

The Optimizer in Scala 2.12

Lukas Rytz, Scala Team @ Typesafe

Why?

- The JVM optimizer (HotSpot) has 17 years of tuning
- More powerful: run-time program statistics, speculative optimization
- javac does not perform *any* optimizations

HotSpot is not Perfect

- HotSpot fails to optimize well-known patterns
- Much more common in Scala – Java 8 catching up
- Recognized by JVM experts

Prominent Issues

- Megamorphic dispatch (JDK-8015416)
 - Click: The inlining problem
 - Rose: Profile pollution (talk on "JVM Challenges")
 - Shipilëv: The Black Magic of Method Dispatch
- Value boxing
 - Project Valhalla: specialization, value types
 - Rose's talk again

Megamorphic Callsites

```
class Range {  
  def foreach(f: Int => Unit) = {  
    while(..) { .. f.apply(i) .. }  
  }  
}
```

```
(1 to 10) foreach (x => foo)  
(2 to 20) foreach (x => bar)  
(3 to 30) foreach (x => baz)
```

Megamorphic Callsites

```
class Range {  
  def foreach(f: Int => Unit) = {  
    while(..) { .. f.apply(i) .. }  
  }  
}
```

Virtual call:

- Run-time type of `f` defines which code to run
- Megamorphic callsite, varying targets
- Method lookup on every loop iteration

```
(1 to 10) foreach (x => foo)  
(2 to 20) foreach (x => bar)  
(3 to 30) foreach (x => baz)
```

Solution: Inlining

```
(1 to 10) foreach (x => foo)
```

Scala optimizer inlines foreach

```
val _this = 1 to 10  
val _f = (x: Int) => foo  
while(..) { .. _f.apply(i) .. }
```

Solution: Inlining

```
(1 to 10) foreach (x => foo)
```


Scala optimizer inlines foreach

```
val _this = 1 to 10  
val _f = (x: Int) => foo  
while(..) { .. _f.apply(i) .. }
```

Monomorphic callsite enables JVM optimizations:

- Skip method lookup
- Inlining `apply` enables further optimizations

Megamorphic Callsites



```
while(..) {  
  .. (x => foo).apply() ..  
}  
while(..) {  
  .. (x => bar).apply() ..  
}  
...
```


github.com/lrytz/benchmarks

A Scala Problem?

```
class A(val x: Int) // similar: B, C, ..  
object AC extends Comparator[A] { .. }  
java.util.Arrays.sort(aArr, 0, N, AC)  
// java.util.Arrays.sort(bArr, 0, N, BC)
```

A Scala Problem?

```
class A(val x: Int) // similar: B, C, ..  
object AC extends Comparator[A] { .. }  
java.util.Arrays.sort(aArr, 0, N, AC)  
// java.util.Arrays.sort(bArr, 0, N, BC)
```



Value Boxing

```
var r = 0  
(1 to 10000) foreach { x => r += x }
```

```
val r = IntRef(0)  
val f = new anonfun(r)  
(1 to 10000) foreach f
```

```
class anonfun(r: IntRef) {  
  def apply(x: Int) {  
    r.elem += x  
  }  
}
```

Value Boxing

```
var r = 0  
(1 to 10000) foreach { x => r += x }
```

```
val r = IntRef(0)  
val f = new anonfun(r)  
(1 to 10000) foreach f
```

Slow

- Why? Not obvious..

```
class anonfun(r: IntRef) {  
  def apply(x: Int) {  
    r.elem += x  
  }  
}
```

Inlining

```
val r = IntRef(0)
val f = new anonfun(r)
(1 to 10000) foreach f
```

Inline foreach and function body

```
val r = IntRef(0)
val f = new anonfun(r)
var x = 0
while (x < 10000) {
  r.elem += x
}
```

Inlining

```
val r = IntRef(0)
val f = new anonfun(r)
(1 to 10000) foreach f
```

Inline foreach and function body

```
val r = IntRef(0)
val f = new anonfun(r)
var x = 0
while (x < 10000) {
  r.elem += x
}
```

Still slow (same as before)!

- Why? IntRef
- Escape analysis fails..

Closure Elimination

```
val r = IntRef(0)
val f = new anonfun(r)
var x = 0
while (x < 100000) {
  r.elem += x
}
```

Eliminate the closure allocation

```
val r = IntRef(0)
var x = 0
while (x < 100000) {
  r.elem += x
}
```


Closure Elimination

```
val r = IntRef(0)
val f = new anonfun(r)
var x = 0
while (x < 100000) {
  r.elem += x
}
```

Eliminate the closure allocation

```
val r = IntRef(0)
var x = 0
while (x < 100000) {
  r.elem += x
}
```

Fast! JVM escape analysis kicks in.

Box Elimination

```
val r = IntRef(0)
var x = 0
while (x < 10000) {
  r.elem += x
}
```

Local var instead of IntRef

```
var r = 0
var x = 0
while (x < 10000) {
  r += x
}
```

Box Elimination


```
val r = IntRef(0)
var x = 0
while (x < 10000) {
  r.elem += x
}
```

Local var instead of IntRef

```
var r = 0
var x = 0
while (x < 10000) {
  r += x
}
```

Same as before!
JVM optimizes the IntRef just fine.

Bars



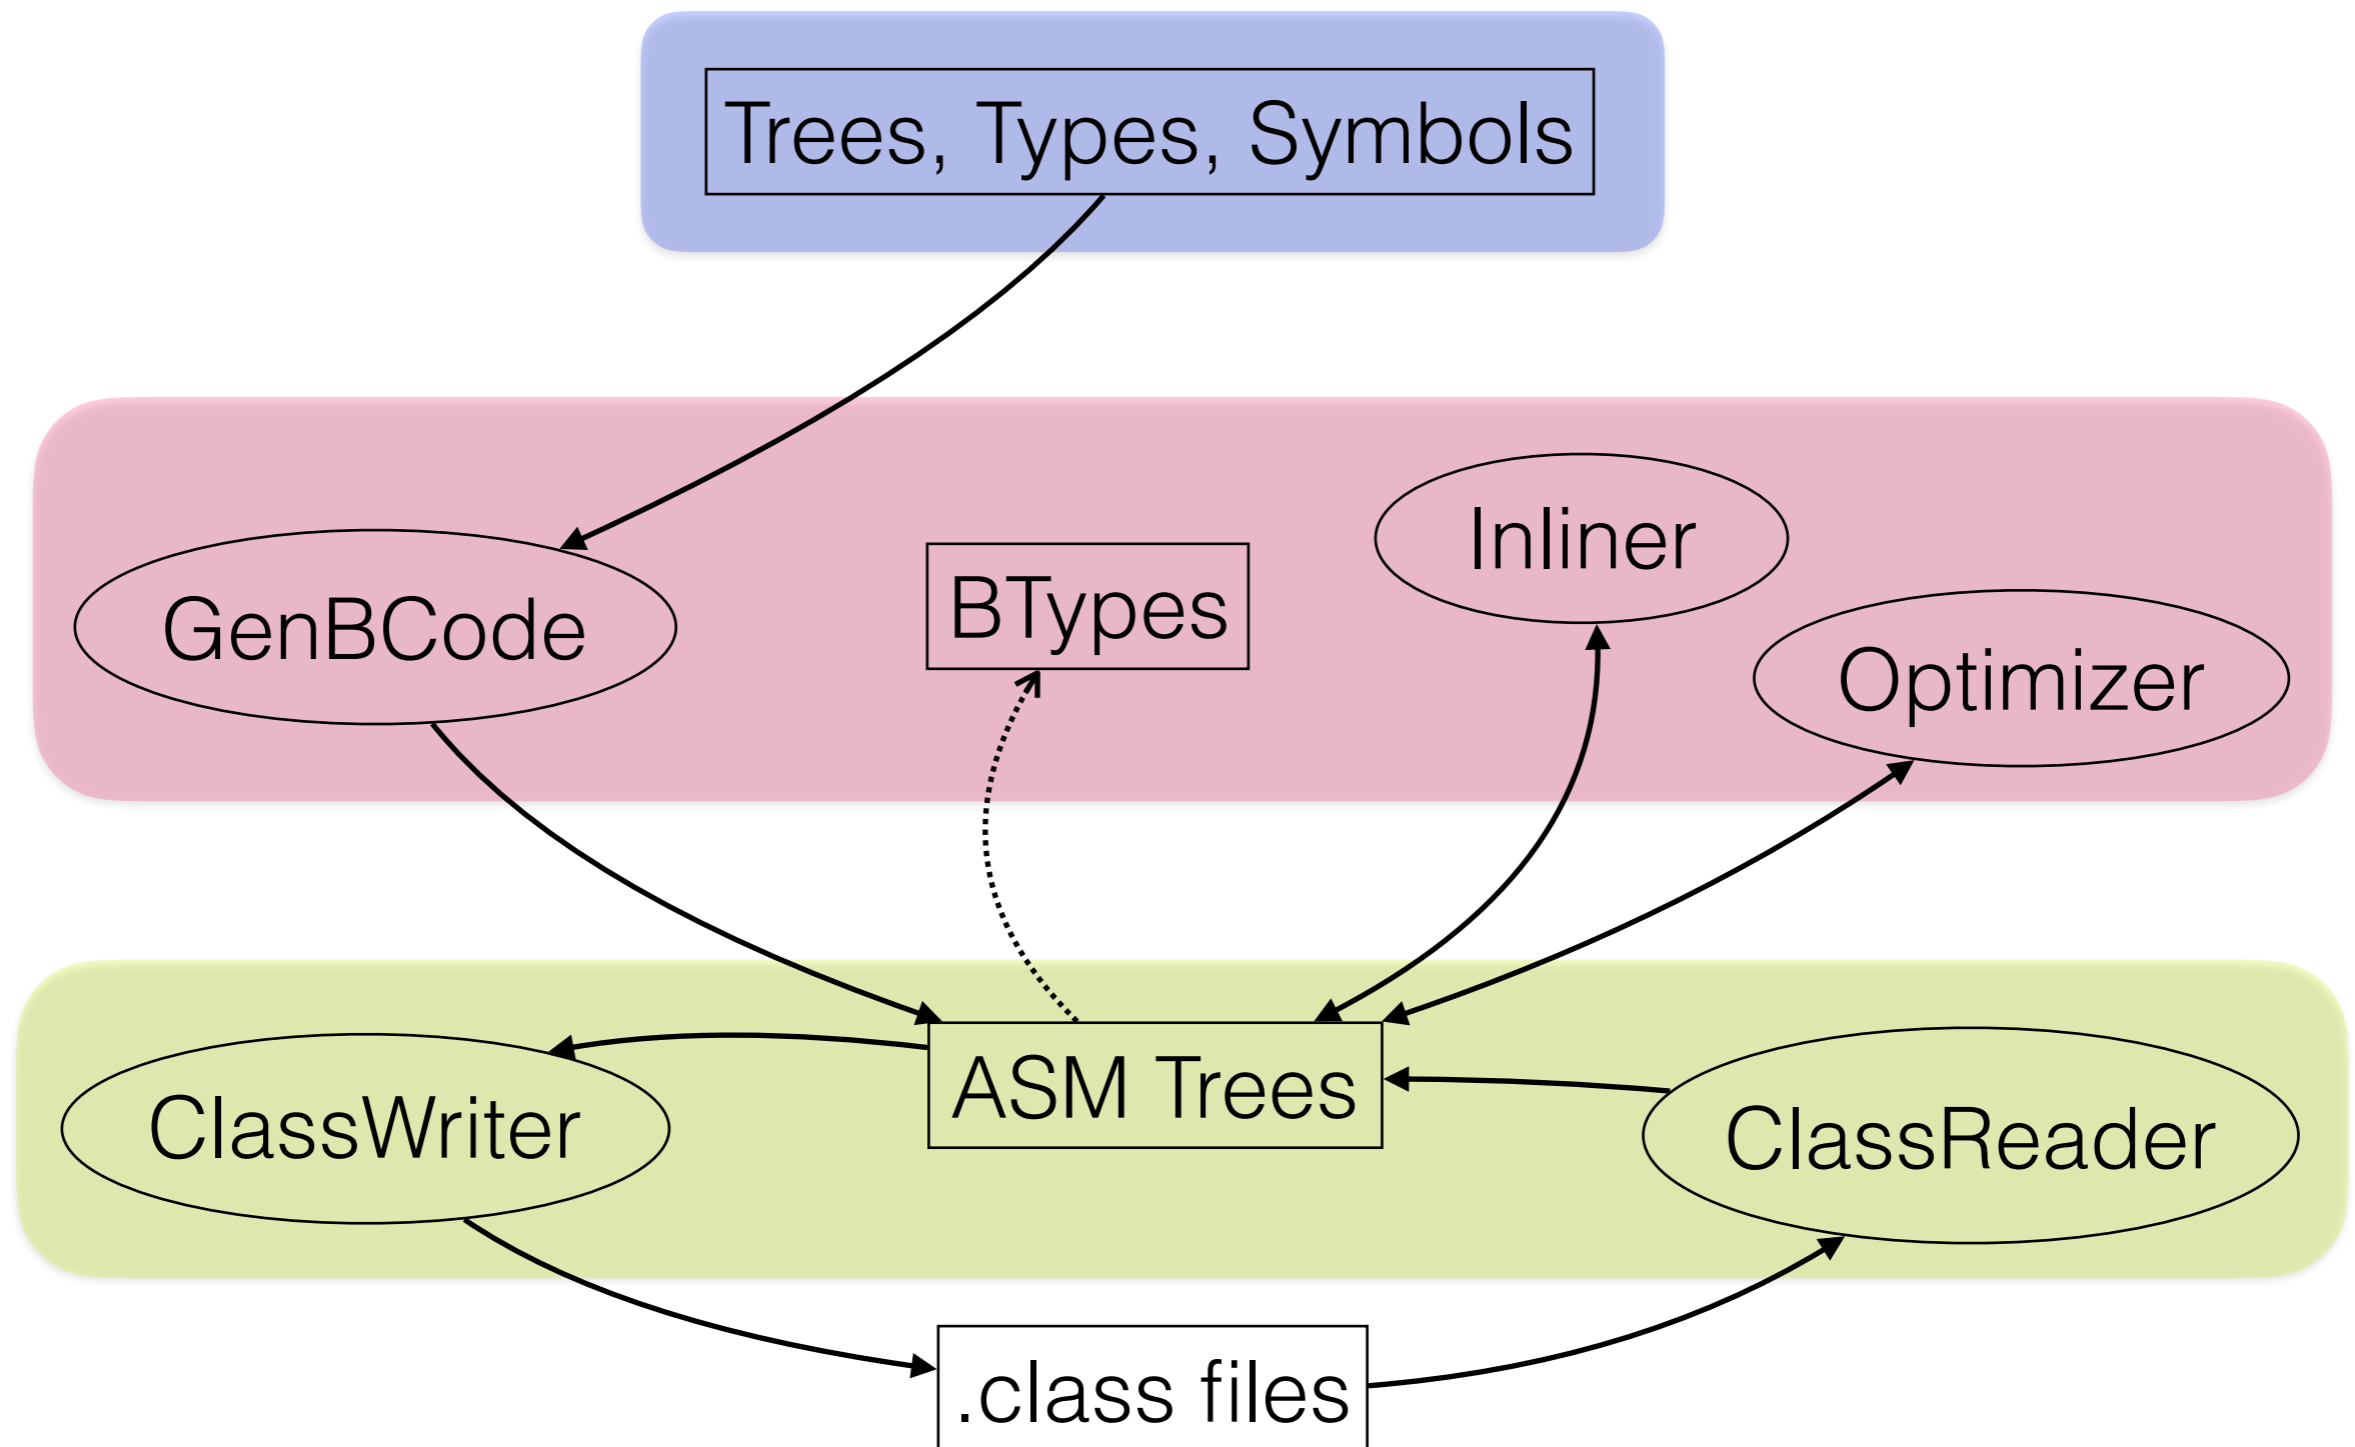
Compile-time Optimizer

- Goal: transform the code to make it please the JVM
- Don't perform optimizations that the JVM does well
- Avoid fruitless inlining: degrades performance
 - JVM optimizer is sensitive to method size

Agenda

- Compilation Pipeline Overview
- Local Optimizations
- Inlining and Heuristics
- Limitations
- Outlook, comparison with Scala.js

2.12 Backend: GenBCode



Local Optimizations

```
trait BooleanOrdering extends Ordering[Boolean] {  
  def compare(x: Boolean, y: Boolean) =  
    (x, y) match {  
      case (false, true) => -1  
      case (true, false) => 1  
      case _ => 0  
    }  
}
```


Nullness

```
Tuple2.mcZZ.sp sp2 = new Tuple2.mcZZ.sp(x, y);
if (sp2 != null) {
    boolean b1 = sp2._1$mcZ$sp();
    boolean b12 = sp2._2$mcZ$sp();
    if (!b1) {
        if (true == b12) return -1;
    }
}
if (sp2 == null) return 0;
boolean b1 = sp2._1$mcZ$sp();
boolean b13 = sp2._2$mcZ$sp();
if (!b1) return 0;
if (false == b13) return 1;
return 0;
```

Nullness

```
Tuple2.mcZZ.sp sp2 = new Tuple2.mcZZ.sp(x, y);  
if (sp2 != null) {  
    boolean b1 = sp2._1$mcZ$sp();  
    boolean b12 = sp2._2$mcZ$sp();  
    if (!b1) {  
        if (true == b12) return -1;  
    }  
}  
if (sp2 == null) return 0;  
boolean b1 = sp2._1$mcZ$sp();  
boolean b13 = sp2._2$mcZ$sp();  
if (!b1) return 0;  
if (false == b13) return 1;  
return 0;
```

Box-Unbox

```
int n;  
Tuple2.mcZZ.sp sp2 = new Tuple2.mcZZ.sp(x, y);  
boolean b1 = sp2._1.mcZ.sp();  
boolean b2 = sp2._2.mcZ.sp();  
if (!b1 && b2) {  
    n = -1;  
} else {  
    boolean b3 = sp2._1.mcZ.sp();  
    boolean b4 = sp2._2.mcZ.sp();  
    n = b3 && !b4 ? 1 : 0;  
}  
return n;
```

Box-Unbox

```
int n;  
Tuple2.mcZZ.sp sp2 = new Tuple2.mcZZ.sp(x, y);  
boolean b1 = sp2._1.mcZ.sp();  
boolean b2 = sp2._2.mcZ.sp();  
if (!b1 && b2) {  
    n = -1;  
} else {  
    boolean b3 = sp2._1.mcZ.sp();  
    boolean b4 = sp2._2.mcZ.sp();  
    n = b3 && !b4 ? 1 : 0;  
}  
return n;
```

Redundant Locals

```
int n;
boolean b1 = y;
boolean b12 = x;
Tuple2.mcZZ.sp sp2 = new Tuple2.mcZZ.sp(b12, b1);
boolean b13 = b12;
boolean b14 = b1;
if (!b13 && b14) {
    n = -1;
} else {
    boolean b15 = b12;
    boolean b16 = b1;
    n = b15 && !b16 ? 1 : 0;
}
return n;
```

Redundant Locals

```
int n;  
boolean b1 = y;  
boolean b12 = x;  
Tuple2.mcZZ.sp sp2 = new Tuple2.mcZZ.sp(b12, b1);  
boolean b13 = b12;  
boolean b14 = b1;  
if (!b13 && b14) {  
    n = -1;  
} else {  
    boolean b15 = b12;  
    boolean b16 = b1;  
    n = b15 && !b16 ? 1 : 0;  
}  
return n;
```

Unused Values

```
Tuple2.mcZZ.sp sp2 = new Tuple2.mcZZ.sp(x, y);  
int n = false == x && true == y ? -1 :  
      (true == x && false == y ? 1 : 0);  
return n;
```

Unused Values

```
Tuple2.mcZZ.sp sp2 = new Tuple2.mcZZ.sp(x, y);  
int n = false == x && true == y ? -1 :  
      (true == x && false == y ? 1 : 0);  
return n;
```


Unused Values

```
Tuple2.mcZZ.sp sp2 = new Tuple2.mcZZ.sp(x, y);  
int n = false == x && true == y ? -1 :  
      (true == x && false == y ? 1 : 0);  
return n;
```

```
int n = false == x && true == y ? -1 :  
      (true == x && false == y ? 1 : 0);  
return n;
```

Local Optimizations



Nullness

Unreachable Code

Box-Unbox

Redundant Locals

Redundant Casts

Unused Values

Simplify Jumps

Local Optimizations

Nullness

Unreachable Code

Box-Unbox

Primitive Boxes, Tuples, Refs

Redundant Locals

Redundant Casts

Unused Values

Simplify Jumps

Local Optimizations

Nullness

Unreachable Code

Box-Unbox

Redundant Locals

Redundant Casts

Unused Values

Simplify Jumps

Primitive Boxes, Tuples, Refs

Boxes, Closures

Closure Elimination

```
(1 to 10) foreach (x => foo)
```

```
val _this = 1 to 10  
val _f = (x: Int) => foo  
while(..) { .. _f.apply(i) .. }
```

```
val _this = 1 to 10  
while(..) { .. foo .. }
```

Agenda

- Compilation Pipeline Overview
- Local Optimizations
- Inlining and Heuristics
- Limitations
- Outlook, comparison with Scala.js

Inlining

- Local optimizations often enabled by inlining
 - Inlining `foreach` allows eliminating the closure
- Challenge: when to inline (heuristics)
 - Method size impacts JVM performance

Heuristics

- Goal: predict elimination of megamorphic callsites and value boxing
- How to identify callsites to inline?
 - Argument types (Functions, type classes)
 - Analysis of the callee (is an argument function invoked? passed to another call? captured?)
 - Backtracking
- Other considerations: method size, call frequency

Agenda

- Compilation Pipeline Overview
- Local Optimizations
- Inlining and Heuristics
- Limitations
- Outlook, comparison with Scala.js

Open World Assumption

- Inline only methods that cannot be overridden
- No whole-program analysis, support reflection
 - Don't know what are subclasses, overrides
 - Don't know what types are instantiated

A Magic Wand?

- The optimizer is by no means a magic wand
 - Most likely no speedup for an average program
- Instead: a tool for experts
 - Allows using high-level patterns in performance-critical code

Performance Example


- Richards from Scala.js benchmark suite
 - Scala.js optimizer: 3x speedup (on V8)
 - 2.12 optimizer: no speedup (JVM)
- TODO: find out why!
 - Scala.js heuristics better than ours?
 - Open-world assumption prevents optimizations?
 - JVM better than V8, more to gain for Scala.js?

Binary Compatibility

- Inlining from a library enforces a specific version
 - ➔ Assumption: consistent run-time classpath
- Problematic for library authors: forces specific versions for dependencies


Binary Compatibility

- Inlining from a library enforces a specific version
 - ➔ Assumption: consistent run-time classpath
- Problematic for library authors: forces specific versions for dependencies



Binary Compatibility

- Inlining from a library enforces a specific version
 - ➔ Assumption: consistent run-time classpath
- Problematic for library authors: forces specific versions for dependencies



Binary Compatibility

- Library authors: don't inline from the classpath
 - `Range.foreach` is slow
- Deployed applications: optimize freely
 - Ensure same classpath at runtime
 - Consider building dependencies from source

Agenda

- Compilation Pipeline Overview
- Local Optimizations
- Inlining and Heuristics
- Limitations
- Outlook, comparison with Scala.js

Global Optimization

- Assumptions for whole-program optimization:
 - Entire program available (including libraries)
 - Entry-point(s) known
 - Can change / remove code, no more linking
 - Restrictions in using reflection (may be "none")

Advantages

- Optimizer causes no binary compatibility issues
- Global knowledge: instantiated types, subclasses
 - Enables inlining more non-final callsites
- Examples:
 - Scala.js, dotty linker (in the works), dart2js
 - JVM, but more time-constrained

Thank You!