - [Tucan - Garbage Collection](./Tucan\ -\ Garbage\ Collection.md)
15
```



**Thorsten Ball**  @thorstenball · Jun 12, 2022



good news: it's not in the optimization pass that I finished this week  
bad news: found out my compiler is flaky







**Thorsten Ball**  @thorstenball · Jun 12, 2022



good news: it's not in the optimization pass that I finished this week  
bad news: found out my compiler is flaky



**Thorsten Ball**  @thorstenball · Jun 14, 2022



oh no, it's in the register allocator





**Thorsten Ball**  @thorstenball · Jun 12, 2022



good news: it's not in the optimization pass that I finished this week  
bad news: found out my compiler is flaky



1



2



27



**Thorsten Ball**  @thorstenball · Jun 14, 2022



oh no, it's in the register allocator



3



2



19



**Ben L. Titzer** @TitzerBL · Jun 14, 2022



The ninth circle of compiler hell has been reached



1



1



10



Time to bring out the  
big guns

# Tree-sitter grammars

```
example.tuc
examples/example.tuc
1 funk hello_world() → bool {
2 print_string("hello world!");
3 return true;
4 }
5 |
6 // magic_funk is a magical function
7 funk magic_funk(a: int, b: int, c: int, d: bool, e: bool) → bool {
8 // comment here
9 let res = a + b + c;
10 // comment there
11 return res == d == e;
12 }
13
14 funk x(a: int, b: bool, c: string) → int {
15 let a: int = 18;
16 let b: bool = a == 10;
17 let c: bool = a ≠ b;
18 let d = true ≠ false;
19 // comment
20 let s = "foobar" = "foo bar";
21
22 1 + 2;
23 1 - 2;
24 1 > 2;
25 1 < 2;
26 1 ⇒ 2;
```

```
example.tucir
examples/example.tucir
1 block_funk #1(node_id: #33, entry: #1)
2 block #1 (→ #6):
3 %1_0 ← 5
4 %2_0 ← 0
5 terminator: goto(#6)
6
7 block #6 (← #1, #4) (→ #2, #3):
8 %2_1 ← phi(var(#1, %2_0), var(#4, %2_2))
9 %8_0 ← %2_1 < 10
10 terminator: if(%8_0, (#3, #2))
11
12 block #2 (← #6):
13 %18_0 ← true
14 %19_0 ← false
15 terminator: return(0)
16
17 block #3 (← #6) (→ #4, #5):
18 %3_0 ← %2_1 < %1_0
19 terminator: if(%3_0, (#5, #4))
20
21 block #4 (← #3, #5) (→ #6):
22 %6_0 ← call print_num, args: (1)
23 %7_0 ← %2_1 + 1
24 %2_2 ← %7_0
```

# Tree-sitter grammars

crates/languages/src/rust/injections.scm

```
1 (macro_invocation
2 (token_tree) @content
3 (#set! "language" "rust"))
4
5 (macro_rule
6 (token_tree) @content
7 (#set! "language" "rust"))
8
9
10 ;; Inject the tucan grammar into Rust strings that are bound to `let tucan =`
11 ((let_declaration
12 pattern: (identifier) @_new
13 (#eq? @_new "tucan")
14 value: [(raw_string_literal
15 (string_content) @content)
16 (string_literal
17 (string_content) @content)])
18 (#set! "language" "tucan"))
```

Works in Zed and Neovim

```
182 #[test]
183 fn for_loop() {
184 let tucan = r#"
185 funk main() → int {
186 let global_one = 0;
187 let global_two = 99;
188
189 for (let i = 0; i < 10; i = i + 1) {
190 global_one = global_one + i;
191 global_two = global_two + 0;
192 print_num(i);
193 }
194
195 print_string("-");
196 print_num(global_one);
197 print_string("-");
198 print_num(global_two);
199 return 0;
200 }"#;
201
202 run("for_loop", tucan, 0, "0123456789-45-99");
203 }
204
205 #[test]
206 fn loop_containing_conditionals() {
207 let tucan = r#"
208 funk main() → int {
209 let global = 0;
210 for (let i = 0; i < 10; i = i + 1) {
211 if (i < 5) {
212 global = global + i;
213 } else {
214 global = global + i - 5;
215 }
216 }
217 print_num(global);
218 return 0;
219 }"#;
```



# Integration tests

```
#[test]
fn big_integration_boi() {
 let tucan: &str = r#"
 funk double_print(num: int) {
 let doubled = num;
 for (let i = num; i > 0; i = i - 1) {
 doubled = doubled + 1;
 }
 print_num(doubled);
 }

 funk my_function(a: int, b: int) -> int {
 let c = a + b + 18 - 5 + 3 + 2;
 let d = 18 - 5 + 3;
 let e = a + b + 30;
 let f = 100 - a - b + 3 + 2;
 let g = c + d + e; // dead code
 let h = a + b + c + e + f;

 let k = 50;

 let result = 0;
 for (let i = 0; i < 10; i = i + 1) {
 result = result + 5;
 if (i < a) {
 if (k >= d) {
 double_print(i);
 } else {
 print_num(i);
 }
 } else {
 result = result + c;
 }
 for (let j = 5; j > 0; j = j - 1) {
 result = result - j;
 }
 print_num(result);
 print_string("-");
 }
 return result;
 }

 funk main() -> int {
 my_function(5, 10);
 my_function(50, 2);
 return 0;
 }
 "#;

 run(
 test_name: "big_integration_boi",
 code: tucan,
 exit_code: 0,
 output: "0-10-2-20-4-30-6-40-8-50--27--4-19-42-65-0-10-2-20-4-30-6-40-8-50-10-60-12-70-14-80-16-90-18-100-",
);
} fn big_integration_boi
```

```
fn run(test_name: &str, code: &'static str, exit_code: i32, output: &'static str) {
 const TOTAL_RUNS: i32 = 12;

 let combinations = [(false, false), (true, true), (true, false), (false, true)];
 for i in 0..TOTAL_RUNS {
 let (optimize, regalloc) = combinations[i as usize % combinations.len()];

 print_test_line(test_name, i, TOTAL_RUNS, regalloc, optimize);

 let (asm, _) = Compiler::for_string(test_name, code)
 .optimize(optimize)
 .regalloc(regalloc)
 .compile()
 .expect("compiler failed");

 let asm_name = format!("integration-test-{{test_name}}-{{i}}");
 let result = AssemblyRunner::new_for_test(&asm_name, &asm)
 .run()
 .expect("failed to create assembly runner");

 assert_eq!(result.0, exit_code, "run #{{i}}: wrong exit code");
 assert_eq!(result.1, output, "run #{{i}}: wrong output");
 }
}
```







# my\_function

## source

```
funk double_print(num: int) {
 let doubled = num;
 for (let i = num; i > 0; i = i - 1) {
 doubled = doubled + 1;
 }
 print_num(doubled);
}

funk my_function(a: int, b: int) → int {
 let c = a + b + 18 - 5 + 3 + 2;
 let d = 18 - 5 + 3;
 let e = a + b + 30;
 let f = 100 - a - b + 3 + 2;
 let g = c + d + e; // dead code
 let h = a + b + c + e + f;

 let k = 50;

 let result = 0;
 for (let i = 0; i < 10; i = i + 1) {
 result = result + 5;
 if (i < a) {
 if (k ≥ d) {
 double_print(i);
 } else {
 print_num(i);
 }
 } else {
 result = result + c;
 }
 }
 for (let j = 5; j > 0; j = j - 1)
 result = result - j;
 print_num(result);
 print_string("-");
}

return result;
}

funk main() → int {
 my_function(5, 10);
 my_function(50, 2);
 return 0;
}
```

pre-ssa
ssa
value_numbering
constant_propagation
unreachable_code_elimination
dead_code_elimination
clean
unreachable_code_elimination2
regalloc
asm

# (Stole it from Go)

File /Users/thorstenball/code/go/src/github.com/mrnugget/monkey/evaluator/ssa.html

Eval,1 help darkmode

<p>sources AST before insertphis</p> <p><b>start</b></p> <pre>b1: v1 (?) = InitMem &lt;mem&gt; v2 (?) = SP &lt;uintptr&gt; v3 (?) = SB &lt;uintptr&gt; v4 (?) = LocalAddr &lt;ast.Node&gt; {node} v2 v1 v5 (?) = LocalAddr &lt;ast.Environment&gt; {env} v2 v1 v6 (?) = LocalAddr &lt;ast.Object&gt; {~r0} v2 v1 v7 (i6) = Arg &lt;ast.Node&gt; {node} (node[ast.Node]) v8 (i6) = Arg &lt;ast.Environment&gt; {env}       (env[*object.Environment]) v9 (?) = ConstInterface &lt;ast.Object&gt; v10 (17) = ITab &lt;uintptr&gt; v7 v11 (?) = ConstNil &lt;uintptr&gt; v12 (17) = EqPtr &lt;bool&gt; v10 v11 v18 (?) = Const64 &lt;uint&gt; [2] v20 (?) = Const32 &lt;uint32&gt; [63] v28 (?) = Addr &lt;uint8&gt;       {type:*github.com/mrnugget/monkey/ast.Identifier}       v3 v29 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.Identifier}       v3 v33 (?) = ConstBool &lt;bool&gt; [true] v34 (?) = ConstNil &lt;ast.Identifier&gt; v35 (?) = ConstBool &lt;bool&gt; [false] (~r0[bool],       ~r0[bool], ~r0[bool], ~r0[bool],       ~r0[bool], ~r0[bool], ~r0[bool], ~r0[bool],       ~r0[bool]) v39 (?) = Addr &lt;uint8&gt;       {type:*github.com/mrnugget/monkey/ast.IndexExpression}       v3 v40 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.IndexExpression}       v3 v44 (?) = ConstNil &lt;ast.IndexExpression&gt; v48 (?) = Addr &lt;uint8&gt;       {type:*github.com/mrnugget/monkey/ast.LetStatement}       v3 v49 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.LetStatement}       v3 v53 (?) = ConstNil &lt;ast.LetStatement&gt; v57 (?) = Addr &lt;uint8&gt;       {type:*github.com/mrnugget/monkey/ast.IntegerLiteral}       v3 v58 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.IntegerLiteral}       v3 v62 (?) = ConstNil &lt;ast.IntegerLiteral&gt; v66 (?) = Addr &lt;uint8&gt;       {type:*github.com/mrnugget/monkey/ast.Boolean} v3 v67 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.Boolean} v3 v71 (?) = ConstNil &lt;ast.Boolean&gt; v75 (?) = Addr &lt;uint8&gt;       {type:*github.com/mrnugget/monkey/ast.ReturnStatement}       v3 v76 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.ReturnStatement}</pre>	<p>number lines early phielim and copyelim early deadcode short circuit decompose user + pre-opt deadcode</p>	<p><b>opt [188500 ns]</b></p> <pre>b1: v1 (?) = InitMem &lt;mem&gt; v2 (?) = SP &lt;uintptr&gt; v3 (?) = SB &lt;uintptr&gt; v7 (i6) = Arg &lt;ast.Node&gt; {node} (node[ast.Node]) v8 (i6) = Arg &lt;ast.Environment&gt; {env}       (env[*object.Environment],       env[*object.Environment],       object.e[*object.Environment],       env[*object.Environment]) v10 (17) = ITab &lt;uintptr&gt; v7 v18 (?) = Const64 &lt;uint&gt; [2] v20 (?) = Const32 &lt;uint32&gt; [63] v29 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.Identifier}       v3 v35 (?) = ConstBool &lt;bool&gt; [false] (~r0[bool],       ~r0[bool], ~r0[bool], ~r0[bool],       ~r0[bool], ~r0[bool], ~r0[bool], ~r0[bool],       ~r0[bool]) v40 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.IndexExpression}       v3 v49 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.LetStatement}       v3 v58 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.IntegerLiteral}       v3 v67 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.Boolean} v3 v76 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.ReturnStatement}       v3 v85 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.HashLiteral}       v3 v94 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.BlockStatement}       v3 v103 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.FunctionLiteral}       v3 v112 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.IfExpression}       v3 v121 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.CallExpression}       v3 v130 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.InfixExpression}       v3 v139 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.Program} v3 v148 (?) = Addr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.StringLiteral}       v3 v157 (?) = Addr &lt;uint8&gt;</pre>	<p>zero arg cse opt deadcode generic cse phi opt + gcse deadcode nichekelim prove early fuse expand calls decompose builtin softfloat + late opt dead auto elim + scop + generic deadcode check bce + branchelim + late fuse dse memcombine writebarrier</p>	<p><b>lower [179875 ns]</b></p> <pre>b1: v1 (?) = InitMem &lt;mem&gt; v2 (?) = SP &lt;uintptr&gt; v3 (?) = SB &lt;uintptr&gt; v8 (17) = ArgIntReg &lt;ast.Environment&gt; {env+0} [2]       (env[*object.Environment],       env[*object.Environment],       env[*object.Environment],       object.e[*object.Environment]) v29 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.Identifier}       v3 v35 (?) = MOVDConst &lt;bool&gt; [0] (~r0[bool], ~r0[bool],       ~r0[bool], ~r0[bool], ~r0[bool], ~r0[bool],       ~r0[bool], ~r0[bool], ~r0[bool], ~r0[bool]) v40 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.IndexExpression}       v3 v49 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.LetStatement}       v3 v58 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.IntegerLiteral}       v3 v67 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.Boolean} v3 v76 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.ReturnStatement}       v3 v85 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.HashLiteral}       v3 v94 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.BlockStatement}       v3 v103 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.FunctionLiteral}       v3 v112 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.IfExpression}       v3 v121 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.CallExpression}       v3 v130 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.InfixExpression}       v3 v139 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.Program} v3 v148 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.StringLiteral}       v3 v157 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.ArrayLiteral}       v3 v166 (?) = MOVDAddr &lt;uint8&gt;       {go:itab.*github.com/mrnugget/monkey/ast.PrefixExpression}</pre>	<p>addressing modes + late lower</p>
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------------------------------------



Two lessons:

1. Keep thinking “how will I debug this?”
2. Invest in debug tooling

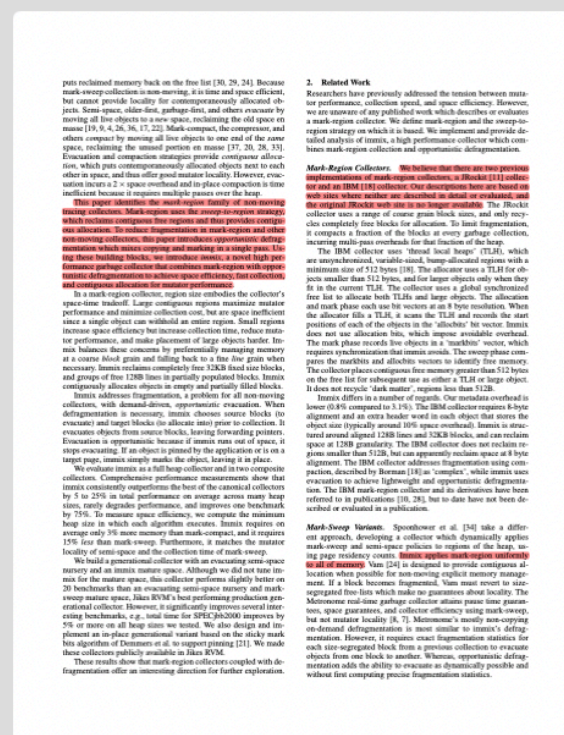
# Runtime & GC



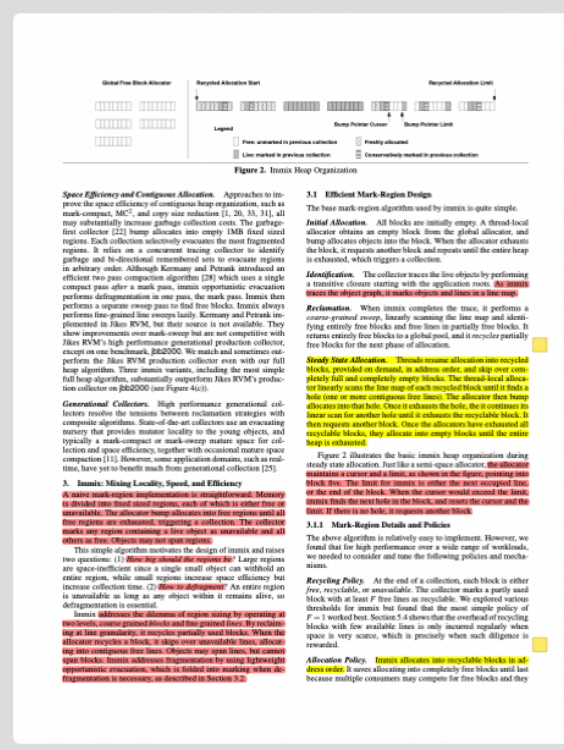
# Immix



1



2



3

# Immix: A Mark-Region Garbage Collector with Space Efficiency, Fast Collection, and Mutator Performance\*

Stephen M. Blackburn  
Australian National University  
Steve.Blackburn@anu.edu.au

Kathryn S. McKinley  
The University of Texas at Austin  
mckinley@cs.utexas.edu

## Abstract

Programmers are increasingly choosing managed languages for modern applications, which tend to allocate many short-to-medium lived small objects. The garbage collector therefore directly determines program performance by making a classic space-time trade-off that seeks to provide space efficiency, fast reclamation, and mutator performance. The three canonical tracing garbage collectors: semi-space, mark-sweep, and mark-compact each sacrifice one objective. This paper describes a collector family, called *mark-region*, and introduces *opportunistic* defragmentation, which mixes copying and marking in a single pass. **Combining both, we implement *immix*, a novel high performance garbage collector that achieves all three performance objectives. The key insight is to allocate and reclaim memory in contiguous regions, at a coarse *block grain* when possible and otherwise in groups of finer *grain lines*.** We show that *immix* outperforms existing canonical algorithms, improving total application performance by 7 to 25% on average across 20 benchmarks. As the mature space in a generational collector, *immix* matches or beats a highly tuned generational collector, e.g. it improves *jbb2000* by 5%. These innovations and the identification of a new family of collectors open new opportunities for garbage collector design.

**Categories and Subject Descriptors** D.3.4 [Programming Languages]: Processors—Memory management (garbage collection)  
**General Terms** Algorithms, Experimentation, Languages, Performance, Measurement  
**Keywords** Fragmentation, Free-List, Compact, Mark-Sweep, Semi-Space, Mark-Region, Immix, Sweep-To-Region, Sweep-To-Free-List

## 1. Introduction

Modern applications are increasingly written in managed languages and make conflicting demands on their underlying memory managers. For example, real-time applications demand pause-time guarantees, embedded systems demand space efficiency, and servers demand high throughput. In seeking to satisfy these demands, the literature includes reference counting collectors and three canonical tracing collectors: semi-space, mark-sweep, and mark-compact. These collectors are typically used as building blocks for more sophisticated algorithms. Since reference counting is incomplete, we omit it from further consideration here. Unfortunately, the tracing collectors each achieve only two of: *space*

\* This work is supported by ARC DP0666059, NSF CNS-0719966, NSF CCF-0429859, NSF EIA-0303609, DARPA F33615-03-C-4106, Microsoft, Intel, and IBM. Any opinions, findings and conclusions expressed herein are the authors' and do not necessarily reflect those of the sponsors.

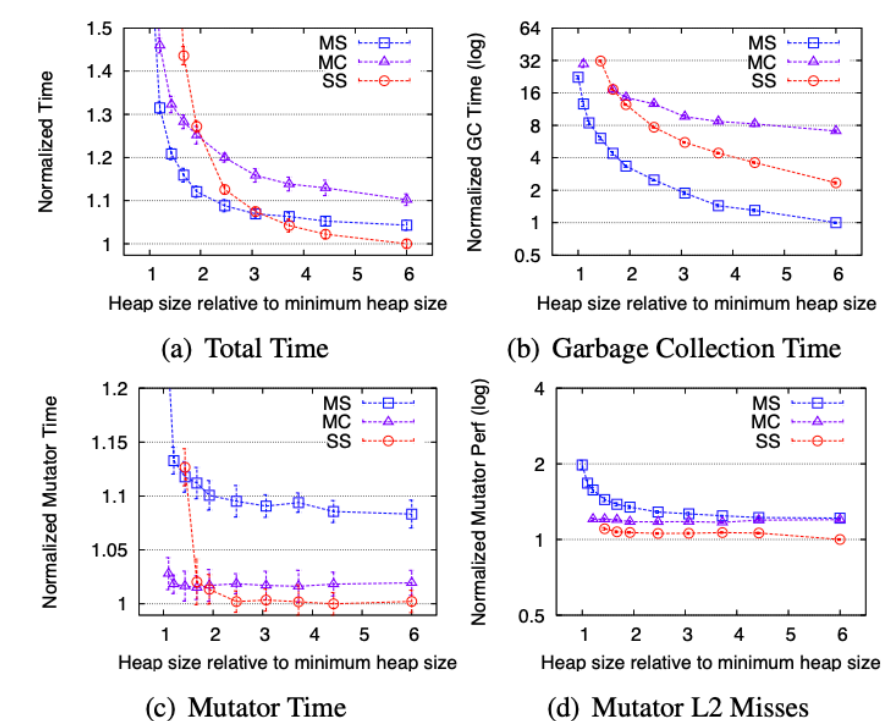
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.  
PLDI'08, June 7–13, 2008, Tucson, Arizona, USA.  
Copyright © 2008 ACM 978-1-59593-860-2/08/06...\$5.00

*efficiency, fast collection, and mutator performance* through contiguous allocation of contemporaneous objects.

Figure 1 starkly illustrates this dichotomy for full heap versions of mark-sweep (MS), semi-space (SS), and mark-compact (MC) implemented in MMTk [12], running on a Core 2 Duo. It plots the geometric mean of total time, collection time, mutator time, and mutator cache misses as a function of heap size, normalized to the best, for 20 DaCapo, SPECjvm98, and SPECjbb2000 benchmarks. The crossing lines in Figure 1(a) illustrate the classic space-time trade-off at the heart of garbage collection. Mark-compact is uncompetitive in this setting due to its overwhelming collection costs. In smaller heap sizes, the space and collector efficiency of mark-sweep perform best since the overheads of garbage collection dominate total performance. Figures 1(c) and 1(d) show that the primary advantage for semi-space is 10% better mutator time compared with mark-sweep, due to better cache locality. Once the heap size is large enough, garbage collection time reduces, and the locality of the mutator dominates total performance so semi-space performs best.

To explain this tradeoff, we need to introduce and slightly expand memory management terminology. **A tracing garbage collector performs *allocation* of new objects, *identification* of live objects, and *reclamation* of free memory. The canonical collectors all identify live objects the same way, by marking objects during a transitive closure over the object graph.**

Reclamation strategy dictates allocation strategy, and the literature identifies just three strategies: (1) *sweep-to-free-list*, (2) *evacuation*, and (3) *compaction*. For example, mark-sweep collectors allocate from a free list, mark live objects, and then *sweep-to-free-list*



**Figure 1.** Performance Tradeoffs For Canonical Collectors: Geometric Mean for 20 DaCapo and SPEC Benchmarks.




🔍 Search or go to...

mu / immix-rust

- Project
- immix-rust
  - Manage >
  - Plan >
  - Code >
  - Build >
  - Deploy >
  - Operate >
  - Monitor >
  - Analyze >

# immix-rust

master immix-rust History Find file Code

 **add apache 2.0 license**  
qinsoon authored 5 years ago db2bf8bc

Name	Last commit	Last update
.settings	allows compile GC as library, added a C...	8 years ago
c_src	moved from internal mercurial repo	8 years ago
rust_c_interface	put fastpaths in C, fixed a synchronizati...	8 years ago
src	removed use of perf.c (used to use it fo...	7 years ago
.gitignore	put fastpaths in C, fixed a synchronizati...	8 years ago
Cargo.toml	changed dependency of time create ve...	8 years ago
LICENSE	add apache 2.0 license	5 years ago
README.md	removed use of perf.c (used to use it fo...	7 years ago
build.rs	removed use of perf.c (used to use it fo...	7 years ago
run-mt-trace.py.example	allows compile GC as library, added a C...	8 years ago
run.py.example	allows compile GC as library, added a C...	8 years ago
test_mac.sh	put fastpaths in C, fixed a synchronizati...	8 years ago

README.md



tucan\_runtime/src/heap.rs

```
95
96 > pub struct ImmixHeap {...
114 }
115
116 impl ImmixHeap {
117 > pub fn new() -> ImmixHeap {...
136 }
137
138 > pub fn new_with_auto_gc(stack_bottom: *const *const c_void) -> ImmixHeap {...
143 }
144
145 > pub fn enable_trace_gc(&mut self) {...
147 }
148
149 > pub fn stats(&self) -> Stats {...
172 }
173
174 pub fn alloc<T>(&mut self, object: T) -> Result<RawPtr<T>, AllocError>
175 where
176 T: TucanObject,
177 {
178 let header_size: usize = size_of::<ObjectHeader>();
179 let object_size: usize = size_of::<T>();
180 let total_size: usize = header_size + object_size;
181
182 let request: AllocRequest = AllocRequest::new(total_size)?;
183 let space: *const u8 = self.find_space(request)?;
184
185 self.allocated_objects_space += request.size;
186 self.live_bytes += request.size;
187
188 unsafe { write_header::<T>(space, size_bytes: object_size as u32, size_class: request.class) };
189
190 let object_space: *const u8 = unsafe { space.add(count: header_size) };
191 unsafe {
192 std::ptr::write_bytes(dst: object_space as *mut T, val: 0, count: 1);
193 std::ptr::write(dst: object_space as *mut T, src: object);
194 }
195
196 Ok(RawPtr::new(ptr: object_space as *const T))
197 }
198
199 pub fn alloc_array(&mut self, size_bytes: u32) -> Result<TucanByteArray, AllocError> {
200 let header_size: usize = size_of::<ObjectHeader>();
201 let total_size: usize = header_size + size_bytes as usize;
202
203 let request: AllocRequest = AllocRequest::new(total_size)?;
204 let space: *const u8 = self.find_space(request)?;
205
206 self.allocated_arrays_space += request.size;
207 self.live_bytes += request.size;
208
209 unsafe { write_header::<TucanByteArray>(space, size_bytes, size_class: request.class) };
210
211 let array_space: *const u8 = unsafe { space.add(count: header_size) };
212 let array: &mut [u8] =
213 unsafe { std::slice::from_raw_parts_mut(data: array_space as *mut u8, len: size_bytes as usize) };
214 array.iter_mut().for_each(|b: &mut u8| *b = 0);
215
216 Ok(RawPtr::new(ptr: array_space))
217 }
```

Lesson:  
Reference implementations are  
amazing



**Thorsten Ball**  @thorstenball · Nov 20, 2022

Now if that's not a Bob Dylan line I don't know



encapsulation of `Address` and `ObjectReference` types, *(ii)* managing ownership of address blocks, *(iii)* managing global ownership of thread-local allocations, and *(iv)* utilizing Rust libraries to support efficient parallel collection.

#### 4.1 Encapsulating Address Types


Memory managers manipulate raw memory, conjuring language-level objects from raw memory. Experience shows the importance of abstracting over both arbitrary raw addresses and references to user-level objects [4][11]. Such abstraction offers type safety and disambiguation with respect to implementation-language (Rust) references. Among the alternatives, raw pointers can be misleading and dereferencing an untyped arbitrary pointer may yield un-



**Steve Blackburn (@steveblackburn@discuss.systems)**

@stevemblackburn



 My usual coauthor was not involved with that paper. No way that line would have slipped past her and her common subexpression elimination.

10:47 AM · Nov 20, 2022



**Thorsten Ball**  @thorstenball · Dec 11, 2022

Hmmm, so is it possible to have an Immix heap without compaction?

I kinda want to avoid forwarding pointers as long as possible.

More reading to do...



**goyox86** @goyox86 · Dec 11, 2022

I thought the idea of opportunistic defragmentation was kind of part of Immix



**Thorsten Ball**  @thorstenball · Dec 11, 2022

Yeah, opportunistic defragmentation. I'm still trying to grok a lot of parts 😊



**Steve Blackburn (@steveblackburn@discuss.systems)**

@stevemblackburn

Yes. You can treat opportunistic copying as an optimisation.

Start with no copying, the once you have that working, add in the copying.

We're doing this right now for our Julia and Ruby ports (where copying is a bit of a big deal). Non copying is a fine place to start.





Thorsten Ball

@thorstenball



love it

ARE YOU COMING TO BED?

I CAN'T. THIS IS IMPORTANT.

WHAT?

|



world-class compiler experts are replying to my shitpost tweets about my sideproject compiler







**Laurence Tratt**  
@laurencetratt



My experience is similar. I wrote a simple testing library lang\_tester to make this more palatable for different compilers. I have a barely-started successor to lang\_tester to try and generalise things while still maintaining usability because it's such a powerful technique!

3:17 PM · Jun 14, 2022



1



4



Post your reply

Reply



**CF Bolz-Tereick** @cfbolz · Jun 14, 2022



It was probably one of you who pointed me to it, but I found the 'correctness' section on the recent Chris Fallin post very interesting:  
[cfallin.org/blog/2022/06/0...](http://cfallin.org/blog/2022/06/0...)



1



7



**Laurence Tratt** @laurencetratt · Jun 14, 2022



It's a very good post! My (admittedly somewhat limited) experience is that some components are much easier to fuzz than others though and I don't yet have a good intuition about when it's worth the effort or not.



1



**CF Bolz-Tereick** @cfbolz · Jun 14, 2022



I've never really worked on our (poor) register allocator but it's definitely had an enormously long bug tail. The fuzzer kept finding bugs after many days of running, and we keep finding the occasional bug in the wild too



1



3



**Ben L. Titzer** @TitzerBL · Feb 4, 2023



RegAlloc bugs are right up there with GC bugs as the things I most hate to spend time on.



3



55



Lesson:

Put stuff out into the world

The biggest lesson

Thank you