

Verified expression graph optimization in GraalVM

Formal Methods Australia 2022

Brae J. Webb Ian J. Hayes Mark Utting

6th June 2022

Outline

1. Introduce *GraalVM*.
2. Introduce and *motivate* our research.
3. *Live demo* some of our work.
4. Explain the Isabelle/HOL representation.

Question

What is **GraalVM**[™]?

GraalVM

A State-of-the-Art *Optimizing* Compiler.

Key points on GraalVM

1. Open-source and written in Java.¹

¹<https://github.com/oracle/graalvm>

²Multiple languages

Key points on GraalVM

1. Open-source and written in Java.¹
2. Polyglot² compiler through partial evaluation.

¹<https://github.com/oracle/graalvm>

²Multiple languages

Key points on GraalVM

1. Open-source and written in Java.¹
2. Polyglot² compiler through partial evaluation.
3. JIT and AOT compilation modes.

¹<https://github.com/oracle/graalvm>

²Multiple languages

Our goal

Verify the correctness of *optimizations* in the *GraalVM* compiler.

Two questions

1. Why verify the *optimizations*?
2. Why verify *GraalVM*?

Why verify the optimizations?

Compiler bugs are *bad*.

Compiler bugs are bad



Why verify the *optimizations*? [Yang et al., 2011]

GCC $\frac{75}{79} \approx 90\%$

Clang $\frac{183}{202} \approx 95\%$

Two questions

1. Why verify the *optimizations*?
2. Why verify *GraalVM*?

Two questions

1. Why verify the *optimizations*?
2. Why verify *GraalVM*?

Why verify GraalVM?

1. Comprehensive optimization suite.

Why verify GraalVM?

1. Comprehensive optimization suite.
2. Actively developed.

Why verify GraalVM?

1. Comprehensive optimization suite.
2. Actively developed.
3. Unique/novel/challenging IR.

Why verify GraalVM?

1. Comprehensive optimization suite.
2. Actively developed.
3. Unique/novel/challenging IR.
4. Partial evaluation relies heavily on optimizations.

Two questions

1. Why verify the *optimizations*?
2. Why verify *GraalVM*?

Bonus question

Why *verify* compilers at all?

Why verify the *optimizations*? [Yang et al., 2011]

GCC $\frac{75}{79} \approx 90\%$

Clang $\frac{183}{202} \approx 95\%$

Why verify the *optimizations*? [Yang et al., 2011]

GCC $\frac{75}{79} \approx 90\%$

Clang $\frac{183}{202} \approx 95\%$

CompCert [Leroy et al., 2016]
 $\frac{0}{2} \approx 0\%$

Why verify the *optimizations*? [Yang et al., 2011]

GCC $\frac{75}{79} \approx 90\%$

Clang $\frac{183}{202} \approx 95\%$

CompCert [Leroy et al., 2016]

$$\frac{0}{0} \approx \text{undefined}$$

Question

How do we verify an optimization?

Question

How do we verify an optimization?

Answer

Use *Isabelle/HOL* to show that the optimized program refines the unoptimized program, thus the optimization is *semantics preserving*.

Steps to prove an optimization

1. A definition of what *any program* will do in *any state*.
2. Demonstrate that for *all states*: The unoptimized program will *behave the same* as the optimized program.
i.e. semantics are preserved.

* Where we say *program* we mean the *intermediate representation* of the program

Steps to prove an optimization

1. ~~A definition of what *any program* will do in *any state*.~~

[Webb et al., 2021]

2. Demonstrate that for *all states*: The unoptimized program will *behave the same* as the optimized program.

i.e. semantics are preserved.

* Where we say *program* we mean the *intermediate representation* of the program

These proofs were hard...

Question

Why were these proofs hard?

Back to GraalVM...

A simple program


```
» cat PositiveAdd.java
```

```
1 static int positiveAdd(int x, int y) {  
2     if (((x - y) + y) - (x - y)) > 0) {  
3         return x + y;  
4     } else {  
5         return x - y;  
6     }  
7 }
```

A simple program in GraalVM

```
» cat PositiveAdd.java
```


```
1 static int positiveAdd(int x,  
2   int y) {  
3   if (((x - y) + y) - (x - y)  
4     ) > 0) {  
5     return x + y;  
6   } else {  
7     return x - y;  
8   }  
9 }
```



A simple program in GraalVM

```
» cat PositiveAdd.java
```

```
1 static int positiveAdd(int x,
2     int y) {
3     if (((x - y) + y) - (x - y)
4         ) > 0) {
5         return x + y;
6     } else {
7         return x - y;
8     }
9 }
```



A mess!

Question

Why the mess?

Question


Why the mess?

Answer

It's very efficient *[Click and Cooper, 1995]*.

Common subexpressions (or subgraphs) of the program can be *shared*.
Each unique expression is evaluated and stored once.

Important points on the GraalVM IR



- Expressions and control-flow are combined into one graph structure.
- Common sub-expressions are shared.
- Control-flow and data-flow (expressions) can have cycles.

Question

Why were these proofs hard?

Question


Why were these proofs hard?

Answer


Graph updates can affect any expressions that share the updated expression.

We have to explicitly show that potential *cycles and self-reference* maintain semantic preservation.

Optimizing a graph



Optimizing a graph




Optimizing a graph

Every optimization proof requires an *inductive proof* to satisfy *self-reference*.


Solution

Represent expressions as *trees* rather than *graphs*.


Recall our program graph



Removing sharing offers tree structures



$\mapsto$



Added benefit


Trees enable natural infix encoding of optimization rules.

Added benefit

Trees enable natural infix encoding of optimization rules.

e.g. $x + 0 \longmapsto x$

How can we optimize these trees?



How about?

1. $(x - y) + y \longmapsto x$

2. $x - (x - y) \longmapsto y$

Wait...

Are these optimizations *correct*?

Wait...

Are these optimizations *correct*?

What about *integer overflow*?

Wait...

Are these optimizations *correct*?

What about *integer overflow*?

What about *floating-point* arithmetic?

Wait...

Are these optimizations *correct*?

What about *integer overflow*?

What about *floating-point* arithmetic?


What about *side-effecting* operations?

Live demo


1. Show $(x - y) + y \longmapsto x$
2. Show $x - (x - y) \longmapsto y$

Now we have proofs

Let's apply the optimizations!



$$(x - y) + y \mapsto x$$



$$x - (x - y) \mapsto y$$

The fun part

How we represent and prove these optimizations.

Expression Trees

```
datatype IRExpr =  
  UnaryExpr IRUnaryOp IRExpr  
  | BinaryExpr IRBinaryOp IRExpr IRExpr  
  | ConditionalExpr IRExpr IRExpr IRExpr  
  | ParameterExpr nat Stamp  
  | LeafExpr nat Stamp  
  | ConstantExpr Value
```

Expression Trees

datatype *IRExpr* =

UnaryExpr *IRUnaryOp* *IRExpr*

| *BinaryExpr* *IRBinaryOp* *IRExpr* *IRExpr*

| *ConditionalExpr* *IRExpr* *IRExpr* *IRExpr* ...

| *ParameterExpr* *nat* *Stamp*

| *LeafExpr* *nat* *Stamp*

| *ConstantExpr* *Value*

datatype *IRUnaryOp* =

UnaryAbs

| *UnaryNeg*

| *UnaryNot*

datatype *IRBinaryOp* =

BinAdd

| *BinMul*

| *BinSub*

...

Expression Semantics

Evaluating expressions depends on


- p** a list of values of the parameters to the method
- m** a mapping storing already evaluated leaf nodes
 - such as results from method calls

Expression Semantics

Evaluating expressions depends on

- p a list of values of the parameters to the method
- m a mapping storing already evaluated leaf nodes
 - such as results from method calls

Evaluating expression e to value v is represented by



Expression Semantics

$$\frac{\text{valid-value } (\text{constantAsStamp } c) \ c}{[m,p] \vdash \text{ConstantExpr } c \mapsto c} \quad (1)$$

$$\frac{i < |p| \quad \text{valid-value } s \ p[i]}{[m,p] \vdash \text{ParameterExpr } i \ s \mapsto p[i]} \quad (2)$$

$$\frac{[m,p] \vdash xe \mapsto v \quad \text{unary-eval } op \ v \neq \text{UndefVal}}{[m,p] \vdash \text{UnaryExpr } op \ xe \mapsto \text{unary-eval } op \ v} \quad (3)$$

$$\frac{[m,p] \vdash xe \mapsto x \quad [m,p] \vdash ye \mapsto y \quad \text{bin-eval } op \ x \ y \neq \text{UndefVal}}{[m,p] \vdash \text{BinaryExpr } op \ xe \ ye \mapsto \text{bin-eval } op \ x \ y} \quad (4)$$

$$\frac{\text{val} = m \ n \quad \text{valid-value } s \ \text{val}}{[m,p] \vdash \text{LeafExpr } n \ s \mapsto \text{val}} \quad (5)$$

Expression refinement

For $e2$ to be an correct optimization of $e1$ we require that

- whenever $e1$ evaluates to a value v in some context $[m, p]$, so does $e2$.
- $e2$ *refines* $e1$,

$$(e2 \leq e1) = (\forall m \ p \ v. [m, p] \vdash e1 \mapsto v \longrightarrow [m, p] \vdash e2 \mapsto v)$$

- $e2$ may be well formed in more contexts than $e1$, e.g. $x - x \geq 0$

Expression Semantics

$$\frac{\text{valid-value } (\text{constantAsStamp } c) \ c}{[m,p] \vdash \text{ConstantExpr } c \mapsto c} \quad (6)$$

$$\frac{i < |p| \quad \text{valid-value } s \ p[i]}{[m,p] \vdash \text{ParameterExpr } i \ s \mapsto p[i]} \quad (7)$$

$$\frac{[m,p] \vdash xe \mapsto v \quad \text{unary-eval } op \ v \neq \text{UndefVal}}{[m,p] \vdash \text{UnaryExpr } op \ xe \mapsto \text{unary-eval } op \ v} \quad (8)$$

$$\frac{[m,p] \vdash xe \mapsto x \quad [m,p] \vdash ye \mapsto y \quad \text{bin-eval } op \ x \ y \neq \text{UndefVal}}{[m,p] \vdash \text{BinaryExpr } op \ xe \ ye \mapsto \text{bin-eval } op \ x \ y} \quad (9)$$

$$\frac{\text{val} = m \ n \quad \text{valid-value } s \ \text{val}}{[m,p] \vdash \text{LeafExpr } n \ s \mapsto \text{val}} \quad (10)$$

Value Semantics

fun *intval-add* :: *Value* $\Rightarrow$ *Value* $\Rightarrow$ *Value* **where**

intval-add (*IntVal32* *v1*) (*IntVal32* *v2*) = (*IntVal32* (*v1*+*v2*)) |

intval-add (*IntVal64* *v1*) (*IntVal64* *v2*) = (*IntVal64* (*v1*+*v2*)) |

intval-add - - = *UndefVal*

fun *intval-xor* :: *Value* $\Rightarrow$ *Value* $\Rightarrow$ *Value* **where**

intval-xor (*IntVal32* *v1*) (*IntVal32* *v2*) = (*IntVal32* (*v1 XOR v2*)) |

intval-xor (*IntVal64* *v1*) (*IntVal64* *v2*) = (*IntVal64* (*v1 XOR v2*)) |

intval-xor - - = *UndefVal*


fun *bin-eval* :: *IRBinaryOp* $\Rightarrow$ *Value* $\Rightarrow$ *Value* $\Rightarrow$ *Value* **where**

bin-eval *BinAdd* *v1* *v2* = *intval-add* *v1* *v2* |

bin-eval *BinXor* *v1* *v2* = *intval-xor* *v1* *v2* |

...

Our Proof Approach



With those definitions

We can start proving!

In summary

- enables more *concise specification* of optimization rules,

In summary

- enables more *concise specification* of optimization rules,
- *reduces the proof burden* by a factor of 5 (by line count), while proving a significantly stronger property, and

In summary

- enables more *concise specification* of optimization rules,
- *reduces the proof burden* by a factor of 5 (by line count), while proving a significantly stronger property, and
- includes *termination* as a proof-obligation.

Future work

- Expand on *proof automation*.

Future work

- Expand on *proof automation*.
- Extend the DSL to *factor out common side conditions*.

Future work

- Expand on *proof automation*.
- Extend the DSL to *factor out common side conditions*.
- Resume work *control-flow optimizations*.

Thank you!

Any questions?

References

- [Click and Cooper, 1995] Click, C. and Cooper, K. D. (1995).
Combining analyses, combining optimizations.
TOPLAS, 17(2):181–196.
- [Leroy et al., 2016] Leroy, X., Blazy, S., Kästner, D., Schommer, B., Pister, M.,
and Ferdinand, C. (2016).
CompCert - A Formally Verified Optimizing Compiler.
In *ERTS 2016*, Toulouse, France. SEE.
- [Webb et al., 2021] Webb, B. J., Utting, M., and Hayes, I. J. (2021).
A formal semantics of the GraalVM intermediate representation.
In Hou, Z. and Ganesh, V., editors, *Automated Technology for Verification and Analysis*, volume 12971 of *Lecture Notes in Computer Science*, pages 111–126,
Cham. Springer International Publishing.

[Yang et al., 2011] Yang, X., Chen, Y., Eide, E., and Regehr, J. (2011).

Finding and understanding bugs in C compilers.

In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '11, page 283–294, New York, NY, USA. Association for Computing Machinery.