

Clojure Compilation, Backwards

- Nicola Mometto (Bronsa)

brobronsa@gmail.com

<https://github.com/Bronsa>

clojure compilation overview

- reader (text -> data)
- macroexpander (code -> code)
- analyzer (code -> AST)
- emitter (AST -> bytecode)

JVM bytecode

Classfile Structure

- Constant Pool (class/method refs, constant values)
- Class Fields
- Method Bodies

JVM bytecode

Method Bodies

- Local Variable Table
 - start, end offset, slot, name, signature
- Exception Table
 - start, end, catch handler, exception type
- Line Number Table
- Bytecode

tools.decompiler

- bytecode loader/parser (bytecode -> symbolic bytecode)
- analyzer (bytecode -> AST)
- sweetener (AST -> sugared AST)
- compiler (sugared AST -> code)
- (macro)compactor (code -> compacted code)
- pprinter (compacted code -> formatted text)

bytecode loader/parser

- BCEL to load bytecode
- parse class fields
- parse class methods
 - build exception table
 - build local variable table
 - build jump table
 - build symbolic bytecode vector

bytecode loader/parser

```
#:class{:name "test$foo",
  :super "clojure.lang.AFunction",
  ...
  :methods [#:method{:name "invokeStatic",
    :return-type "java.lang.Object",
    :arg-types [],
    :bytecode [#:insn{:name "ldc",
      :length 2,
      :pool-element #:insn{:target-value "hello"
        ..}
      :label 0}
      #:insn{:name "areturn", :length 1, :label 2}],
    :jump-table {0 0, 2 1}
    ..}
  ..]}
```

analyzer

- dispatches on super class/interface/class name
 - AFunction/RestFn -> function
 - IType/IRecord -> deftype/defrecord
 - IObj and name contains "reify" -> reify
 - name ends with "__init" -> namespace initializer

analyzer

- ignore abstract/bridge methods
- initialize static fields
- collect methods to decompile (invoke, load, deftype methods)
- process each method
 - stack machine interpreter to build AST
 - process using ctx map (statements, stack, pc, lvt, etc)

analyzer

- instruction types
 - stack ops (dup, swap, ..)
 - branch ops (ifeq, instanceof, ..) -> conditionals
 - local variable ops (aload, astore, ..) -> lexical blocks
 - math ops (ladd, ldiv, ..) -> math intrinsics
 - other ops (invokevirtual, athrow, putstatic, ..) -> all else

analyzer

```
(defmethod process-instruction :dup [{:keys [stack] :as ctx} _]
  (let [val (peek stack)]
    (-> ctx
        (update :stack conj val)))))

;; simplified version
(defmethod process-instruction :getstatic [ctx {:insn/keys [pool-element]}]
  (let [{:insn/keys [target-class target-name]} pool-element]
    (-> ctx
        (update :stack conj {:op :static-field
                             :target target-class
                             :field target-name}))))
```

analyzer

```
(defmethod process-insn :ifeq [{:keys [stack] :as ctx} insn]
  (let [else-label (goto-label insn)
        goto-endInsn (insn-at ctx {:label else-label :offset -2})
        end-label (goto-label goto-endInsn)
        {then-label :insn/label} (insn-at ctx {:offset 1})
        test (peek stack)]
    (-> ctx
        (update :stack pop)
        (process-if test [then-label (:insn/label goto-endInsn)]
                    [else-label end-label]))))
```

analyzer

```
{:op :fn,
  :name "foo",
  :fn-methods [{:op :fn-method,
    :fn-name "foo",
    :var-args? false,
    :args (),
    :body {:op :do,
      :statements [],
      :ret {:op :const,
        :val "hello"} } }]}}
```

sweetener

- pass over AST to add some syntax sugar
- some overlap with compactor
- e.g. clojure.lang.PersistentVector/EMPTY -> []

compiler

- compiles AST to clojure code
- same(ish) pass as tools.analyzer.jvm

(macro)compactor

- need to undo inlining/macroexpansion
- hardcoded patterns for clojure.core macros
- extension point with DSL for user-defined macros

(macro)compactor

attempt #1

- Kibit
- core.logic as a unification engine
- too slow
- no defined ordering is a problem

(macro)compactor

attempt #2

- core.match to pattern match over EDN
- fast!
- DSL over `match` to remove boilerplate
- no unification, use guards to enforce equality
- compact in postwalk over the source

(macro)compactor

attempt #2

```
(compact expr

[(`let [?t ?x]
  (if ?t ?y ?t))

{?t #(and (symbol? %)
           (-> % name (.startsWith "and_"))))

:-> `(and ~?x ~?y)]

...)
```

pprinter

- trim constant statements
- elide referred namespaces from symbols
- use aliases instead of long namespaces in symbols
- use fipp to pretty print the output in a readable format

down the rabbit hole

**java.io.IOException: File name too long, compiling:
(clojure/tools/decompiler/compact.clj:150:21)**

```
clojure.tools.decompiler.compact$macrocompact_step$fn_6117$fn_6118$fn_6119$fn_
_6120$fn_6121$fn_6122$fn_6123$fn_6124$fn_6125$fn_6126$fn_6127$fn_6128$
fn_6129$fn_6130$fn_6131$fn_6132$fn_6133$fn_6134$fn_6135$fn_6136$fn_613
7$fn_6138$fn_6139$fn_6140$fn_6141$fn_6142$fn_6143$fn_6144$fn_6145$fn_6
166$fn_6167$fn_6168$fn_6195$fn_6200$fn_6201$fn_6202$fn_6203$fn_6204$fn_
_6205$fn_6206$fn_6207$fn_6208$fn_6209$fn_6210$fn_6232$fn_6233$fn_6234$f
n_6235$fn_6238$fn_6241$fn_6244$fn_6247$fn_6250$fn_6251$fn_6254$fn_6255
$fn_6256$fn_6257$fn_6258$fn_6262$fn_6266$fn_6269$fn_6270$fn_6271$fn_62
74$fn_6278$fn_6279$fn_6280$fn_6284$fn_6285$fn_6286$fn_6287$fn_6300$fn_
6301$fn_6302$fn_6303.invoke()
```

down the rabbit hole

- workaround patch to Compiler.java
- it worked!
- but one thread kept 100% CPU

down the rabbit hole

- BCEL
- core.match
- fipp

down the rabbit hole

- ~~BCEL~~
- core.match
- fipp

down the rabbit hole

- ~~BCEL~~
- ~~core.match~~
- fipp

down the rabbit hole

- ~~BCEL~~
- ~~core.match~~
- ~~fipp~~

down the rabbit hole

- ~~BCEL~~
- ~~core.match~~
- ~~fipp~~
- clojure compiler?

down the rabbit hole

- ~~BCEL~~
- ~~core.match~~
- ~~fipp~~
- ~~closure compiler?~~

down the rabbit hole



down the rabbit hole

- ~~BCEL~~
- ~~core.match~~
- ~~fipp~~
- ~~closure compiler?~~
- JVM?

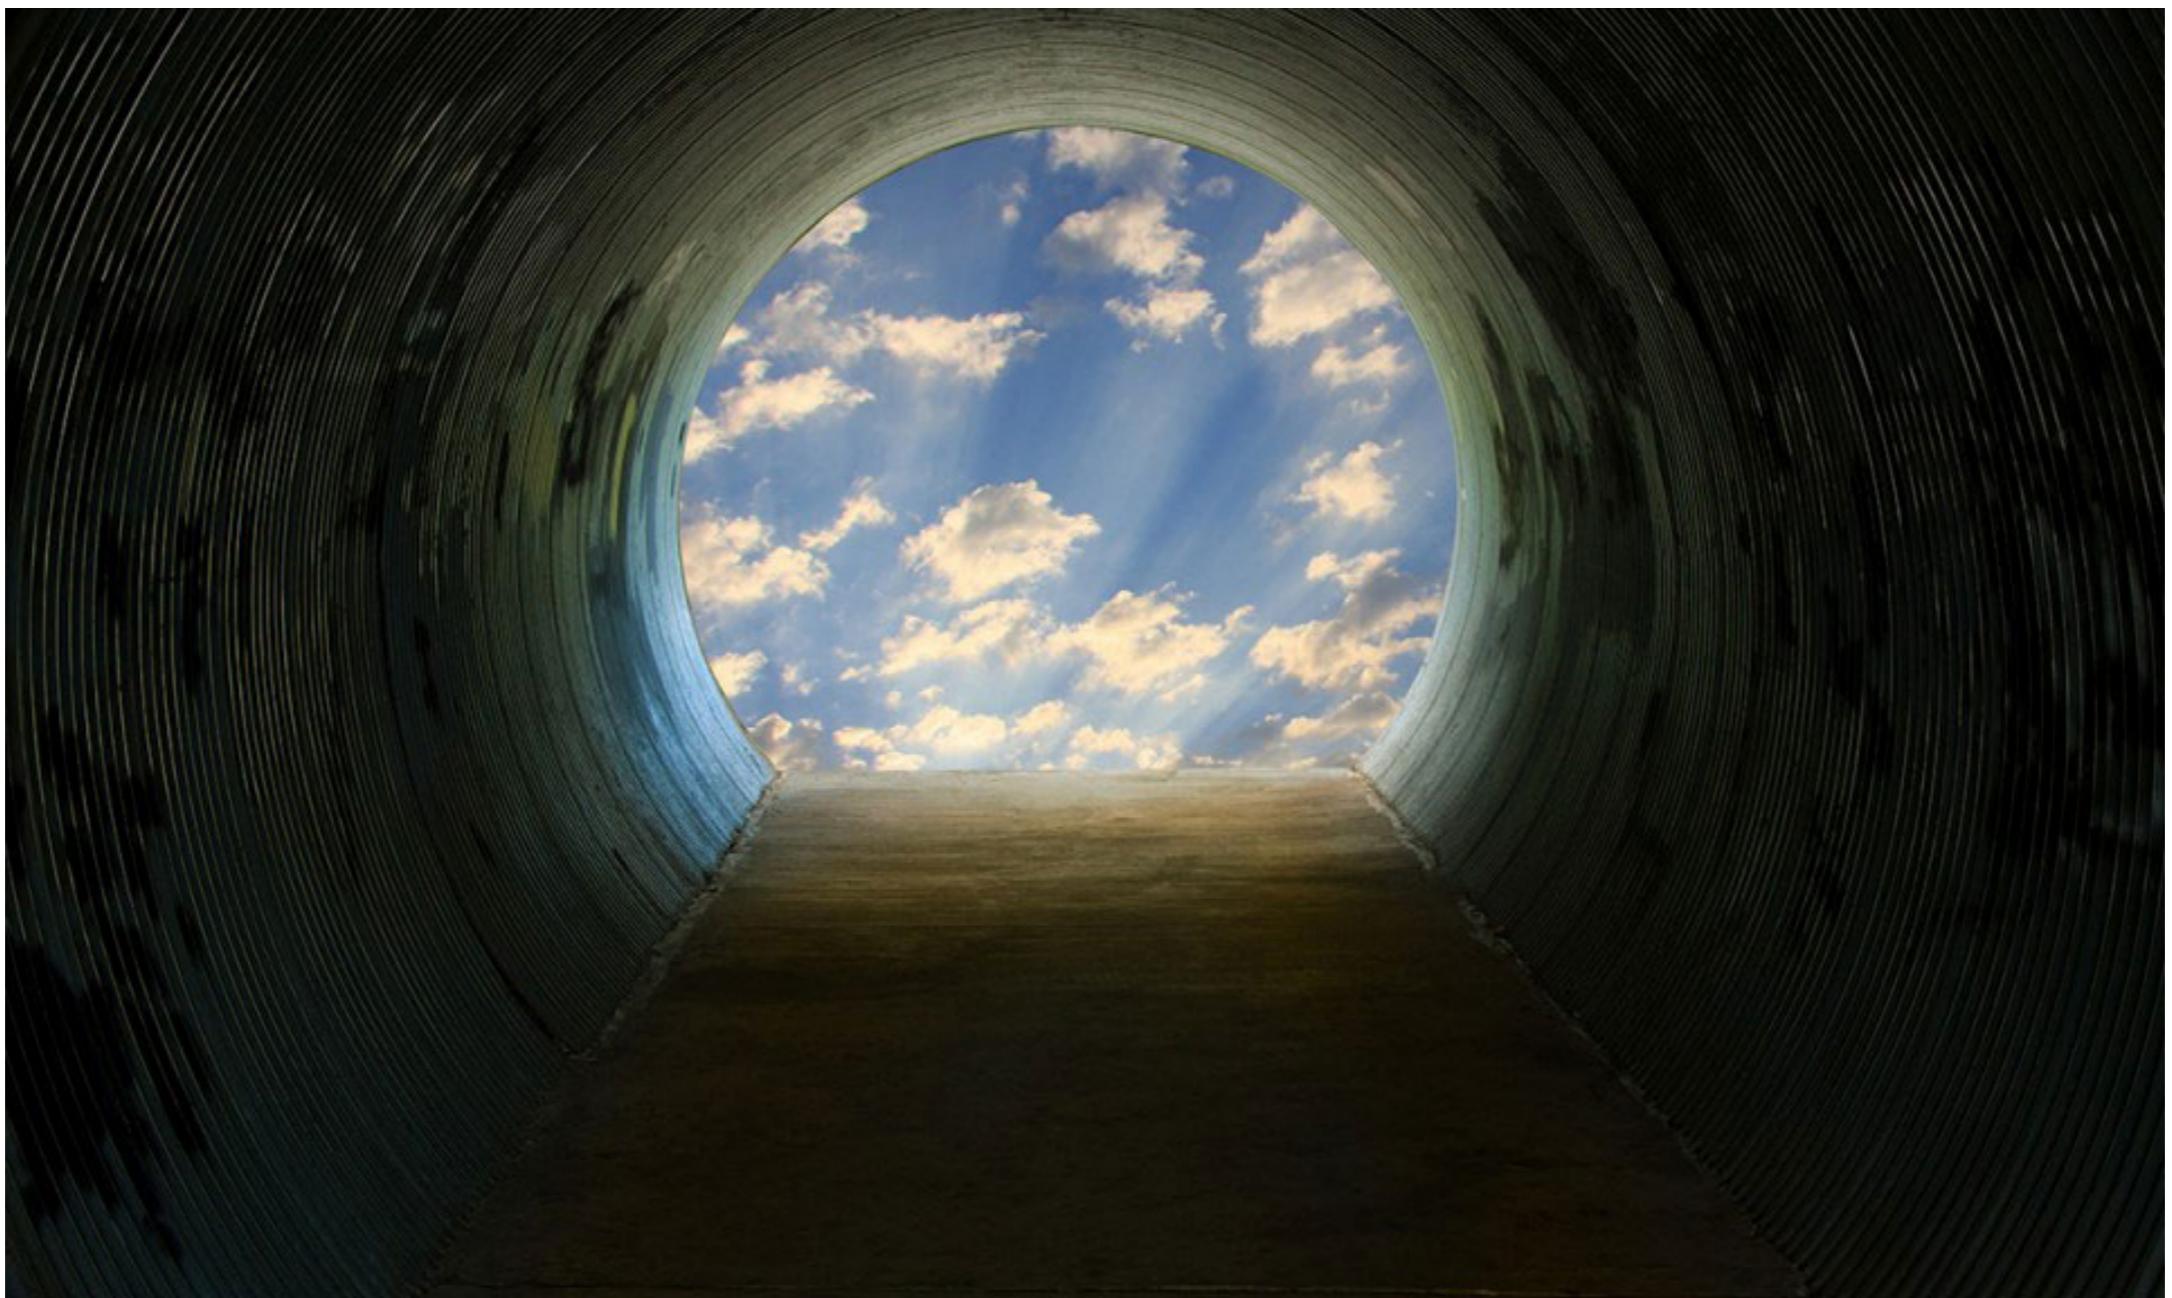
down the rabbit hole

```
2546 Thread_11839119 Java: C1 CompilerThread3
+ 2546 thread_start (in libsystem_thread.dylib) + 13 [0x7fffde25708c]
+ 2546 __pthread_start (in libsystem_pthread.dylib) + 86 [0x7fffde257887]
+ 2546 __pthread_body (in libsystem_pthread.dylib) + 180 [0x7fffde25793b]
+ 2546 java_start(Thread*) (in libjvm.dylib) + 246 [0x10908a5b2]
+ 2546 JavaThread::run() (in libjvm.dylib) + 450 [0x10916c1fc]
+ 2546 JavaThread::thread_main_inner() (in libjvm.dylib) + 155 [0x10916ab0f]
+ 2546 CompileBroker::compiler_thread_loop() (in libjvm.dylib) + 657 [0x108db8d0f]
+ 2546 CompileBroker::invoke_compiler_on_method(CompileTask*) (in libjvm.dylib) + 1458 [0x108db67c8]
+ 2546 Compiler::compile_method(ciEnv*, ciMethod*, int) (in libjvm.dylib) + 144 [0x108ce8e9a]
+ 2546 Compilation::Compilation(AbstractCompiler*, ciEnv*, ciMethod*, int, BufferBlob*) (in libjvm.dylib) + 418 [0x108ce8718]
+ 2546 Compilation::compile_method() (in libjvm.dylib) + 109 [0x108ce8503]
+ 2546 Compilation::compile_java_method() (in libjvm.dylib) + 88 [0x108ce82da]
+ 2546 Compilation::build_hir() (in libjvm.dylib) + 280 [0x108ce81bc]
+ 2546 IR::compute_code() (in libjvm.dylib) + 43 [0x108cf2d43]
+ 2546 ComputeLinearScanOrder::ComputeLinearScanOrder(Compilation*, BlockBegin*) (in libjvm.dylib) + 482 [0x108cf1fc]
+ 2546 ComputeLinearScanOrder::compute_order(BlockBegin*) (in libjvm.dylib) + 441 [0x108cfcd4d]
+ 2546 ComputeLinearScanOrder::compute_dominator(BlockBegin*, BlockBegin*) (in libjvm.dylib) + 121 [0x108cf36b]
+ 2546 ComputeLinearScanOrder::compute_dominator(BlockBegin*, BlockBegin*) (in libjvm.dylib) + 121 [0x108cf36b]
+ 2546 ComputeLinearScanOrder::compute_dominator(BlockBegin*, BlockBegin*) (in libjvm.dylib) + 121 [0x108cf36b]
+ 2546 ComputeLinearScanOrder::compute_dominator(BlockBegin*, BlockBegin*) (in libjvm.dylib) + 121 [0x108cf36b]
+ 2544 ComputeLinearScanOrder::compute_dominator(BlockBegin*, BlockBegin*) (in libjvm.dylib) + 121 [0x108cf36b]
+ ! 2540 ComputeLinearScanOrder::compute_dominator(BlockBegin*, BlockBegin*) (in libjvm.dylib) + 121 [0x108cf36b]
+ ! : 2525 ComputeLinearScanOrder::compute_dominator(BlockBegin*, BlockBegin*) (in libjvm.dylib) + 121 [...]
+ ! : | 2487 ComputeLinearScanOrder::compute_dominator(BlockBegin*, BlockBegin*) (in libjvm.dylib) + 121 [...]
+ ! : | + 2403 ComputeLinearScanOrder::compute_dominator(BlockBegin*, BlockBegin*) (in libjvm.dylib) + 121 [...]
+ ! : | + ! 2259 ComputeLinearScanOrder::compute_dominator(BlockBegin*, BlockBegin*) (in libjvm.dylib) + 121 [...]
```

down the rabbit hole

- core.match is implemented as a transpiler
- wrote new backend, continuations instead of exceptions for backtracking
- no more C1 issues
- after no-op continuation elision, 3x faster than default backend

light at the end of the tunnel



DEMO

limitations

- locals are munged, symbols could collide
- no type hints/mutable deftype fields printing support
- casts/widening missing or extra ones inserted
- genclass/definterface not properly supported
- not all macros are compacted
- some extra garbage is produced
- some known bugs when exceptions thrown in return context

future work

- detect and sugar syntax-quote expressions
- compact for/doseq/ns
- properly support type-hints/mutable deftype field decls
- use line table to guide pprinting

<https://github.com/Bronsa/tools.decompiler>

<https://github.com/Bronsa/tools.decompiler-aot>

fin.