

A Language for Describing Optimization Strategies

Bastian Hagedorn¹ | Johannes Lenfers¹ | Thomas Koehler² | Sergei Gorlatch¹ | Michel Steuwer² ¹University of Münster | ²University of Glasgow

The Landscape of Optimizing Compilers



The Landscape of Optimizing Compilers



The Landscape of Optimizing Compilers



The Landscape of Optimizing Compilers



"Somewhere between -00 and -02"

The Landscape of Optimizing Compilers



General Purpose Compilers



-targetlibinfo -tti -tbaa -scoped-noalias -assumption-cache-tracker -profile-summarv-info -forceattrs -inferattrs -callsite-splitting psccp -called-value-propagation -globalopt -domtree -mem2reg -deadargelim -domtree -basicaa -aa -loops -lazy-branch-prob -lazy-block-fre q -opt-remark-emitter -instcombine -simplifycfg -basiccg -globals-aa -prune-eh -inline -functionattrs -argpromotion -domtree -sroa -basi caa -aa -memoryssa -early-cse-memssa -speculative-execution -basicaa -aa -lazy-value-info -jump-threading -correlated-propagation -simpl ifycfg -domtree -aggressive-instcombine -basicaa -aa -loops -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -instcombine -libcal s-shrinkwrap -loops -branch-prob -block-freq -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -pgo-memop-opt -basicaa -aa -loops lazy-branch-prob -lazy-block-freq -opt-remark-emitter -tailcallelim -simplifycfg -reassociate -domtree -loops -loop-simplify -lcssa-veri fication -lcssa -basicaa -aa -scalar-evolution -loop-rotate -licm -loop-unswitch -simplifycfg -domtree -basicaa -aa -loops -lazy-branchprob -lazy-block-freq -opt-remark-emitter -instcombine -loop-simplify -lcssa-verification -lcssa -scalar-evolution -indvars -loop-idiom -loop-deletion -loop-unroll -mldst-motion -phi-values -basicaa -aa -memdep -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -gvn phi-values -basicaa -aa -memdep -memopyopt -scop -demanded-bits -bdce -basicaa -aa -loops -lazy-branch-prob -lazy-block-freg -opt-remar -emitter -instcombine -lazy-value-info -jump-threading -correlated-propagation -basicaa -aa -phi-values -memdep -dse -loops -loop-simpli fy -lcssa-verification -lcssa -basicaa -aa -scalar-evolution -licm -postdomtree -adce -simplifycfg -domtree -basicaa -aa -loops -lazy-br anch-prob -lazy-block-freq -opt-remark-emitter -instcombine -barrier -elim-avail-extern -basiccg -rpo-functionattrs -globalopt -globaldo e -basiccg -globals-aa -float2int -domtree -loops -loop-simplify -lcssa-verification -lcssa -basicaa -aa -scalar-evolution -loop-rotate -loop-accesses -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -loop-distribute -branch-prob -block-freq -scalar-evolution -basi aa -aa -loop-accesses -demanded-bits -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -loop-vectorize -loop-simplify -scalar-evolu tion -aa -loop-accesses -loop-load-elim -basicaa -aa -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -instcombine -simplifycfg -d omtree -loops -scalar-evolution -basicaa -aa -demanded-bits -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -slp-vectorizer -optremark-emitter -instcombine -loop-simplify -lcssa-verification -lcssa -scalar-evolution -loop-unroll -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -instcombine -loop-simplify -lcssa-verification -lcssa -scalar-evolution -licm -alignment-from-assumptions -strip-de ad-prototypes -globaldce -constmerge -domtree -loops -branch-prob -block-freq -loop-simplify -lcssa-verification -lcssa -basicaa -aa -so alar-evolution -branch-prob -block-freq -loop-sink -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -instsimplify -div-rem-pairs simplifycfg -verify

-03

The Landscape of Optimizing Compilers

a

General Purpose Compilers



-targetlibinfo -tti -tbaa -scoped-noalias -assumption-cache-tracker -profile-summar psccp -called-value-propagation -globalopt -domtree -mem2reg -deadargelim -domtree q -opt-remark-emitter -instcombine -simplifycfg -basiccg -globals-aa -prune-eh -inl caa -aa -memoryssa -early-cse-memssa -speculative-execution -basicaa -aa -lazy-valu ifycfg -domtree -aggressive-instcombine -basicaa -aa -loops -lazy-branch-prob -lazy s-shrinkwrap -loops -branch-prob -block-freq -lazy-branch-prob -lazy-block-freq -op lazy-branch-prob -lazy-block-freg -opt-remark-emitter -tailcallelim -simplifycfg fication -lcssa -basicaa -aa -scalar-evolution -loop-rotate -licm -loop-unswitch -s prob -lazy-block-freq -opt-remark-emitter -instcombine -loop-simplify -lcssa-verifi -loop-deletion -loop-unroll -mldst-motion -phi-values -basicaa -aa -memdep -lazy-br phi-values -basicaa -aa -memdep -memcpyopt -sccp -demanded-bits -bdce -basicaa -aa -emitter -instcombine -lazy-value-info -jump-threading -correlated-propagation -bas fy -lcssa-verification -lcssa -basicaa -aa -scalar-evolution -licm -postdomtree -ad anch-prob -lazy-block-freq -opt-remark-emitter -instcombine -barrier -elim-avail-ex e -basiccg -globals-aa -float2int -domtree -loops -loop-simplify -lcssa-verificatio -loop-accesses -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -loop-distrib aa -aa -loop-accesses -demanded-bits -lazy-branch-prob -lazy-block-freq -opt-remark tion -aa -loop-accesses -loop-load-elim -basicaa -aa -lazy-branch-prob -lazy-blockomtree -loops -scalar-evolution -basicaa -aa -demanded-bits -lazy-branch-prob -lazy remark-emitter -instcombine -loop-simplify -lcssa-verification -lcssa -scalar-evolu -opt-remark-emitter -instcombine -loop-simplify -lcssa-verification -lcssa -scalarad-prototypes -globaldce -constmerge -domtree -loops -branch-prob -block-freq -loop alar-evolution -branch-prob -block-freq -loop-sink -lazy-branch-prob -lazy-block-fr simplifycfg -verify

-03

te-splitting -i

lazy-block-fre

ee -sroa -basi

agation -simpl mbine -libcall

aa -aa -loops

fy -lcssa-veri

-lazy-branch

s -loop-idiom

-emitter -gvn

reg -opt-remark

os -loop-simpli

loops -lazy-bi

alopt -globaldo

-loop-rotate

volution -basi

-scalar-evolu

simplifycfg -

ectorizer -opt

zy-block-freq

tions -strip-de basicaa -aa -se

div-rem-pairs

Compiler Passes

The Landscape of Optimizing Compilers

a

General Purpose Compilers



-targetlibinfo -tti -tbaa -scoped-noalias -assumption-cache-tracker -profile-summar psccp -called-value-propagation -globalopt -domtree -mem2reg -deadargelim -domtree g -opt-remark-emitter -<mark>instcombine</mark> -simplifycfg -basiccg -globals-aa -prune-eh -inl caa -aa -memoryssa -early-cse-memssa -speculative-execution -basicaa -aa -lazy-valu ifycfg -domtree -aggressive-<mark>instcombine</mark> -basicaa -aa -loops -lazy-branch-prob -lazy s-shrinkwrap -loops -branch-prob -block-freq -lazy-branch-prob -lazy-block-freq -op lazy-branch-prob -lazy-block-freq -opt-remark-emitter -tailcallelim -simplifycfg -r fication -lcssa -basicaa -aa -scalar-evolution -loop-rotate -licm -loop-unswitch prob -lazy-block-freq -opt-remark-emitter -<mark>instcombine</mark> -loop-simplify -lcssa-verifi -loop-deletion -loop-unroll -mldst-motion -phi-values -basicaa -aa -memdep -lazy-br phi-values -basicaa -aa -memdep -memcpyopt -sccp -demanded-bits -bdce -basicaa -aa -emitter -instcombine -lazy-value-info -jump-threading -correlated-propagation -bas fy -lcssa-verification -lcssa -basicaa -aa -scalar-evolution -licm -postdomtree -ad anch-prob -lazy-block-freq -opt-remark-emitter -instcombine -barrier -elim-avail-ex e -basiccg -globals-aa -float2int -domtree -loops -loop-simplify -lcssa-verificatio -loop-accesses -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -loop-distrib aa -aa -loop-accesses -demanded-bits -lazy-branch-prob -lazy-block-freq -opt-remark tion -aa -loop-accesses -loop-load-elim -basicaa -aa -lazy-branch-prob -lazy-blockomtree -loops -scalar-evolution -basicaa -aa -demanded-bits -lazy-branch-prob -lazy remark-emitter -<mark>instcombine</mark> -loop-simplify -lcssa-verification -lcssa -scalar-evolu -opt-remark-emitter -instcombine -loop-simplify -lcssa-verification -lcssa -scalarad-prototypes -globaldce -constmerge -domtree -loops -branch-prob -block-freg -loop alar-evolution -branch-prob -block-freq -loop-sink -lazy-branch-prob -lazy-block-fr simplifycfg -verify

-03

aa -aa -loops reg -opt-remark -loops -lazy-b alopt -globald volution -basi -simplifycfg ectorizer -opt azy-block-freq tions -strip-de basicaa -aa -s div-rem-pairs

te-splitting -:

lazy-block-fre

ree -sroa -bas:

agation -simp

bine -libcal

ify -lcssa-ver

-lazy-branch

rs -loop-idiom

-emitter -gvn

os -loop-simpl:

1 -loop-rotate

-scalar-evol

Compiler Passes

The Landscape of Optimizing Compilers

 -targetlibinfo -tti -tbaa -scoped-noalias -assumption-cache-tracker -profile-summa psccp -called-value-propagation -globalopt -domtree -mem2reg -deadargelim -domtree q -opt-remark-emitter -<mark>instcombine</mark> -simplifycfg -basiccg -globals-aa -prune-eh -inl caa -aa -memoryssa -early-cse-memssa -speculative-execution -basicaa -aa -lazy-valu ifycfg -domtree -aggressive-<mark>instcombine</mark> -basicaa -aa -loops -lazy-branch-prob -lazy s-shrinkwrap -loops -branch-prob -block-freq -lazy-branch-prob -lazy-block-freq -op lazy-branch-prob -lazy-block-freq -opt-remark-emitter -tailcallelim -simplifycfg -r fication -lcssa -basicaa -aa -scalar-evolution -loop-rotate -licm -loop-unswitch prob -lazy-block-freq -opt-remark-emitter -<mark>instcombine</mark> -loop-simplify -lcssa-verifi -loop-deletion -loop-unroll -mldst-motion -phi-values -basicaa -aa -memdep -lazy-br phi-values -basicaa -aa -memdep -memcpyopt -sccp -demanded-bits -bdce -basicaa -aa -emitter -instcombine -lazy-value-info -jump-threading -correlated-propagation -bas fy -lcssa-verification -lcssa -basicaa -aa -scalar-evolution -licm -postdomtree -ad anch-prob -lazy-block-freq -opt-remark-emitter -<mark>instcombine</mark> -barrier -elim-avail-ex e -basiccg -globals-aa -float2int -domtree -loops -loop-simplify -lcssa-verificatio -loop-accesses -lazy-branch-prob -lazy-block-freq -opt-remark-emitter -loop-distrib aa -aa -loop-accesses -demanded-bits -lazy-branch-prob -lazy-block-freq -opt-remark tion -aa -loop-accesses -loop-load-elim -basicaa -aa -lazy-branch-prob -lazy-blockomtree -loops -scalar-evolution -basicaa -aa -demanded-bits -lazy-branch-prob -lazy remark-emitter -<mark>instcombine</mark> -loop-simplify -lcssa-verification -lcssa -scalar-evolu -opt-remark-emitter <mark>-instcombine</mark> -loop-simplify -lcssa-verification -lcssa -scalar ad-prototypes -globaldce -constmerge -domtree -loops -branch-prob -block-freg -loop alar-evolution -branch-prob -block-freq -loop-sink -lazy-branch-prob -lazy-block-fr simplifycfg -verify

-03

e-splitting lazy-block-fre ree -sroa -bas: agation -simp bine -libcal a -aa -loops fy -lcssa-ver -lazy-branch s -loop-idiom -emitter -gvn reg -opt-remark os -loop-simpl: -loops -lazy-b alopt -globald -loop-rotate volution -basi -scalar-evol -simplifycfg ectorizer -opt azy-block-freq tions -strip-de basicaa -aa -s div-rem-pairs

Compiler Passes



The Landscape of Optimizing Compilers







Experts define how to optimize the program (*algorithm*) in a separate *schedule*

The Landscape of Optimizing Compilers





Schedule



Experts define how to optimize the program (*algorithm*) in a separate *schedule*

The Landscape of Optimizing Compilers



Schedule-Based Compilers



Schedule

The Landscape of Optimizing Compilers



Schedule-Based Compilers



// functional description of matrix multiplication 2 Var x("x"), y("y"); Func prod("prod"); RDom r(o, size); 3 prod(x, y) += A(x, r) * B(r, y);out(x, y) = prod(x, y);5 6 // schedule for Nvidida GPUs const int warp_size = 32; const int vec_size = 2; const int x_tile = 3; const int y_tile = 4; 9 const int y unroll = 8; const int r unroll = 1; Var xi,yi,xio,xii,yii,xo,yo,x_pair,xiio,ty; RVar rxo,rxi; 10 11 out.bound(x, 0, size).bound(y, 0, size) 12 .tile(x, y, xi, yi, x_tile * vec_size * warp_size, 13 y_tile * y_unroll) .split(yi, ty, yi, y_unroll) 14 15 .vectorize(xi, vec size) 16 .split(xi, xio, xii, warp_size) 17 .reorder(xio, yi, xii, ty, x, y) 18 .unroll(xio).unroll(yi) .gpu_blocks(x, y).gpu_threads(ty).gpu_lanes(xii); 19 20 prod.store_in(MemoryType::Register).compute_at(out, x) 21 .split(x, xo, xi, warp_size * vec_size, RoundUp) 22 .split(y, ty, y, y_unroll) .gpu_threads(ty).unroll(xi, vec_size).gpu_lanes(xi) 23 24 .unroll(x0).unroll(y).update() 25 .split(x, xo, xi, warp_size * vec_size, RoundUp) 26 .split(y, ty, y, y_unroll) 27 .gpu threads(ty).unroll(xi, vec size).gpu lanes(xi) 28 .split(r.x, rxo, rxi, warp size) 29 .unroll(rxi, r unroll).reorder(xi, xo, y, rxi, ty, rxo) 30 .unroll(xo).unroll(y); 31 Var Bx = B.in().args()[0], By = B.in().args()[1]; 32 Var Ax = A.in().args()[0], Ay = A.in().args()[1];B.in().compute_at(prod, ty).split(Bx, xo, xi, warp size) 33 34 .gpu_lanes(xi).unroll(xo).unroll(By); 35 A.in().compute_at(prod, rxo).vectorize(Ax, vec_size) 36 .split(Ax,xo,xi,warp_size).gpu_lanes(xi).unroll(xo) 37 .split(Ay,yo,yi,y_tile).gpu_threads(yi).unroll(yo); A.in().in().compute_at(prod, rxi).vectorize(Ax, vec_size) 38 39 .split(Ax, xo, xi, warp_size).gpu_lanes(xi) 40 .unroll(xo).unroll(Ay);

Schedule for Nvidia GPUs



```
Matrix Multiplication
```

out(x, y) = prod(x, y);

.split(x, xo, xi, warp size * vec size, RoundUp)

Func prod("prod");
RDom r(0, size);

A(x, r) * B(r, y);

prod(x, y) +=

24

26



1

Schedules are *much harder to write* than algorithms



Matrix Multiplication



1

Schedules are *much harder to write* than algorithms

Schedules and algorithms are *not really separated*





1

Schedules and algorithms are not really separated

Schedules are *much harder to write* than algorithms



The schedule "language" is a fixed API and *not extensible*

```
2 Var x("x"), y("y"); Func prod("prod"); RDom r(o, size);
  prod(x, y) += A(x, r) + B(r, y);
   out(x, y) = prod(x, y);
  const int warp_size = 32; const int vec_size = 2;
  const int x_tile = 3; const int y_tile = 4;
   const int y_unroll = 8; const int r_unroll = 1;
   out.bound(x, 0, size).bound(y, 0, size)
      .tile(x, y, xi, yi, x_tile * vec_size * warp_size,
           y_tile * y_unroll)
   prod.store_in(MemoryType::Register).compute_at(out, x)
      .split(x, xo, xi, warp_size * vec_size, RoundUp)
24
      .split(x, xo, xi, warp size * vec size, RoundUp)
26
   Func prod("prod");
   RDom r(0, size);
   prod(x, y) +=
   A(x, r) * B(r, y);
   out(x, y) = prod(x, y);
```

Matrix Multiplication



Matrix Multiplication

The Landscape of Optimizing Compilers



Separable Convolution: Sobel Filter

Schedules are *much harder to write* than algorithms



Schedules and algorithms are *not really separated*



The schedule "language" is a fixed API and *not extensible*



5

Schedule primitives are not intuitive and have unclear semantics

Schedules are *not expressive enough*

1	Var x, y; Func out;	<pre>Func in = BC::repeat_edge(input);</pre>
2	out(x, y) = (
3	1.f * in (x-1,y-1)	+ 2.f * $in(x,y-1)$ + 1.f * $in(x+1,y-1)$ +
4	2.f * in (x-1,y)	+ 4.f * $in(x,y)$ + 2.f * $in(x+1,y)$ +
5	1.f * in (x-1,y+1)	+ 2.f * $in(x,y+1)$ + 1.f * $in(x+1,y+1)$
6) * (1.f/16.f);	

2D Convolution <u>Algorithm</u>

```
1 Var x,y; Func b_x,b_y,out; Func in=BC::repeat_edge(input);
2 b_y(x, y) = in(x, y-1) + 2.f * in(x, y) + in(x, y+1);
3 b_x(x, y) = b_y(x-1, y) + 2.f * b_y(x, y) + b_y(x+1, y);
4 out(x, y) = b_x(x, y) * (1.f/16.f);
```

Separated <u>Algorithm</u>

Desirable Properties of a Strategy Language

Wouldn't it be great...



if we could *look behind the curtains* of optimizing compilers and actually understand how optimizations are applied

Desirable Properties of a Strategy Language

Wouldn't it be great...



if we could *look behind the curtains* of optimizing compilers and actually understand how optimizations are applied



to have a *flexible* way of specifying optimizations for your compiler and your programming language

Desirable Properties of a Strategy Language

Wouldn't it be great...



if we could *look behind the curtains* of optimizing compilers and actually understand how optimizations are applied



to have a *flexible* way of specifying optimizations for your compiler and your programming language



to build custom optimizations in an *extensible* language while avoiding to rely on fixed scheduling APIs

Desirable Properties of a Strategy Language

Wouldn't it be great...



if we could *look behind the curtains* of optimizing compilers and actually understand how optimizations are applied



to have a *flexible* way of specifying optimizations for your compiler and your programming language



to build custom optimizations in an *extensible* language while avoiding to rely on fixed scheduling APIs



to have a *scalable* approach that competes with state-of-the-art solutions

Desirable Properties of a Strategy Language

Wouldn't it be great...

if we could *look behind the curtains* of optimizing compilers and actually understand how optimizations are applied

A strategy language should be built with the same standards as a language describing computation



to build custom optimizations in an *extensible* language while avoiding to rely on fixed scheduling APIs



to have a *scalable* approach that competes with state-of-the-art solutions

STRATEGIES

Optimizing Programs like it's 1998 2020

Visser et. al.: Building program optimizers with rewriting strategies (ICFP 1998)

What actually is a "Strategy"?

A *Strategy* encodes a program transformation:

type Strategy[P] = P => RewriteResult[P]

What actually is a "Strategy"?

A *Strategy* encodes a program transformation:

```
type Strategy[P] = P => RewriteResult[P]
```

A *RewriteResult* encodes its success or failure:

What actually is a "Strategy"?

A *Strategy* encodes a program transformation:

type Strategy[P] = P => RewriteResult[P]

A *RewriteResult* encodes its success or failure:

Two naive generic strategies:

def id[P]: Strategy[P] = (p:P) => Success(p)
def fail[P]: Strategy[P] = (p:P) => Failure(fail)

A Language-Specific Strategy

Let's encode an arithmetic simplification as a strategy: $x + 0 \rightarrow x$

A Language-Specific Strategy

Let's encode an arithmetic simplification as a strategy: $x + o \rightarrow x$

val p: ArithExpr = x + 0 // AST: Plus(Var("x"), 0)

Our toy-example target DSL

A Language-Specific Strategy

Let's encode an arithmetic simplification as a strategy: $x + 0 \rightarrow x$

val p: ArithExpr = x + 0 // AST: Plus(Var("x"), 0)



A Language-Specific Strategy

Let's encode an arithmetic simplification as a strategy: $x + 0 \rightarrow x$

val p: ArithExpr = x + 0 // AST: Plus(Var("x"), 0)



A Language-Specific Strategy

Let's encode an arithmetic simplification as a strategy: $x + 0 \rightarrow x$

val p: ArithExpr = x + 0 // AST: Plus(Var("x"), 0)



def plus0: Strategy[ArithExpr] =
 (p:ArithExpr) => p match {

Simplification rule expressed in ELEVATE:

A Language-Specific Strategy

Let's encode an arithmetic simplification as a strategy: $x + 0 \rightarrow x$

val p: ArithExpr = x + 0 // AST: Plus(Var("x"), 0)



```
def plus0: Strategy[ArithExpr] =
  (p:ArithExpr) => p match {
    case Plus(Var(x),0) => Success( Var(x) )
  }
}
```

Simplification rule expressed in ELEVATE:

A Language-Specific Strategy

Let's encode an arithmetic simplification as a strategy: $x + 0 \rightarrow x$

val p: ArithExpr = x + 0 // AST: Plus(Var("x"), 0)



```
def plus0: Strategy[ArithExpr] =
  (p:ArithExpr) => p match {
    case Plus(Var(x),0) => Success( Var(x) )
    case _ => Failure( plus0 )
  }
```

Simplification rule expressed in ELEVATE:

A More Interesting Language-Specific Strategy

Let's encode another language specific strategy: $map(f) \circ map(g) \rightarrow map(f \circ g)$

val p: Lift = fun(xs => map(f)(map(g)(xs)))
A More Interesting Language-Specific Strategy

Let's encode another language specific strategy: $map(f) \circ map(g) \rightarrow map(f \circ g)$

```
val p: Lift = fun(xs => map(f)(map(g)(xs)))
```



A More Interesting Language-Specific Strategy

Let's encode another language specific strategy: $map(f) \circ map(g) \rightarrow map(f \circ g)$

val p: Lift = fun(xs => map(f)(map(g)(xs)))



 $map(f)(map(g)(xs)) \rightarrow map(fun(x \Rightarrow f(g(x))))(xs)$

A More Interesting Language-Specific Strategy

Let's encode another language specific strategy: $map(f) \circ map(g) \rightarrow map(f \circ g)$

```
val p: Lift = fun(xs \Rightarrow map(f)(map(g)(xs)))
```



A More Interesting Language-Specific Strategy

Let's encode another language specific strategy: $map(f) \circ map(g) \rightarrow map(f \circ g)$

val p: Lift = $fun(xs \Rightarrow map(f)(map(g)(xs)))$



A More Interesting Language-Specific Strategy

Let's encode another language specific strategy: $map(f) \circ map(g) \rightarrow map(f \circ g)$

```
val p: Lift = fun(xs \Rightarrow map(f)(map(g)(xs)))
```



A More Interesting Language-Specific Strategy

Let's encode another language specific strategy: $map(f) \circ map(g) \rightarrow map(f \circ g)$

val p: Lift = $fun(xs \Rightarrow map(f)(map(g)(xs)))$



Program Transformations as Strategies

Essentially: myTransformation: $lhs \rightarrow rhs$

ELEVATE

```
def myTransformation: Strategy[MyLanguage] =
  (p:MyLanguage) => p match {
    case lhs => Success( rhs )
    case _ => Failure( myTransformation )
  }
```

How to Build More Powerful Strategies

The seq combinator applies two strategies in sequence



How to Build More Powerful Strategies

The seq combinator applies two strategies in sequence

```
def seq[P]: Strategy[P] => Strategy[P] => Strategy[P] =
      fs => ss => p => fs(p) >>= ss
```

The IChoice combinator applies the second Strategy only if the first one failed



How to Build More Powerful Strategies

The seq combinator applies two strategies in sequence

```
def seq[P]: Strategy[P] => Strategy[P] => Strategy[P] =
      fs => ss => p => fs(p) >>= ss
```

The IChoice combinator applies the second Strategy only if the first one failed



How to Build More Powerful Strategies

```
def seq[P]: Strategy[P] => Strategy[P] => Strategy[P] =
        fs => ss => p => fs(p) >>= ss
```

```
def lChoice[P]: Strategy[P] => Strategy[P] => Strategy[P] =
      fs => ss => p => fs(p) <|> ss(p)
```

The try combinator tries to apply a strategy and in case of Failure returns the input unchanged

```
def try[P]: Strategy[P] => Strategy[P] =
    s => p => (s <+ id)(p)</pre>
```

Observation: *try* never fails!

How to Build More Powerful Strategies

```
def seq[P]: Strategy[P] => Strategy[P] => Strategy[P] =
      fs => ss => p => fs(p) >>= ss
```

```
def lChoice[P]: Strategy[P] => Strategy[P] => Strategy[P] =
    fs => ss => p => fs(p) <|> ss(p)
```

The try combinator tries to apply a strategy and in case of Failure returns the input unchanged

```
def try[P]: Strategy[P] => Strategy[P] =
    s => p => (s <+ id)(p)</pre>
```

The *repeat* combinator applies a strategy until it's no longer applicable

def repeat[P]: Strategy[P] => Strategy[P] =
 s => p => try(s ; repeat(s))(p)

Describing Precise Locations in the AST

Another simple Lift program: (*map(f)* • *map(g)* • *map(h)*)(xs)

val threeMaps: Lift = fun(xs => map(f)(map(g)(map(h)(xs))))

Describing Precise Locations in the AST

Another simple Lift program: (*map(f)* • *map(g)* • *map(h)*)(xs)

val threeMaps: Lift = fun(xs => map(f)(map(g)(map(h)(xs))))



Describing Precise Locations in the AST

Another simple Lift program: (*map(f)* • *map(g)* • *map(h)*)(xs)

val threeMaps: Lift = fun(xs => map(f)(map(g)(map(h)(xs))))



```
def mapFusion: Strategy[Lift] =
  (p:Lift) => p match {
    case app(app(map, f),
        app(app(map, g), xs)) =>
        Success( map(fun(x => f(g(x))))(xs) )
        case _ => Failure( mapFusion )
    }
```

mapFusion: $map(f) \circ map(g) \rightarrow map(f \circ g)$

Describing Precise Locations in the AST

Another simple Lift program: (*map(f)* • *map(g)* • *map(h)*)(xs)

val threeMaps: Lift = fun(xs => map(f)(map(g)(map(h)(xs))))



```
def mapFusion: Strategy[Lift] =
  (p:Lift) => p match {
    case app(app(map, f),
        app(app(map, g), xs)) =>
        Success( map(fun(x => f(g(x))))(xs) )
        case _ => Failure( mapFusion )
    }
```

mapFusion: $map(f) \circ map(g) \rightarrow map(f \circ g)$

Describing Precise Locations in the AST

Another simple Lift program: (*map(f)* • *map(g)* • *map(h)*)(xs)

val threeMaps: Lift = fun(xs => map(f)(map(g)(map(h)(xs))))



def mapFusion: Strategy[Lift] =
 (p:Lift) => p match {
 case app(app(map, f),
 app(app(map, g), xs)) =>
 Success(map(fun(x => f(g(x))))(xs))
 case _ => Failure(mapFusion)
 }

mapFusion: $map(f) \circ map(g) \rightarrow map(f \circ g)$

Describing Precise Locations in the AST

Another simple Lift program: (*map(f) • map(g) • map(h))(xs)*

val threeMaps: Lift = fun(xs => map(f)(map(g)(map(h)(xs))))



def mapFusion: Strategy[Lift] =
 (p:Lift) => p match {
 case app(app(map, f),
 app(app(map, g), xs)) =>
 Success(map(fun(x => f(g(x))))(xs))
 case _ => Failure(mapFusion)
 }

mapFusion: $map(f) \circ map(g) \rightarrow map(f \circ g)$

mapFusion(threeMaps) == Failure(mapFusion)

Describing Precise Locations in the AST

Another simple Lift program: (*map(f) • map(g) • map(h))(xs)*

val threeMaps: Lift = fun(xs => map(f)(map(g)(map(h)(xs))))



def mapFusion: Strategy[Lift] =
 (p:Lift) => p match {
 case app(app(map, f),
 app(app(map, g), xs)) =>
 Success(map(fun(x => f(g(x))))(xs))
 case _ => Failure(mapFusion)
 }

mapFusion: $map(f) \circ map(g) \rightarrow map(f \circ g)$

mapFusion(threeMaps) == Failure(mapFusion)

Describing Precise Locations in the AST



A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)



Describing Precise Locations in the AST

A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

```
def all[P] : Strategy[P] => Strategy[P]
def one[P] : Strategy[P] => Strategy[P]
def some[P]: Strategy[P] => Strategy[P]
```



Describing Precise Locations in the AST

A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

<pre>def all[P] :</pre>	Strategy [P]	=>	<pre>Strategy[P]</pre>
<pre>def one[P] :</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
<pre>def some[P]:</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>

Let's try...

all(mapFusion)(threeMaps)



Describing Precise Locations in the AST

A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

<pre>def all[P] :</pre>	Strategy [P]	=>	Strategy [P]
<pre>def one[P] :</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
<pre>def some[P]:</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>

Let's try...

all(mapFusion)(threeMaps)

all fails if the strategy is not applicable to all children



Describing Precise Locations in the AST

A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

<pre>def all[P] :</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
<pre>def one[P] :</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
<pre>def some[P]:</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>

Let's try...

all(mapFusion)(threeMaps)

all fails if the strategy is not applicable to all children



all(mapFusion)

un

Describing Precise Locations in the AST

A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

<pre>def all[P] :</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
<pre>def one[P] :</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
<pre>def some[P]:</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>

Let's try...

one(mapFusion)(threeMaps)

one fails if the strategy is not applicable to any child



Describing Precise Locations in the AST

A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

<pre>def all[P] :</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
<pre>def one[P] :</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
<pre>def some[P]:</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>

Let's try...

one(mapFusion)(threeMaps)

one fails if the strategy is not applicable to any child



Describing Precise Locations in the AST



A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

<pre>def all[P] :</pre>	Strategy [P]	=>	<pre>Strategy[P]</pre>
<pre>def one[P] :</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
<pre>def some[P]:</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>

Let's try...

one(mapFusion)(threeMaps)

one fails if the strategy is not applicable to any child

Describing Precise Locations in the AST

A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

<pre>def all[P] :</pre>	Strategy [P]	=>	Strategy [P]
<pre>def one[P] :</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
<pre>def some[P]:</pre>	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>

Let's try...

one(mapFusion)(threeMaps)

one fails if the strategy is not applicable to any child



Describing Precise Locations in the AST

A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

def	<pre>all[P]</pre>	•	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
def	one[P]	•	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
def	<pre>some[P]</pre>	•	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>

or we define our own *domain-specific traversals*:

```
def body: Strategy[Lift] => Strategy[Lift] =
   s => p => p match {
```



Describing Precise Locations in the AST

A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

def	<pre>all[P]</pre>	•	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
def	one[P]	•	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>
def	<pre>some[P]</pre>	•	<pre>Strategy[P]</pre>	=>	<pre>Strategy[P]</pre>

or we define our own *domain-specific traversals*:

```
def body: Strategy[Lift] => Strategy[Lift] =
    s => p => p match {
    case fun(x,b) =>
        s(b).mapSuccess( nb => fun(x,nb) )
    case _ => Failure( body ) }
```



map

 $fun(xs \Rightarrow map(f)(map(g)(map(h)(xs))))$

un

map

Describing Precise Locations in the AST

mapFusion app id app app app map app app map XS map

body(mapFusion)

 $fun(xs \Rightarrow map(f)(map(g)(map(h)(xs))))$

A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

```
def all[P] : Strategy[P] => Strategy[P]
def one[P] : Strategy[P] => Strategy[P]
def some[P]: Strategy[P] => Strategy[P]
```

or we define our own *domain-specific traversals*:

```
def body: Strategy[Lift] => Strategy[Lift] =
    s => p => p match {
    case fun(x,b) =>
        s(b).mapSuccess( nb => fun(x,nb) )
    case _ => Failure( body ) }
```

Describing Precise Locations in the AST



 $fun(xs \Rightarrow map(f)(map(g)(map(h)(xs))))$

A strategy is generally always applied at the *root* of the AST

mapFusion(threeMaps)

...but we can use *generic one-level traversals* to push strategy applications down the AST

```
def all[P] : Strategy[P] => Strategy[P]
def one[P] : Strategy[P] => Strategy[P]
def some[P]: Strategy[P] => Strategy[P]
```

or we define our own *domain-specific traversals*:

```
def arg: Strategy[Lift] => Strategy[Lift] =
    s => p => p match {
      case app(f,e) =>
        s(e).mapSuccess( ne => app(f,ne) )
      case _ => Failure( arg ) }
```

COMPLETE TRAVERSALS

Go Down More Than One Step

The *topDown* traversal traverses the tree until it finds a successful location



COMPLETE TRAVERSALS

Go Down More Than One Step

The topDown traversal traverses the tree until it finds a successful location

def topDown[P]: Strategy[P] => Strategy[P] =
 s => p => (s <+ one(topDown(s)))(p)</pre>

def bottomUp[P]: Strategy[P] => Strategy[P] =
 s => p => (one(bottomUp(s)) <+ s)(p)</pre>

def allTopDown[P]: Strategy[P] => Strategy[P] =
 s => p => (s ; one(allTopDown(s)))(p)

COMPLETE TRAVERSALS

Go Down More Than One Step

The topDown traversal traverses the tree until it finds a successful location

def topDown[P]: Strategy[P] => Strategy[P] =
 s => p => (s <+ one(topDown(s)))(p)</pre>

def bottomUp[P]: Strategy[P] => Strategy[P] =
 s => p => (one(bottomUp(s)) <+ s)(p)</pre>

def allTopDown[P]: Strategy[P] => Strategy[P] =
 s => p => (s ; one(allTopDown(s)))(p)

or we could also *normalize* an AST

def normalize[P]: Strategy[P] => Strategy[P] =
 s => p => (repeat(topDown(s))(p)

RECAP

What have we seen so far?

With ELEVATE we are able to...



to define *language-specific transformations* as strategies



to compose strategies using generic strategy combinators



to describe precise locations in the AST using *generic* and *language-specific one-step traversals*



to compose one-step traversals to define *whole-tree traversals* including *normalization*
CASE STUDIES

Put it into Practice

Optimizing F-Smooth using Elevate

ICFP'19: Efficient Differentiable Programming in a Functional Array-Processing Language

Efficient Differentiable Programming in a Functional Array-Processing Language

AMIR SHAIKHHA, University of Oxford, United Kingdom ANDREW FITZGIBBON, Microsoft Research, United Kingdom DIMITRIOS VYTINIOTIS, DeepMind, United Kingdom SIMON PEYTON IONES, Microsoft Research, United Kingdom

We present a system for the automatic differentiation (AD) of a higher-order functional array-processing language. The core functional language underlying this system simultaneously supports both source-loss source forexan-broads AD and global optimizations such as a loop transformations. In combinisting, gradient computition with forward-roads AD can be as efficient a serverse mode, and that the Jacobian natives retrongeneous section of the section of the section of the section of the efficient of comp

 $\label{eq:CCS} Concepts:= Mathematics of computing \rightarrow Automatic differentiation:= Software and its encering \rightarrow Functional languages: Donais specific languages.$

Additional Key Words and Phrases: Linear Algebra, Differentiable Programming, Optimising i Fusion, Code Motion. ACM Reference Format: Amir Shakiba, Andrew Fizgibbon, Dimitriss Vytiniotis, and Simen Peyton Jones. 2019. Efficie

Juni makani, Antrew magnosi, Junima vyinioni, ani sinon reytori jons. 2019. Encont i innermatore Programming in a Functional Array-Processing Language. Proc. ACM Programs. Lang. 3, ICEP, Article 97 (August 2019), 30 pages. https://doi.org/10.1143/3341701

... in the summer of 1958 John McCarthy decided to investigate differentiation as an interesting symbolic computation problem, which was difficult to express in the primitive programming languages of the day. This investigation led hint to ree the imperatore of functional arguments and recurrive functions in the field of symbolic computation. From Norvig [Norvig: 1992, pp. 248].

INTRODUCTION

For easi - doub. Automatic Differentiation is relatively straightforward, both as a minimum behaping and implement as a source-to-source program transformation. However, forward-mode AD is unally considered widdly infinition as a way to compare the gradient of a function, because the moder of any first forward mode AD is unally been explored and a AD formation times — and in any hevery large (e.g. $a = 10^{\circ}$). Transformation, reverse mode AD is characterised by the receasity to maintain temporary unalless transformation, reverse-mode AD is characterised by the necessity to maintain temporary unalless. Modern

Anthone' addressee: Anzie Shalddha, Uziversity of Oxford, United Klagdom, amizshaldzhaiglecoreac alc' Andrew Titrgfebor, Microsoft Besoneth, United Klagdom, awdfmicrosoft Genes Dimitrico Vytiniotis, DoepMind, United Klagdom, dvytiad gogale com, Strom Dysten, Jones, Missenell Besearch, Mindel Klagdom, simorpgi@macrosoft.com.

This work is licensed under a Creative Comment Attribution 4.9 International License. 2015 Copyright Solid by the overactuation(s). 2015-1012/014-04079 Mayerscheerg 14:1512301311 Arbitrary F-Smooth expressions are differentiable

They achieve efficiency by *rewriting* differentiated code

The strategy for applying rewrite rules can become tricky

Use ELEVATE for optimizing F-Smooth programs

Optimizing F-Smooth using Elevate



length (build $e_0 e_1$) $\rightarrow e_0$

Fig. 8. Transformation Rules for \widetilde{F} . Even though none of these rules are AD-specific, the rules of Figure 8f and Figure 8g are more useful in the AD context.

F-Smooth Rewrite Rules

Optimizing F-Smooth using Elevate



Fig. 8. Transformation Rules for \widetilde{F} . Even though none of these rules are AD-specific, the rules of Figure 8f and Figure 8g are more useful in the AD context.

F-Smooth Rewrite Rules

Optimizing F-Smooth using Elevate



Fig. 8. Transformation Rules for \widetilde{F} . Even though none of these rules are AD-specific, the rules of Figure 8f and Figure 8g are more useful in the AD context. We are able to *trace* the rule applications: Here, 12 steps

F-Smooth Rewrite Rules





We are able to **trace** the rule applications: Here, 12 steps

Fig. 8. Transformation Rules for \widetilde{F} . Even though none of these rules are AD-specific, the rules of Figure 8f and Figure 8g are more useful in the AD context.



Expressing Separable Convolution with Elevate and Lift

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \times \begin{bmatrix} -1 & 0 & 1 \end{bmatrix}$$

Separable Convolution: Sobel Filter

```
1 Var x, y; Func out; Func in = BC::repeat_edge(input);
2 out(x, y) = (
3    1.f * in(x-1,y-1) + 2.f * in(x,y-1) + 1.f * in(x+1,y-1) +
4    2.f * in(x-1,y) + 4.f * in(x,y) + 2.f * in(x+1,y) +
5    1.f * in(x-1,y+1) + 2.f * in(x,y+1) + 1.f * in(x+1,y+1)
6  ) * (1.f/16.f);
```

Halide: 2D Convolution

```
no schedule for this optimization
```

```
1 Var x,y; Func b_x,b_y,out; Func in=BC::repeat_edge(input);
2 b_y(x, y) = in(x, y-1) + 2.f * in(x, y) + in(x, y+1);
3 b_x(x, y) = b_y(x-1, y) + 2.f * b_y(x, y) + b_y(x+1, y);
4 out(x, y) = b_x(x, y) * (1.f/16.f);
```

Halide: Separated Convolution

Expressing Separable Convolution with Elevate and Lift

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \times \begin{bmatrix} -1 & 0 & 1 \end{bmatrix}$$

Separable Convolution: Sobel Filter

```
1 Var x, y; Func out; Func in = BC::repeat_edge(input);
2 out(x, y) = (
3   1.f * in(x-1,y-1) + 2.f * in(x,y-1) + 1.f * in(x+1,y-1) +
4   2.f * in(x-1,y) + 4.f * in(x,y) + 2.f * in(x+1,y) +
5   1.f * in(x-1,y+1) + 2.f * in(x,y+1) + 1.f * in(x+1,y+1)
6 ) * (1.f/16.f);
```

Halide: 2D Convolution



no schedule for this optimization

1 Var x,y; Func b_x,b_y,out; Func in=BC::repeat_edge(input); 2 b_y(x, y) = in(x, y-1) + 2.f * in(x, y) + in(x, y+1); 3 b_x(x, y) = b_y(x-1, y) + 2.f * b_y(x, y) + b_y(x+1, y); 4 out(x, y) = b x(x, y) * (1.f/16.f);

Halide: Separated Convolution

img |>
 pad2D(1) |>
 slide2D(3)(1) |>
 map2D(fun(nbh => nbh |> // 2 1D stencils
 map(dot(weightsH)) |> map(dot(weightsV))))

Lift: Separated Convolution

Expressing Separable Convolution with Elevate and Lift



Expressing Separable Convolution with Elevate and Lift



map(dot(weightsH)) |> map(dot(weightsV))))

Lift: Separated Convolution

Expressing Separable Convolution with Elevate and Lift



Lift: Separated Convolution

Implementing a Scheduling Language using Strategies

Stvm *Tutorial*: How to optimize GEMM



In this tutorial, we will demonstrate how to use TVM to optimize square matrix multiplication and achieve **200 times faster** than baseline by simply adding **18 extra lines of code**.



Implementing a Scheduling Language using Strategies



TVM: Matrix Multiplication





```
val dot = fun((a,b) => zip(a,b) |> map(*) |> reduce(+,0))
val mm = fun(a :: M.K.float => fun(b :: K.N.float =>
  map( fun(arow => // iterating over M
      map( fun(bcol => // iterating over N
      dot(arow, bcol) // iterating over K
     )(transpose(b))
  )(a)
```



Implementing a Scheduling Language using Strategies



and therefore customizable



Implementing a Scheduling Language using Strategies



<pre>val blocking = (topDown(tile(32,32));</pre>	
<pre>topDown(isReduce ; split(4)) ;</pre>	
topDown (<i>reorder</i> (Seq(1,2,5,6,3,4))))
<pre>(blocking ; lowerToC)(mm)</pre>	

ELEVATE: blocking strategy

Implementing a Scheduling Language using Strategies



Our strategies achieve the **same trend** in performance \rightarrow they encode the **same optimizations** as described by the schedules





Rewriting the input program requires less than 60 seconds per version

Implementing a Scheduling Language using Strategies



Rewriting the baseline version requires less than 60 seconds per version

THANK YOU

for your attention!

b.hagedorn@wwu.de ELEVATE is OpenSource: github.com/elevate-lang