HIGH PERFORMANCE STENCIL CODE GENERATION WITH LIFT

Bastian Hagedorn | Larisa Stoltzfus | Michel Steuwer | Sergei Gorlatch | Christophe Dubach

WWU MÜNSTER

THE UNIVERSITY of EDINBURGH

University of Glasgow

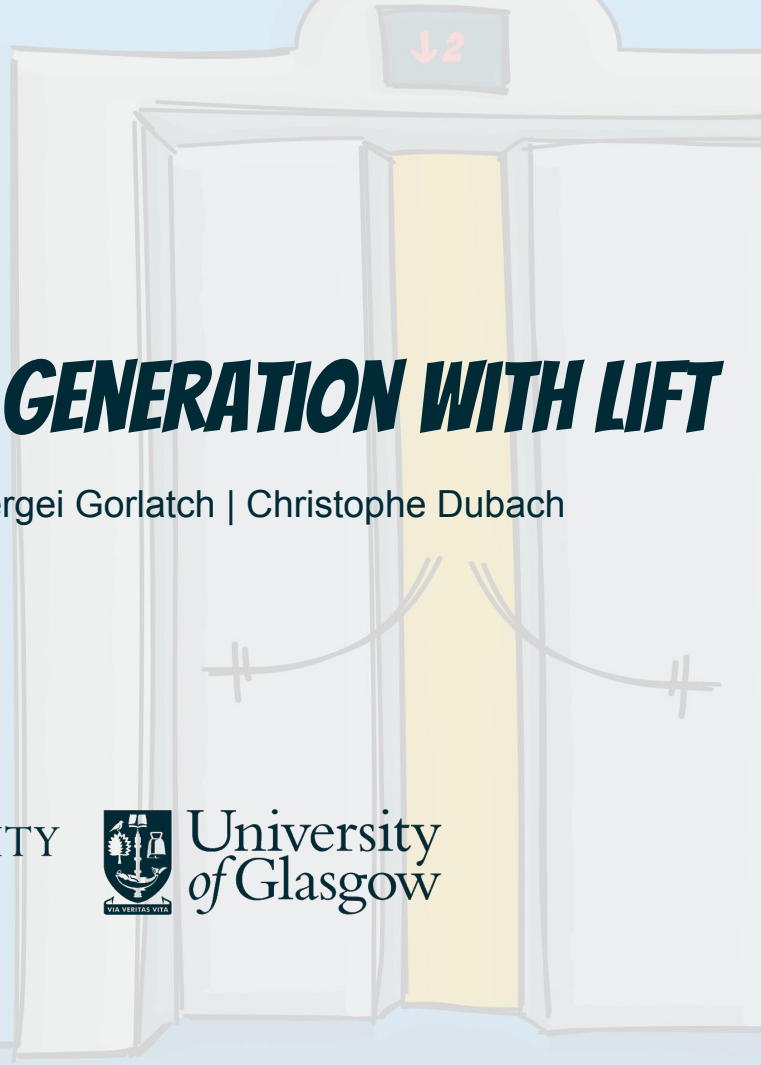

WHY STENCIL COMPUTATIONS?

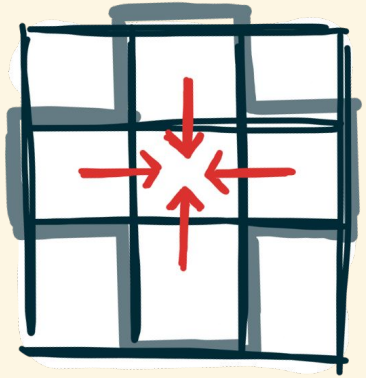

Stencil computations are a class of kernels which update *neighboring* array elements according to a fixed pattern, called *stencil*.

Frequently occur in:

- Medical Imaging
- Physics Simulations
- Machine Learning
- PDE Solvers

WHY STENCIL COMPUTATIONS?

Stencil compute time:

HPC Center München

2017

HPC Center Stuttgart

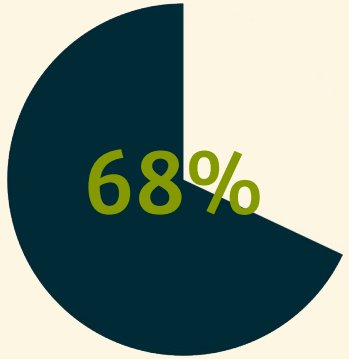

2016

Frequently occur in:

- Medical Imaging
- Physics Simulations
- Machine Learning
- PDE Solvers

YET ANOTHER STENCIL PAPER?

2005 2007 2009 2011 2013 2015 2018

ICS'05 ICS'09 CGO'12 CGO'15 CLUSTER'17

PLDI'07 SC'10 CLUSTER'13 WOLFHPC'16 CGO'18

DOMAIN SPECIFIC LANGUAGES

PATUS | Pochoir | PARTANS | Halide | ... | *DSL*

EXPLOITING DOMAIN KNOWLEDGE

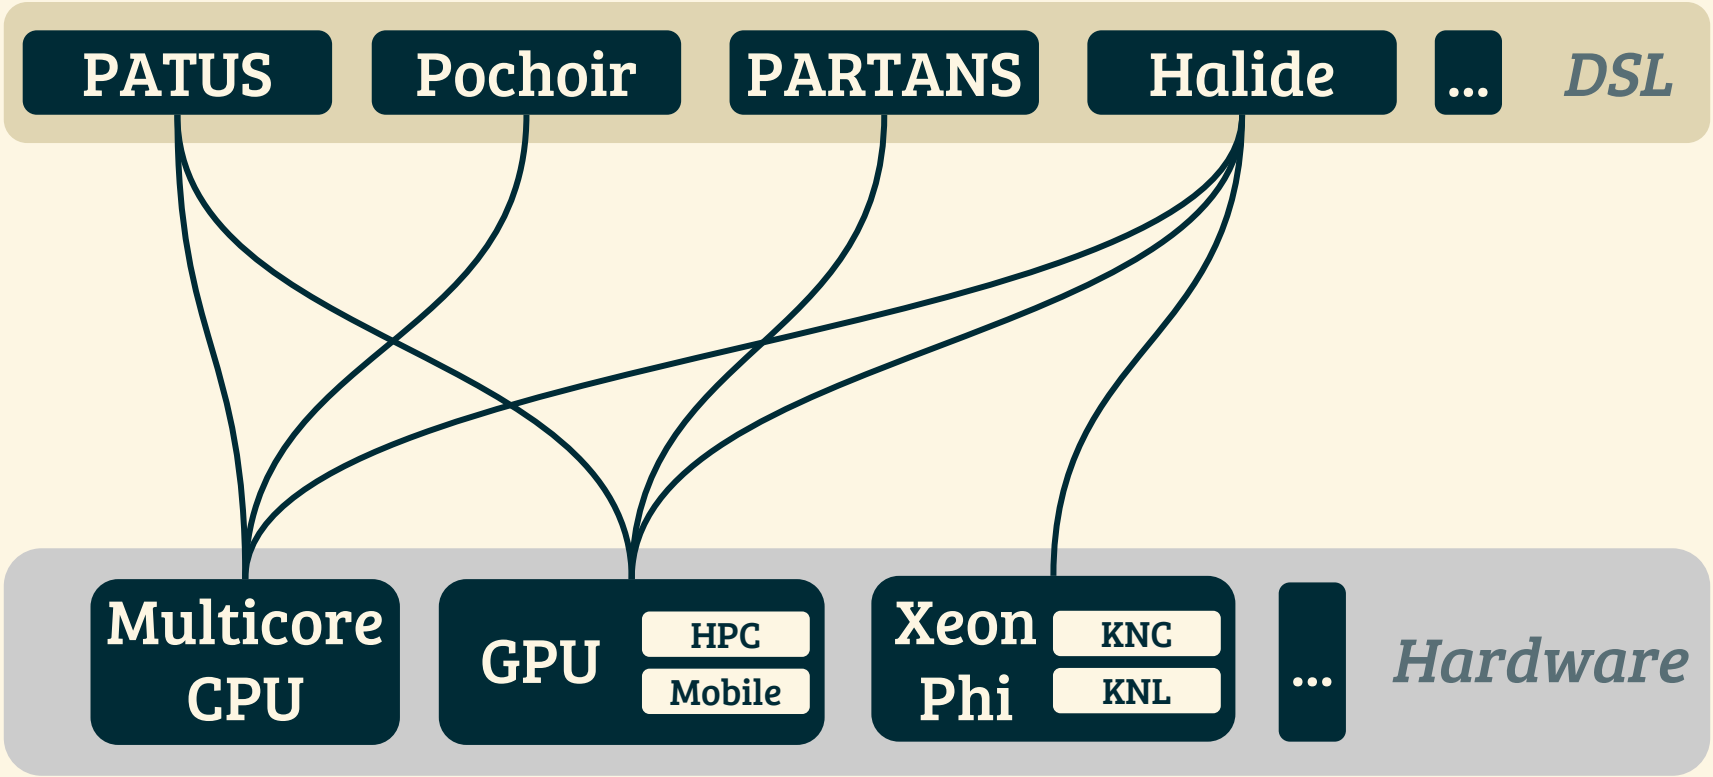

EXPLOITING DOMAIN KNOWLEDGE

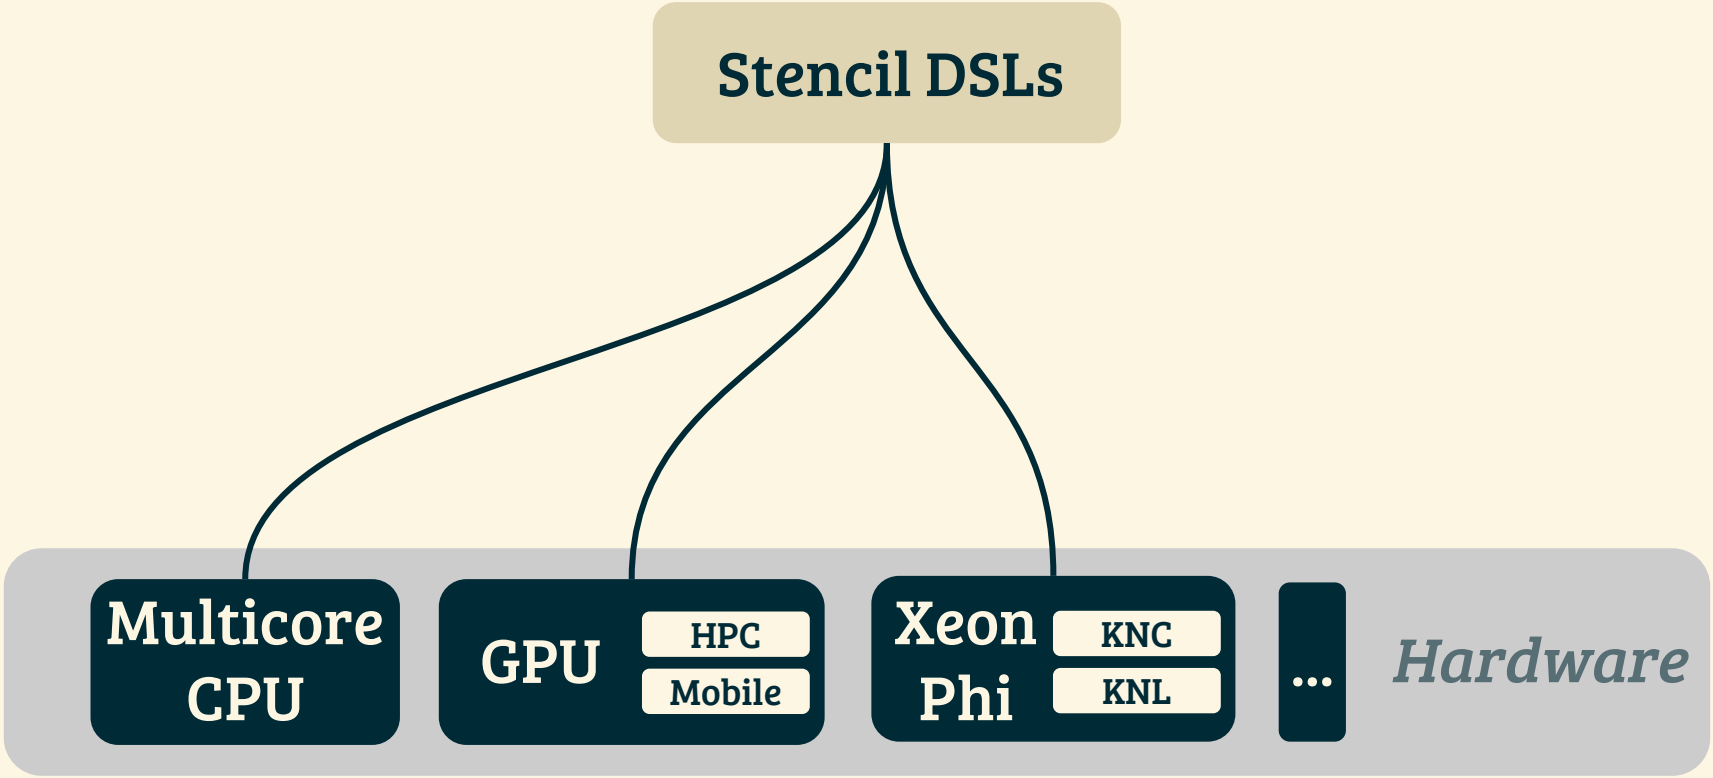

EXPLOITING DOMAIN KNOWLEDGE

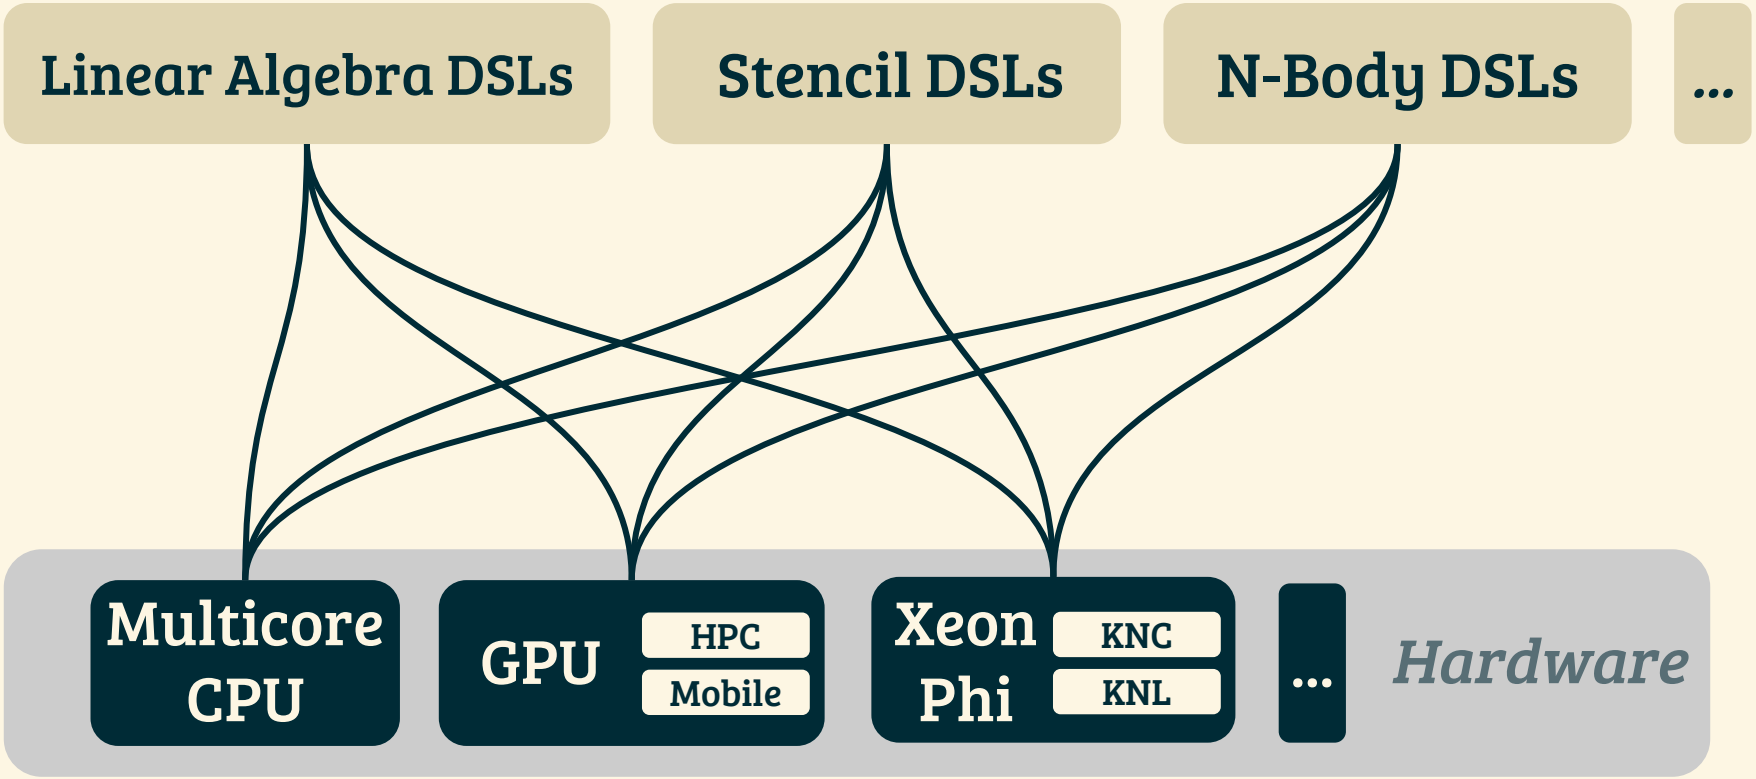

APPROACHING PERFORMANCE PORTABILITY

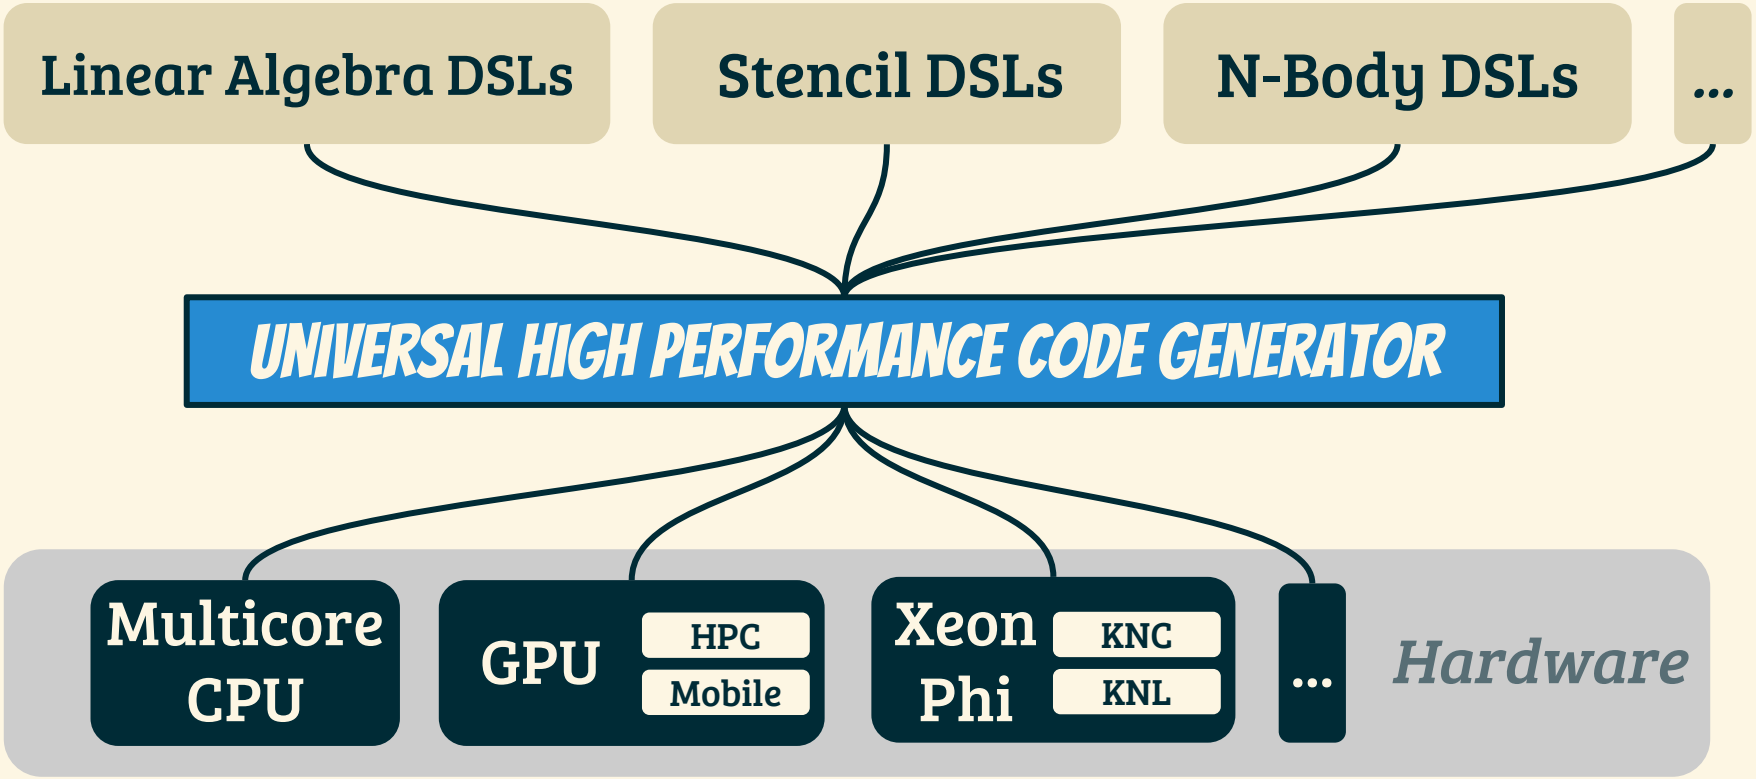

Approaching Performance Portability

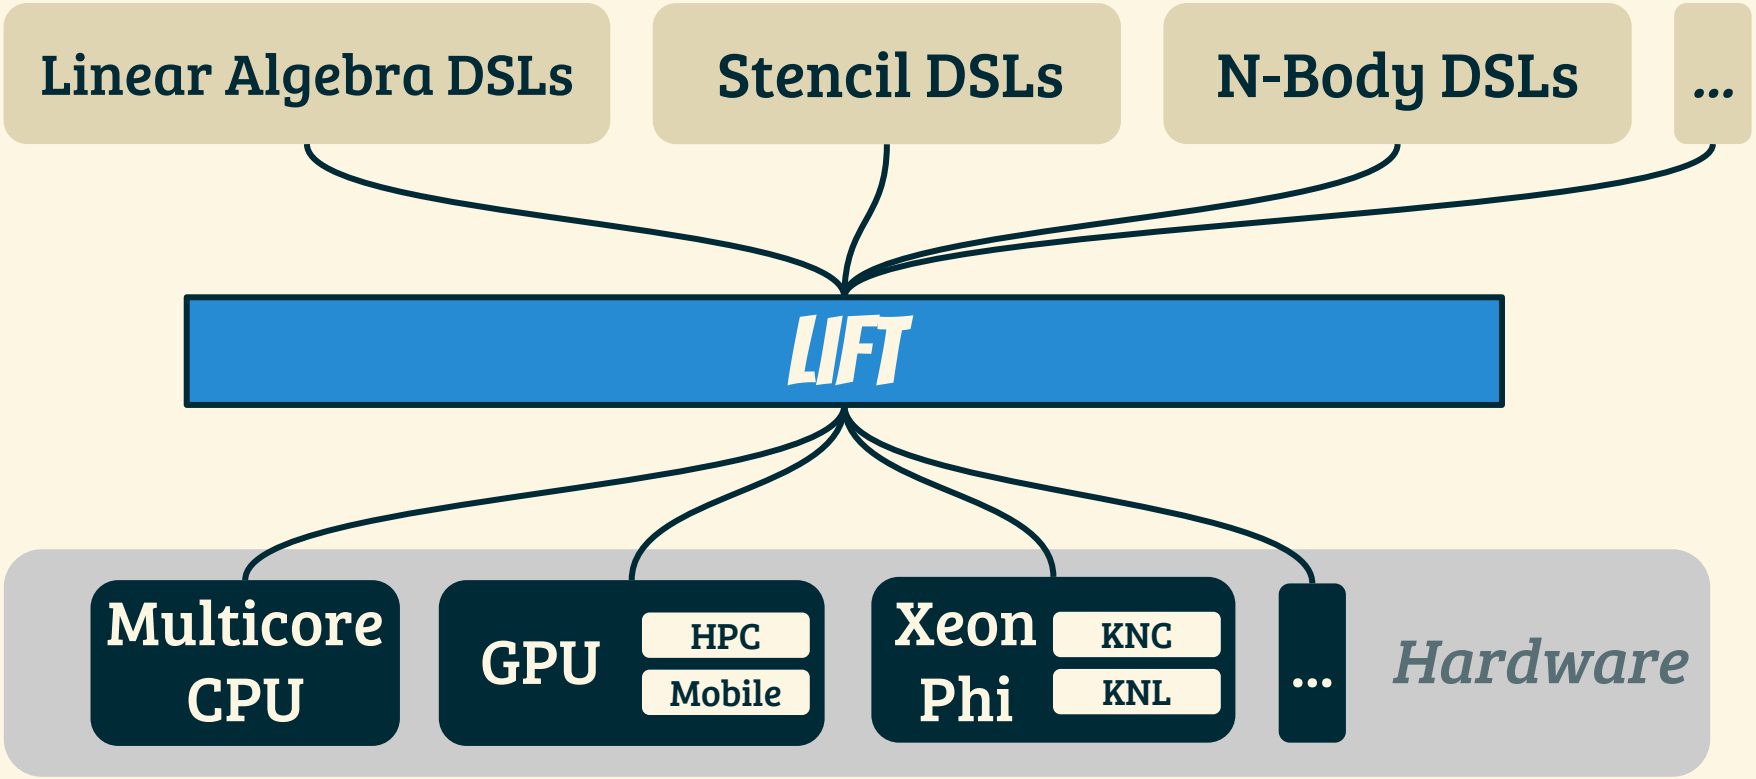

DSL DSL DSL

LIFT

Multicore CPU

GPU HPC Mobile

Xeon Phi KNC KNL

...

Hardware

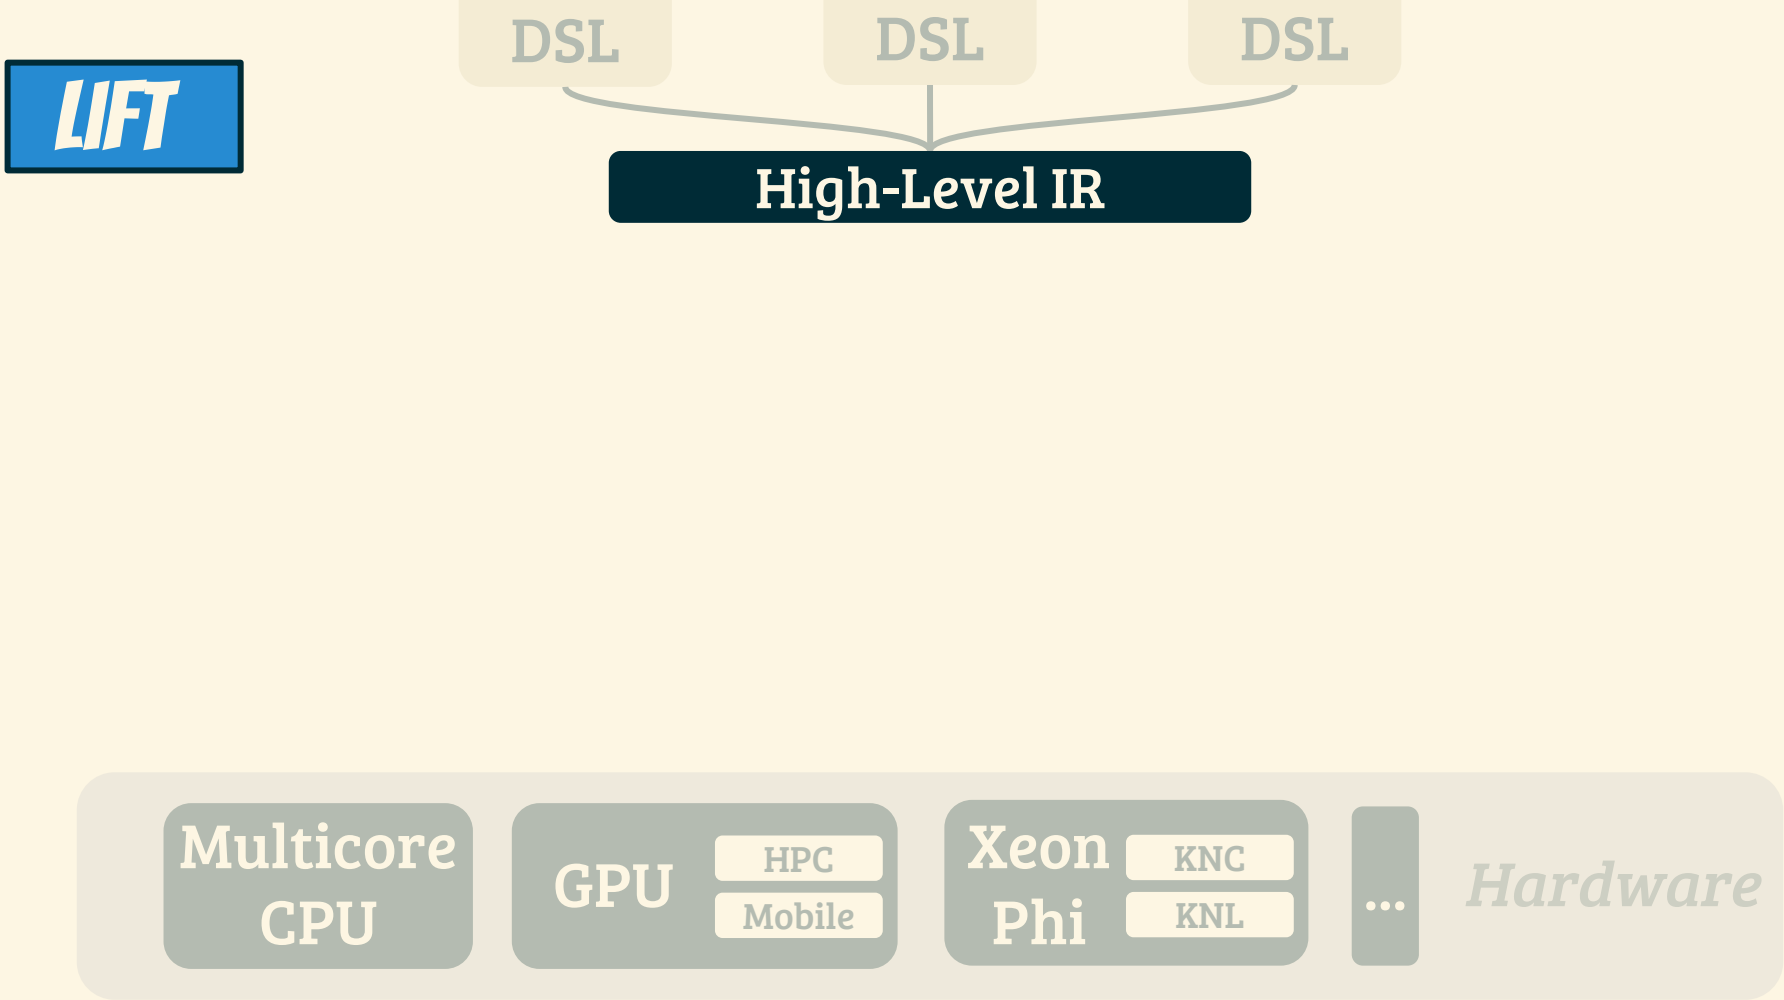
LIFT
DSL
DSL
DSL
High-Level IR
Multicore CPU
GPU
HPC
Mobile
Xeon Phi
KNC
KNL
...
Hardware

LIFT

DSL DSL DSL

High-Level IR

Explore Optimizations
by rewriting

Multicore CPU

GPU HPC Mobile

Xeon Phi KNC KNL

...

Hardware

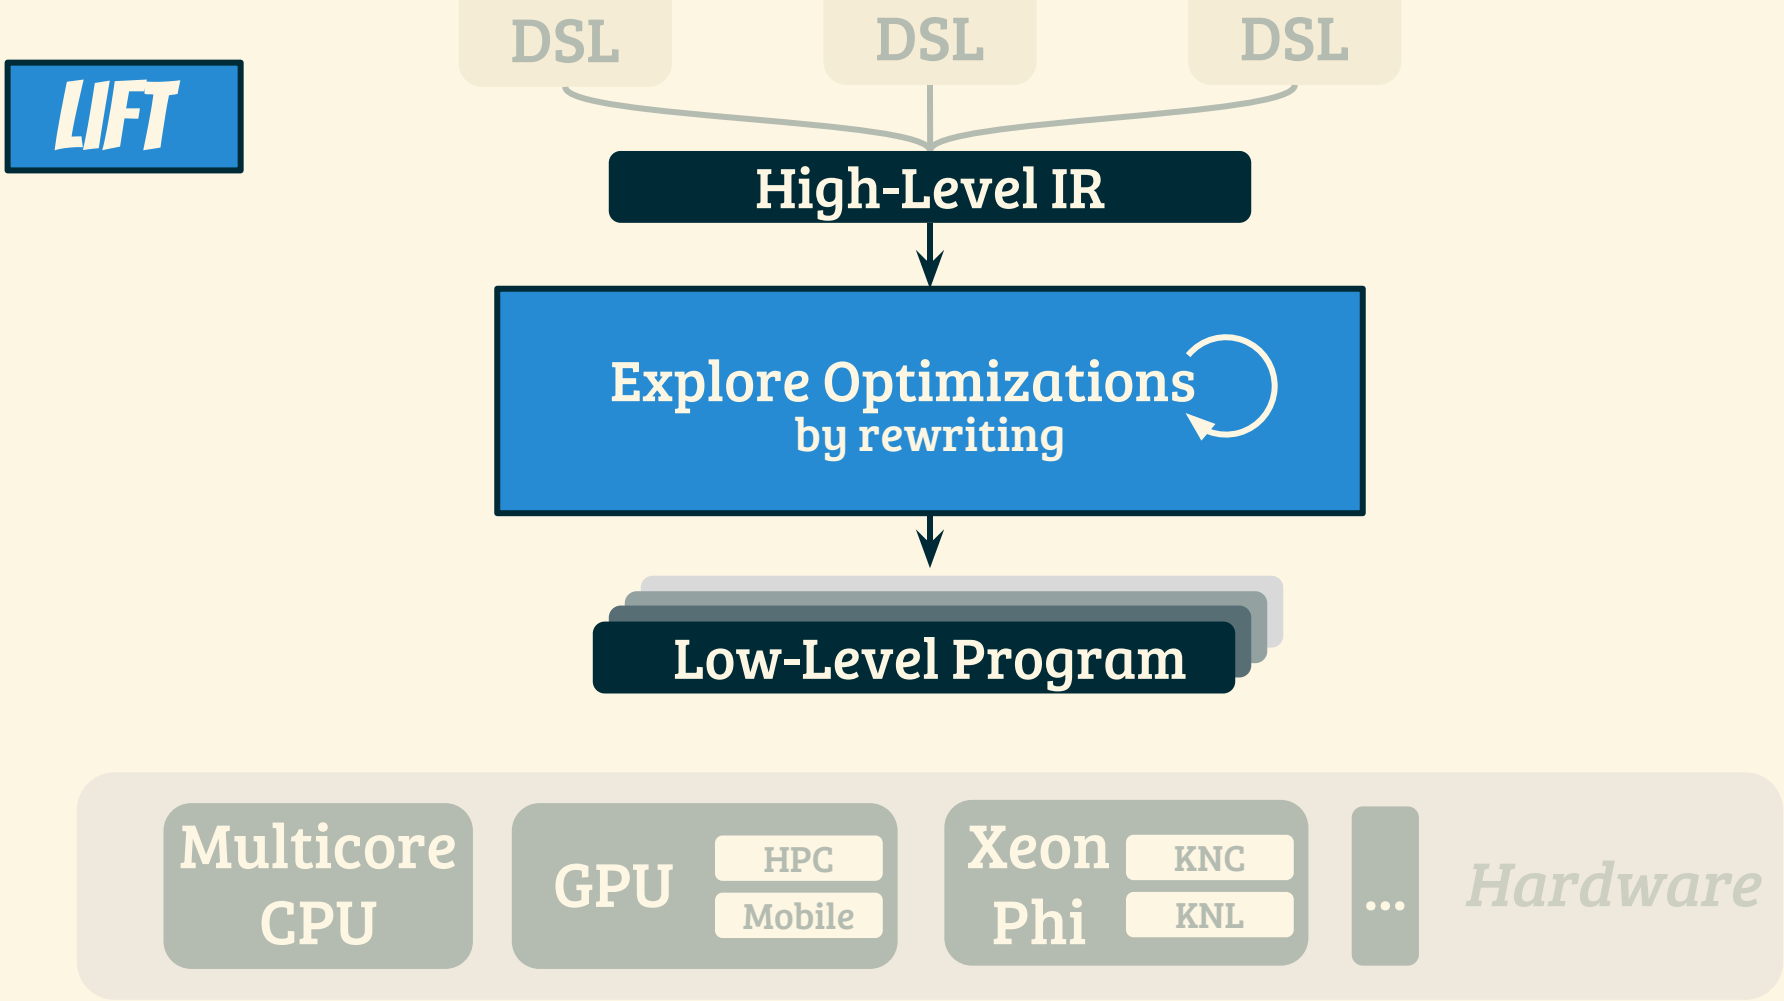
LIFT
DSL
DSL
DSL
High-Level IR
Explore Optimizations
by rewriting
Low-Level Program
Multicore CPU
GPU
HPC
Mobile
Xeon Phi
KNC
KNL
...
Hardware

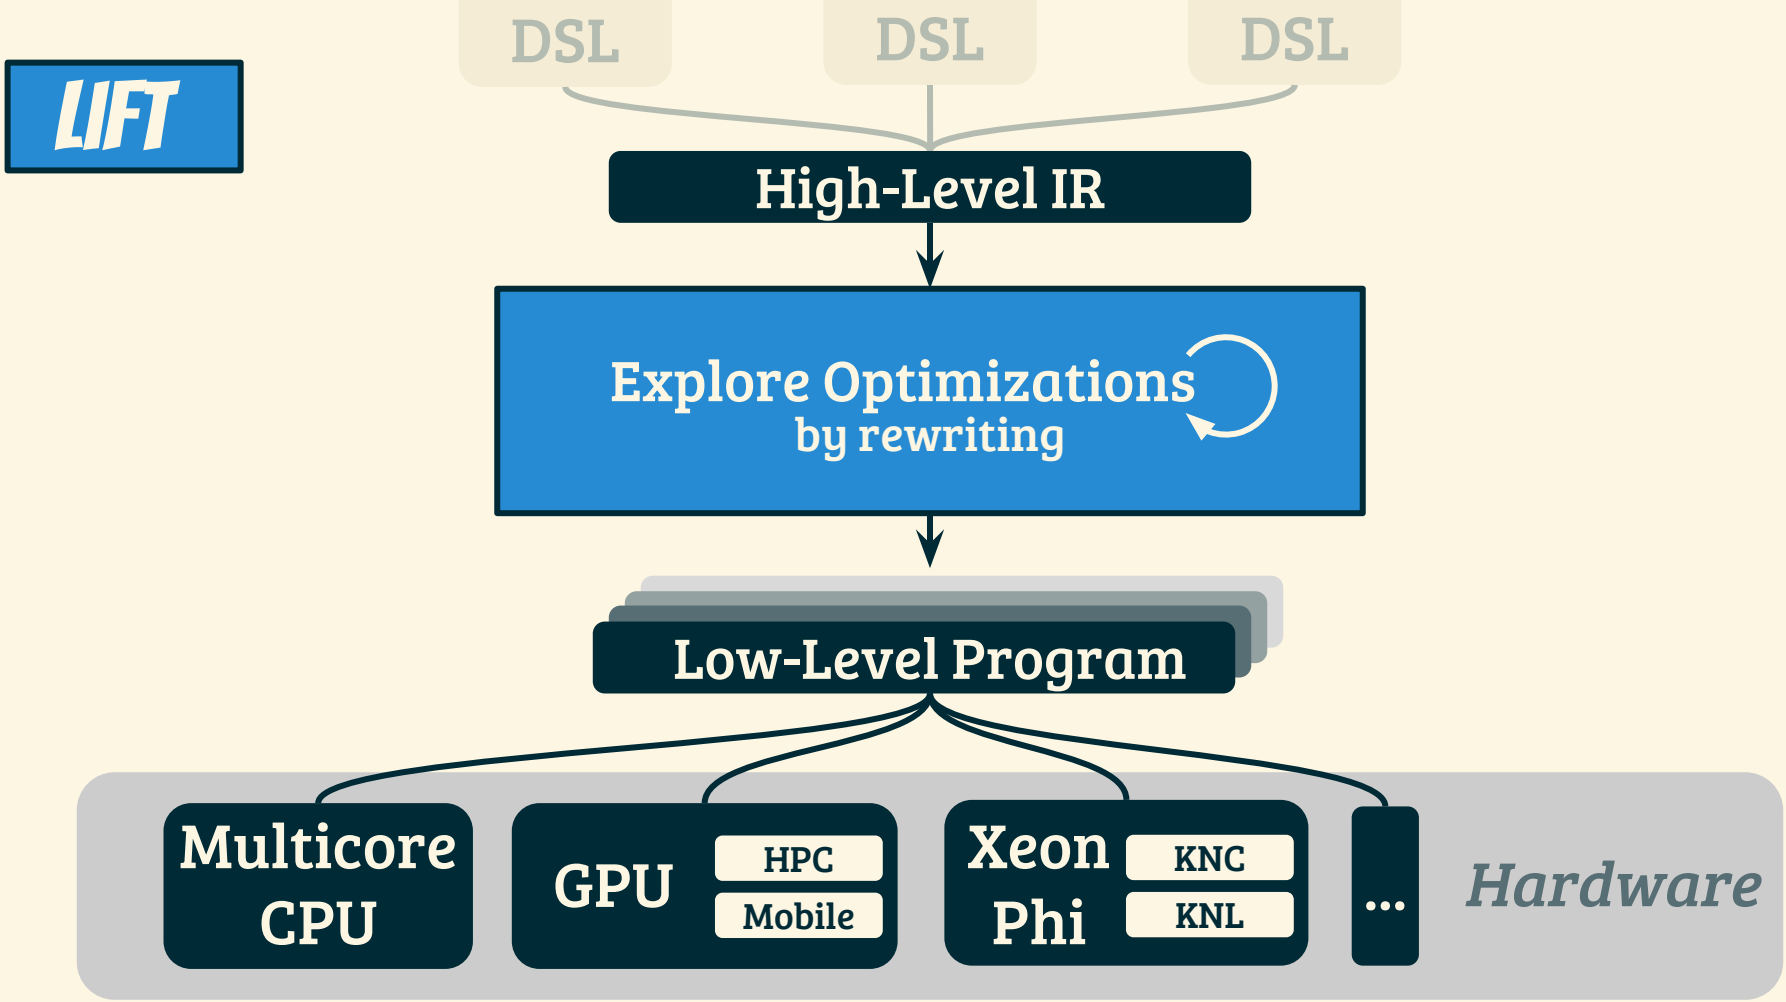
LIFT
DSL
DSL
DSL
High-Level IR
Explore Optimizations
by rewriting
Low-Level Program
Multicore CPU
GPU
HPC
Mobile
Xeon Phi
KNC
KNL
...
Hardware

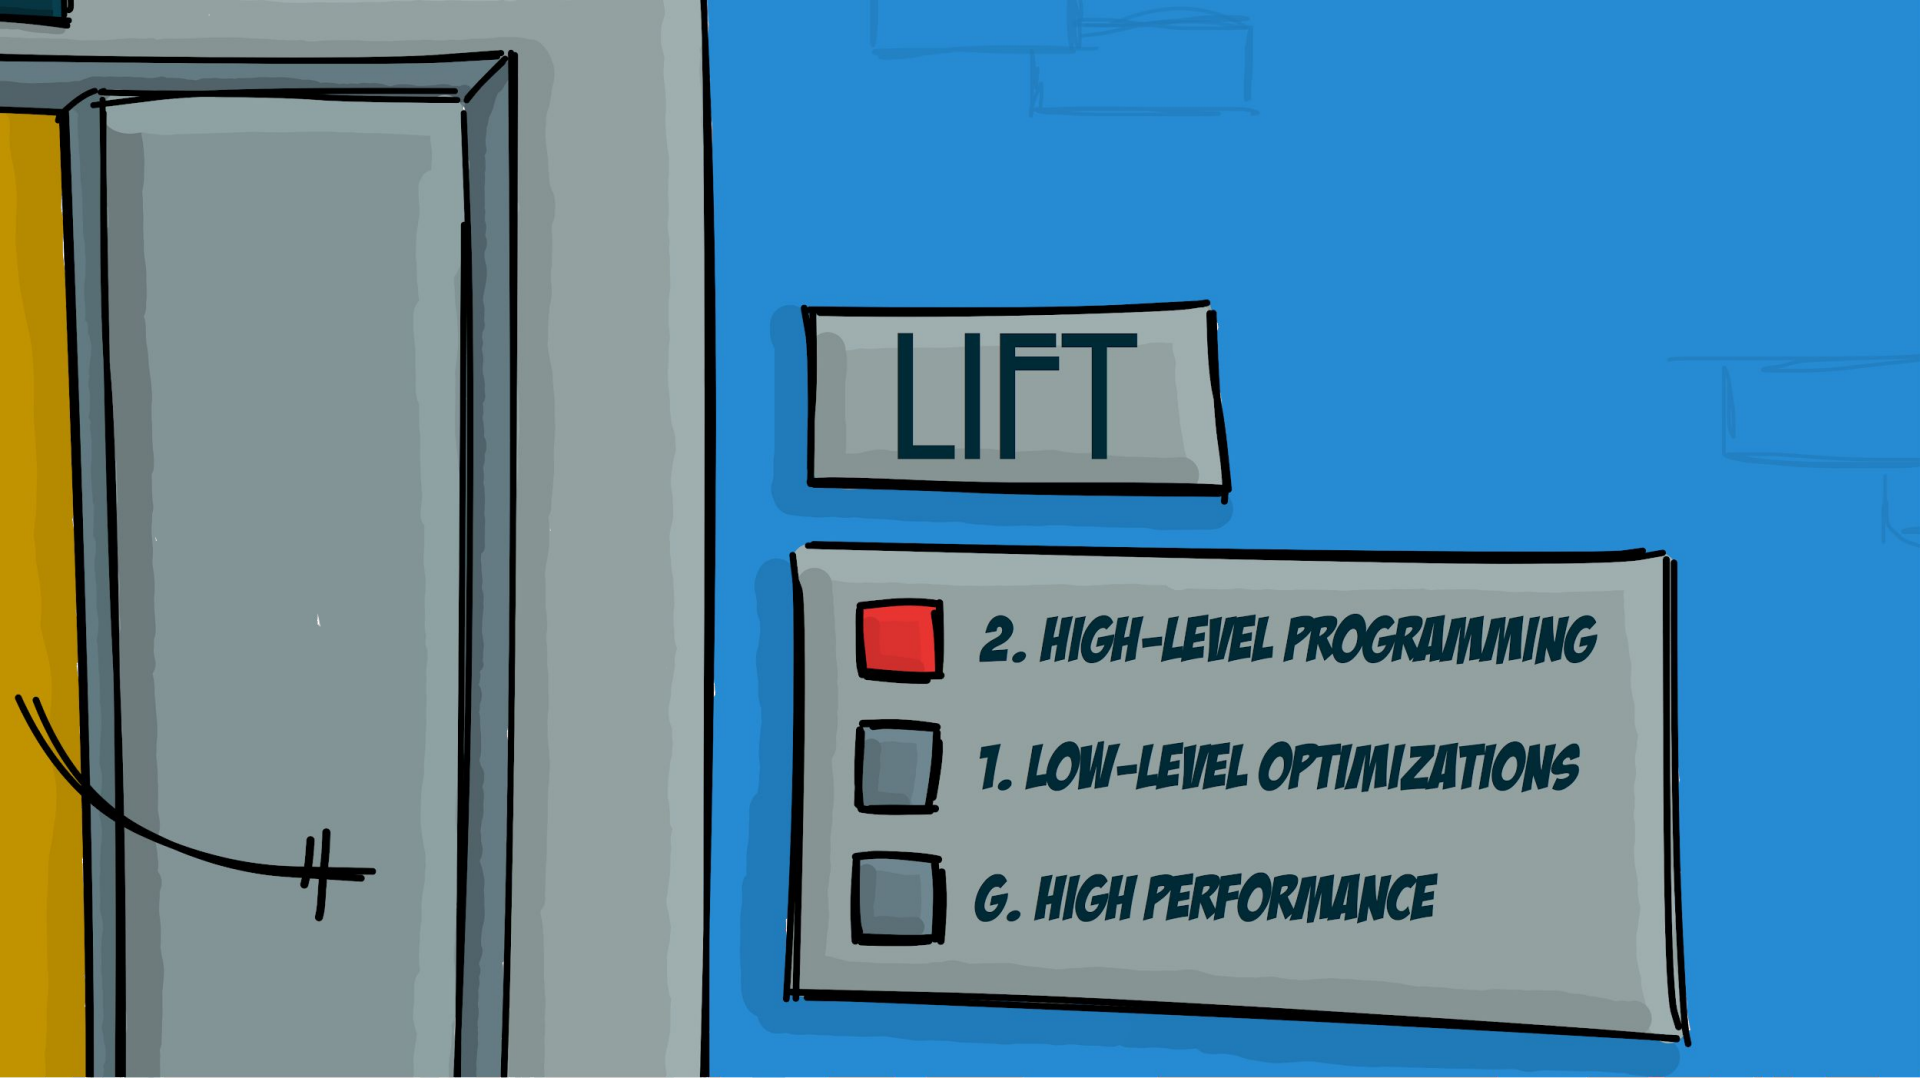
LIFT
2. HIGH-LEVEL PROGRAMMING
1. LOW-LEVEL OPTIMIZATIONS
G. HIGH PERFORMANCE

LIFT'S HIGH-LEVEL PRIMITIVES

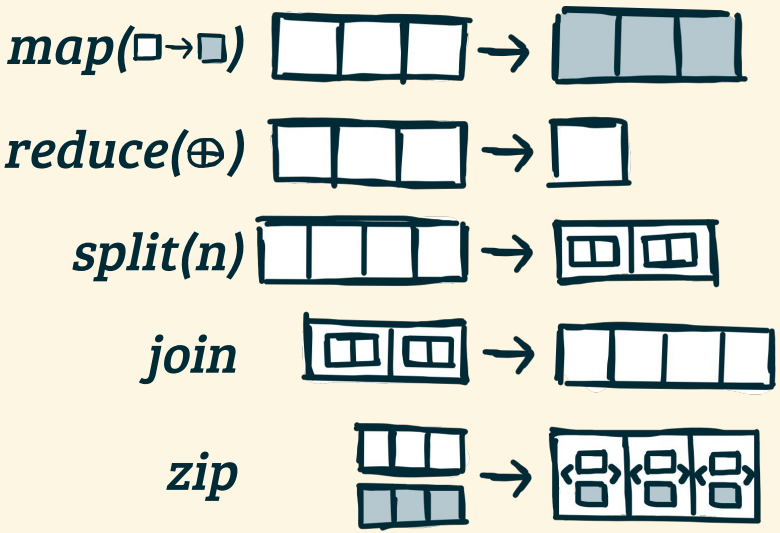

LIFT'S HIGH-LEVEL PRIMITIVES

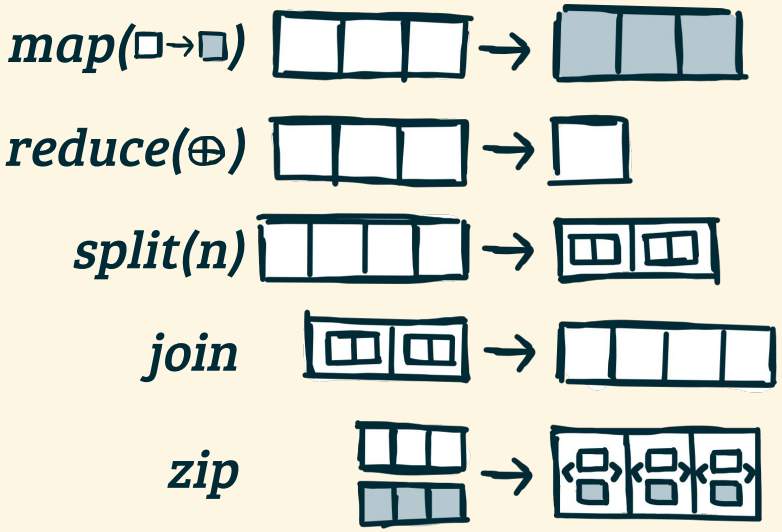

dotproduct.lift

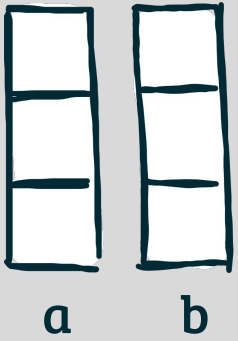

LIFT'S HIGH-LEVEL PRIMITIVES

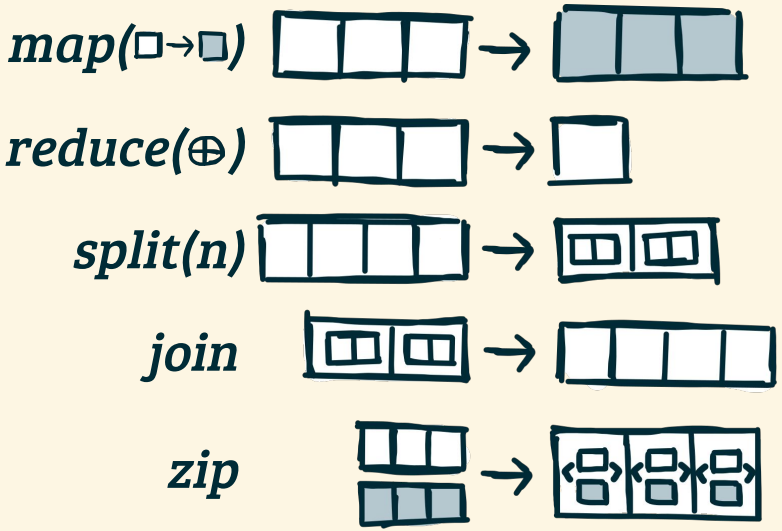

dotproduct.lift

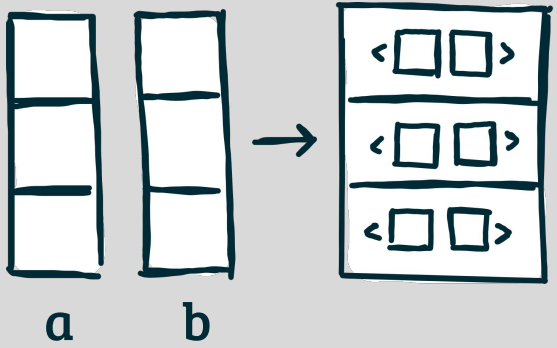

`zip(a,b)`

LIFT'S HIGH-LEVEL PRIMITIVES

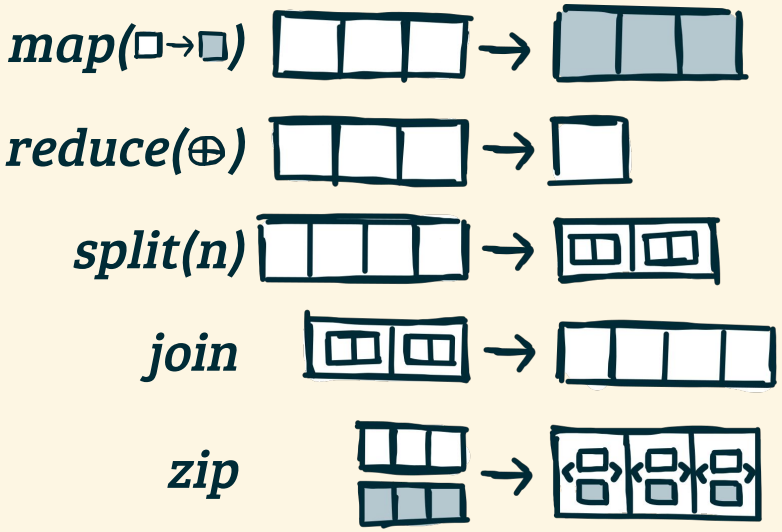

dotproduct.lift

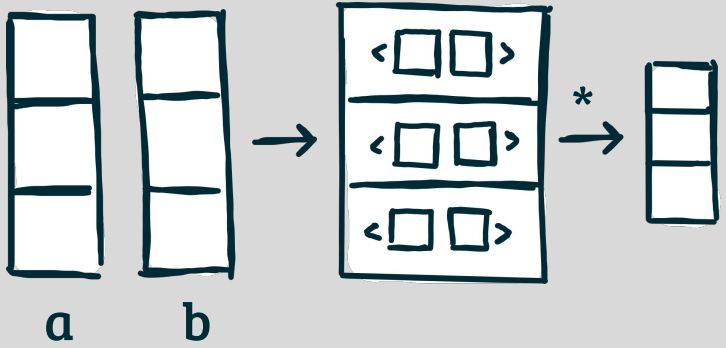


```
map(*, zip(a,b))
```

LIFT'S HIGH-LEVEL PRIMITIVES

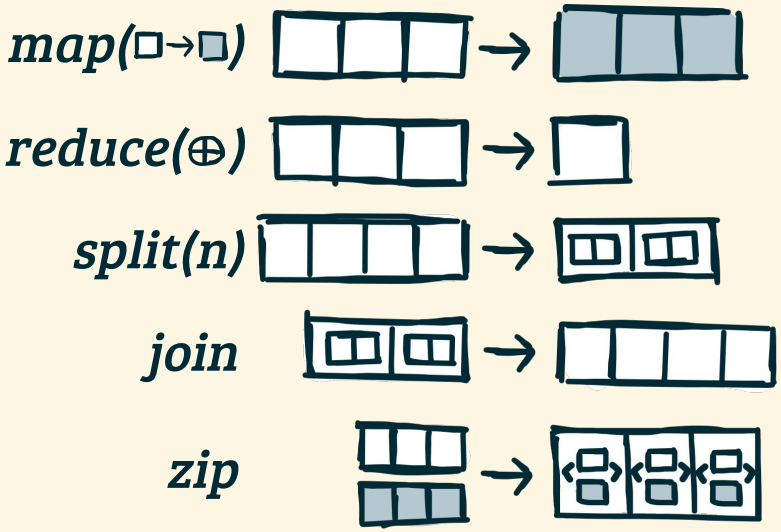

dotproduct.lift

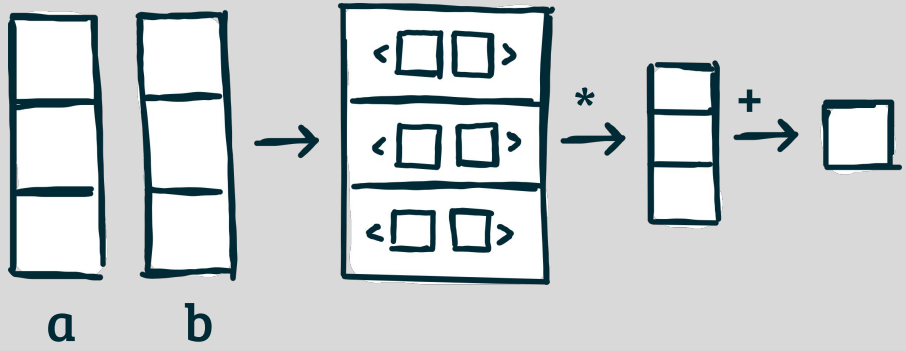


```
reduce(+,0, map(*, zip(a,b)))
```

LIFT'S HIGH-LEVEL PRIMITIVES

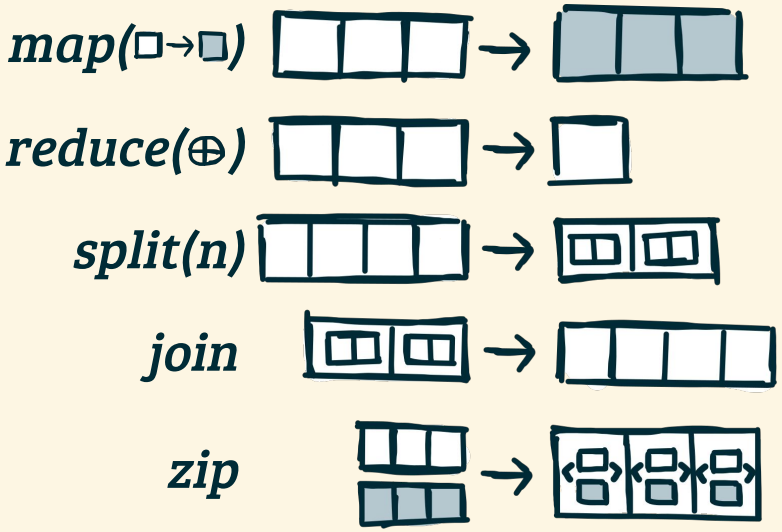

stencil.lift?

Can we express stencil computations in Lift?

Let's look at a simple stencil example...

WHAT ARE STENCIL COMPUTATIONS?

3-point-stencil.c

```
for (int i = 0; i < N ; i ++) {
    int sum = 0;
    for ( int j = -1; j <= 1; j ++) {
        int pos = i + j;
        pos = pos < 0 ? 0 : pos;
        pos = pos > N - 1 ? N - 1 : pos;
        sum += A[ pos ]; }
 B[ i ] = sum ; }
```

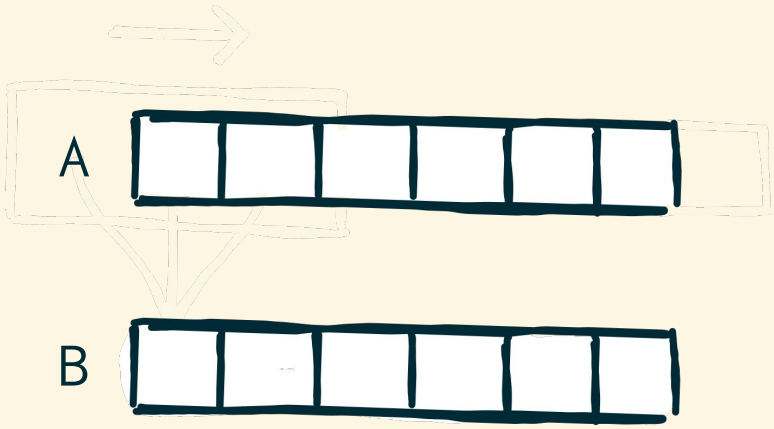

WHAT ARE STENCIL COMPUTATIONS?

3-point-stencil.c

```
for (int i = 0; i < N ; i ++) {
    int sum = 0;
    for ( int j = -1; j <= 1; j ++) {
        int pos = i + j;
        pos = pos < 0 ? 0 : pos;
        pos = pos > N - 1 ? N - 1 : pos;
        sum += A[ pos ]; }
 B[ i ] = sum ; }
```

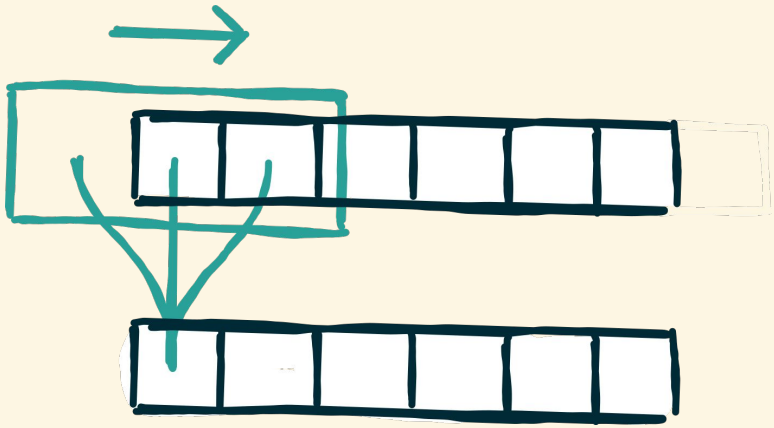

WHAT ARE STENCIL COMPUTATIONS?

3-point-stencil.c

```
for (int i = 0; i < N ; i ++) {
    int sum = 0;
    for ( int j = -1; j <= 1; j ++) {
        int pos = i + j;
        pos = pos < 0 ? 0 : pos;
        pos = pos > N - 1 ? N - 1 : pos;
        sum += A[ pos ]; }
 B[ i ] = sum ; }
```

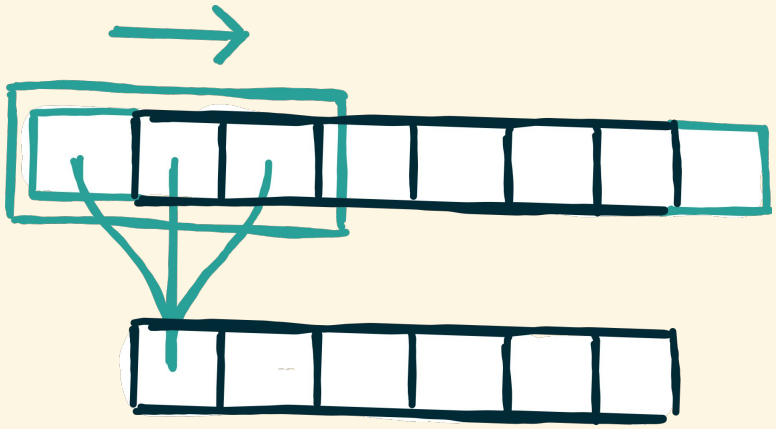

STENCIL COMPUTATIONS IN LIFT

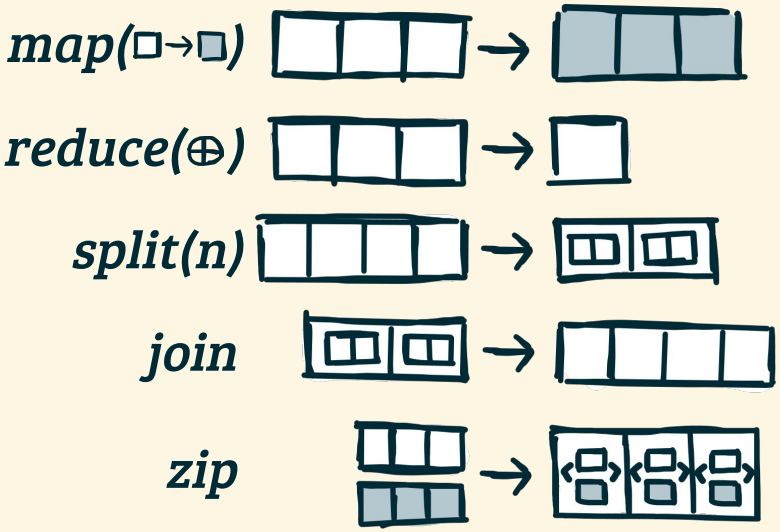

3-point-stencil.lift

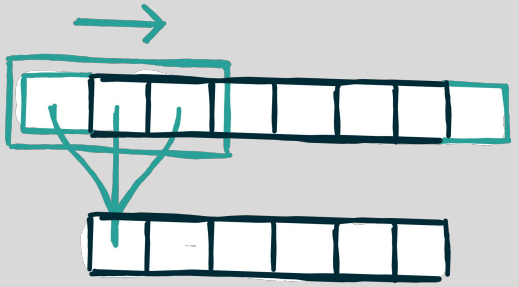

STENCIL COMPUTATIONS IN LIFT

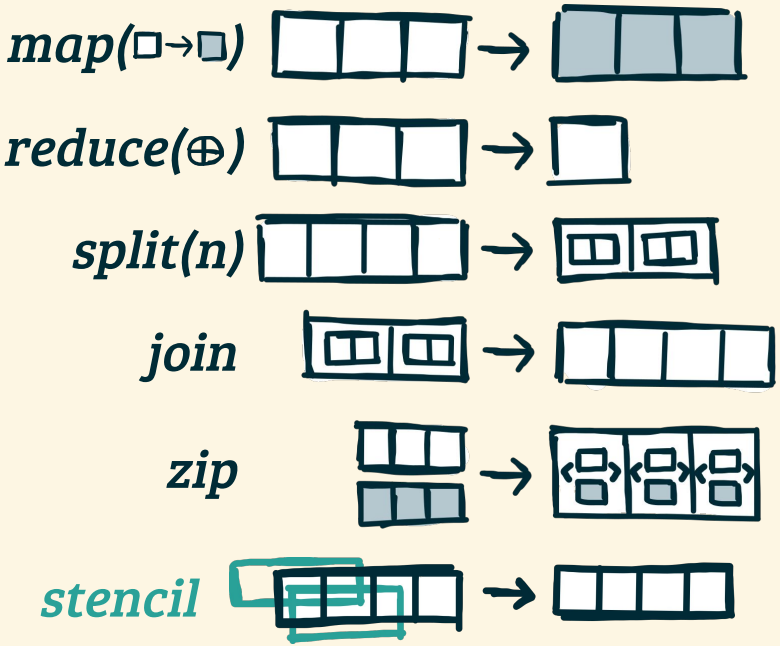

3-point-stencil.lift

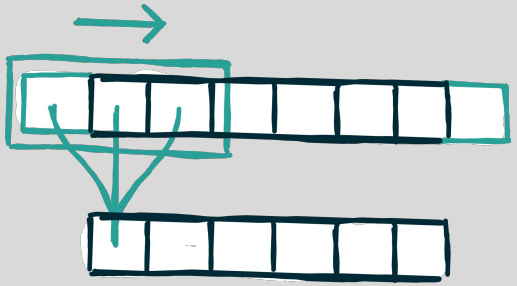

Add specialized primitive: Job done?

STENCIL COMPUTATIONS IN LIFT

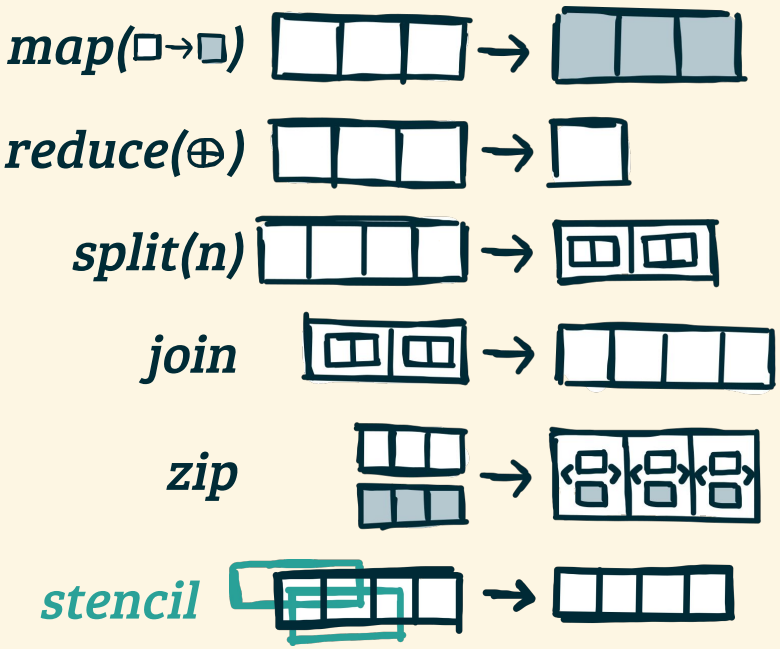

3-point-stencil.lift

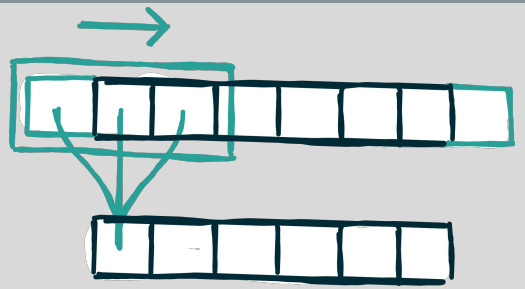

Add specialized primitive: Job done?

- ⊘ **No Reuse**
  of existing primitives and optimizations
- ⊘ **Domain-specific**
  rather than generic
- ⊘ **Multidimensional?**
  is it composable?

DECOMPOSING STENCIL COMPUTATIONS

3-point-stencil.c

```
for (int i = 0; i < N ; i ++) {
    int sum = 0;
    for ( int j = -1; j <= 1; j ++) {
        int pos = i + j;
        pos = pos < 0 ? 0 : pos;
        pos = pos > N - 1 ? N - 1 : pos;
        sum += A[ pos ]; }
 B[ i ] = sum ; }
```

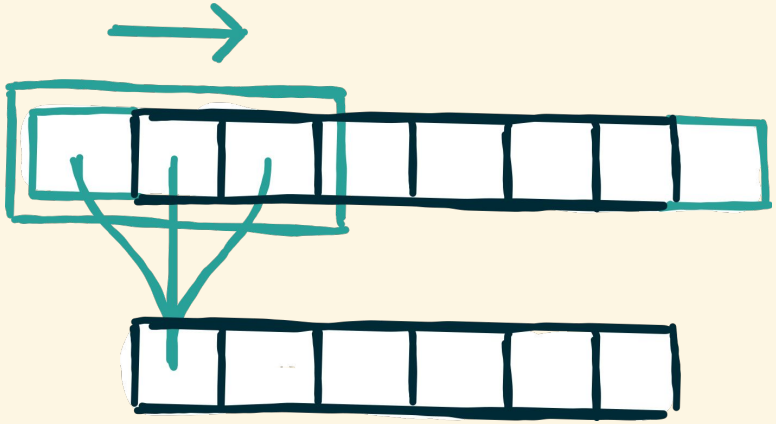

DECOMPOSING STENCIL COMPUTATIONS

3-point-stencil.c

```
for (int i = 0; i < N ; i ++) {
    int sum = 0;
    for ( int j = -1; j <= 1; j ++) {   // ( a )
        int pos = i + j;
        pos = pos < 0 ? 0 : pos;
        pos = pos > N - 1 ? N - 1 : pos;
        sum += A[ pos ]; }
 B[ i ] = sum ; }
```

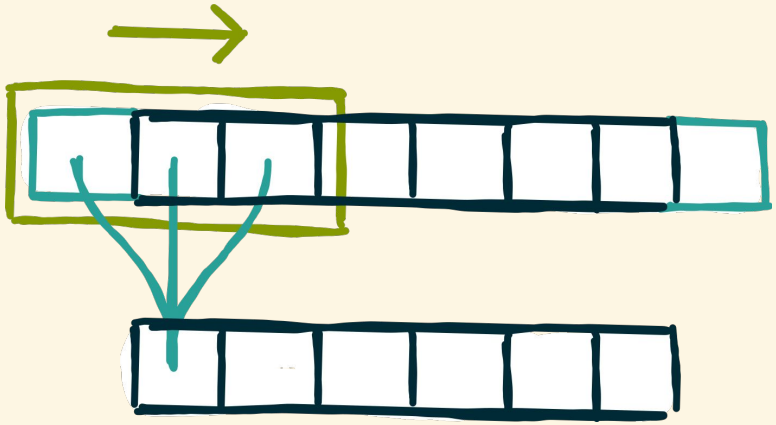

(a) access neighborhoods for every element

DECOMPOSING STENCIL COMPUTATIONS

3-point-stencil.c

```
for (int i = 0; i < N ; i ++) {
    int sum = 0;
    for ( int j = -1; j <= 1; j ++) {   // ( a )
        int pos = i + j;
        pos = pos < 0 ? 0 : pos;        // ( b )
        pos = pos > N - 1 ? N - 1 : pos;
        sum += A[ pos ]; }
 B[ i ] = sum ; }
```

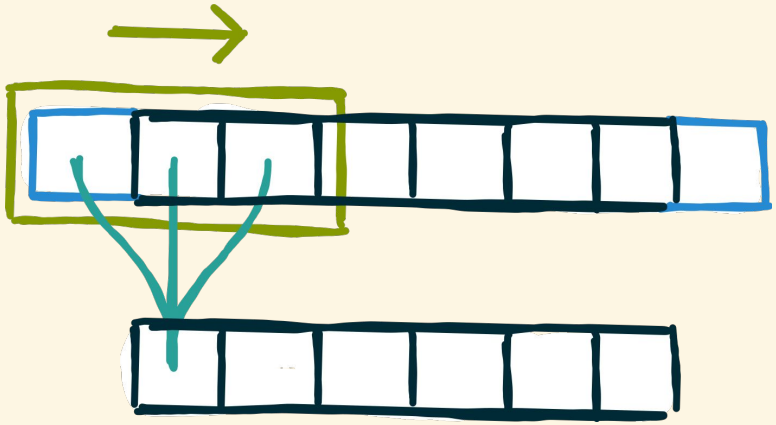

(a) access neighborhoods for every element

(b) specify boundary handling

DECOMPOSING STENCIL COMPUTATIONS

3-point-stencil.c

```
for (int i = 0; i < N ; i ++) {
    int sum = 0;
    for ( int j = -1; j <= 1; j ++) {   // ( a )
        int pos = i + j;
        pos = pos < 0 ? 0 : pos;        // ( b )
        pos = pos > N - 1 ? N - 1 : pos;
        sum += A[ pos ]; }              // ( c )
 B[ i ] = sum ; }
```

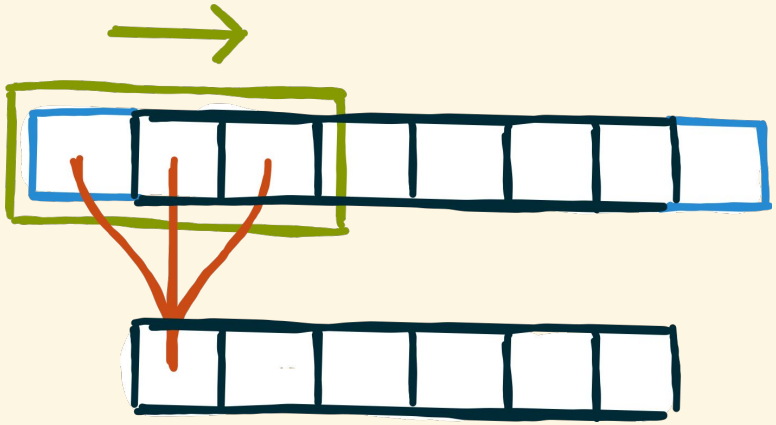

(a) access neighborhoods for every element
(b) specify boundary handling
(c) apply stencil function to neighborhoods

DECOMPOSING STENCIL COMPUTATIONS

3-point-stencil.c

```
for (int i = 0; i < N ; i ++) {
    int sum = 0;
    for ( int j = -1; j <= 1; j ++) {   // ( a )
        int pos = i + j;
        pos = pos < 0 ? 0 : pos;        // ( b )
        pos = pos > N - 1 ? N - 1 : pos;
        sum += A[ pos ]; }              // ( c )
 B[ i ] = sum ; }
```

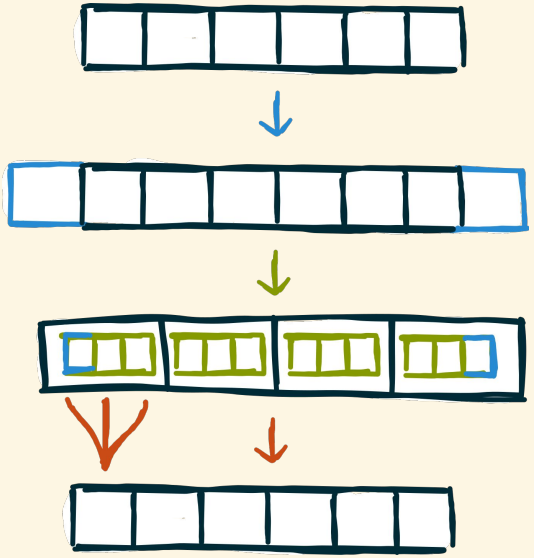

(a) access neighborhoods for every element
(b) specify boundary handling
(c) apply stencil function to neighborhoods

STENCIL COMPUTATIONS IN LIFT

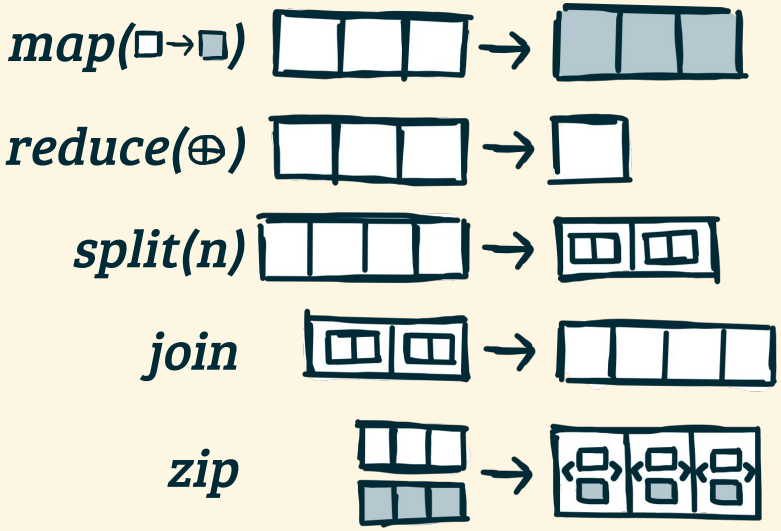

3-point-stencil.lift

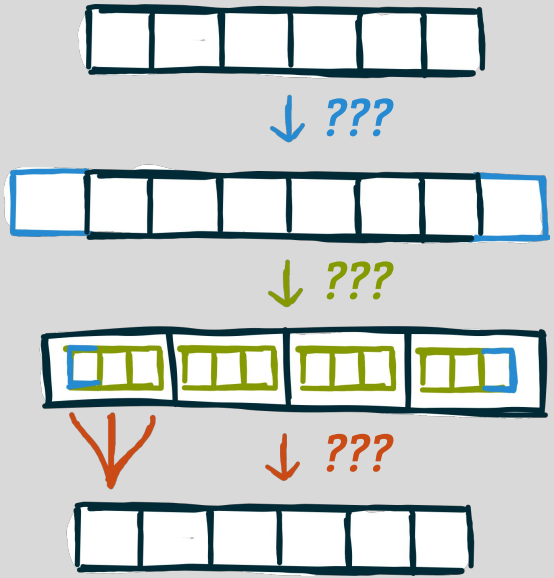

STENCIL COMPUTATIONS IN LIFT

map(□→□) 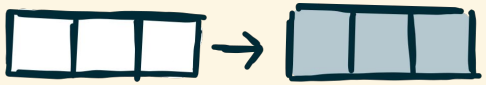

reduce(⊕) 

split(n) 

join 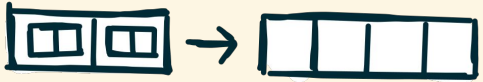

zip 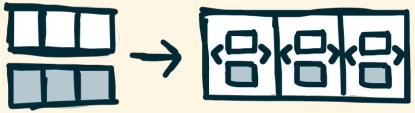

3-point-stencil.lift

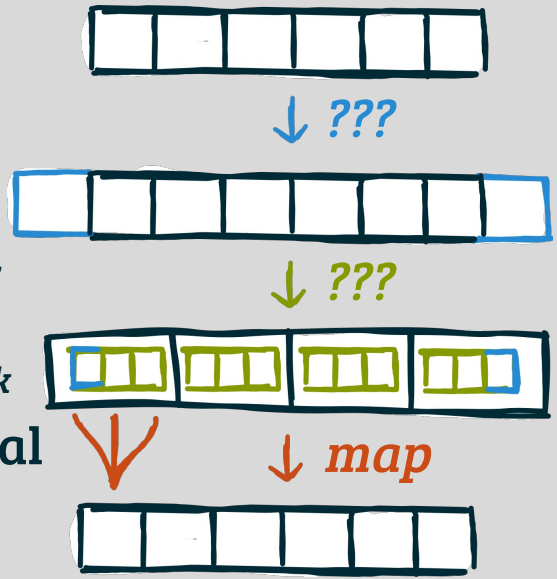

- **Reuse *map***
  allows to reuse existing rewrite rules
- **Simplicity**
  one primitive per task
- **Multidimensional**
  easily composable

BOUNDARY HANDLING USING PAD

pad (reindexing)

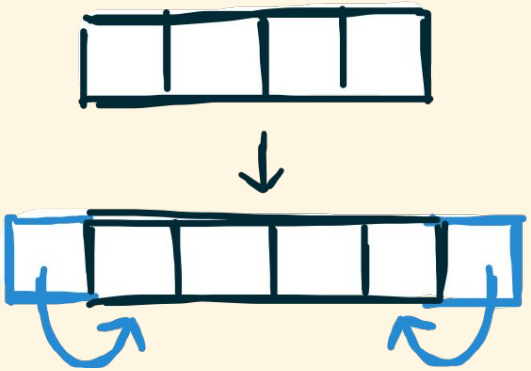

pad (constant)

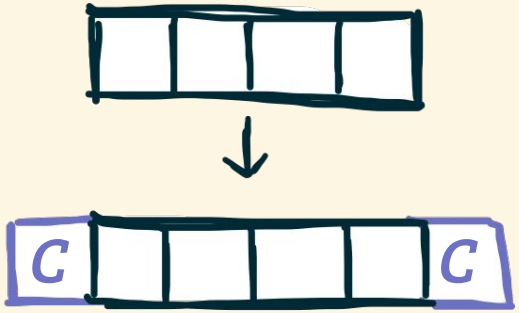

pad-reindexing.lift

```
clamp(i, n) = (i <  0) ? 0  :
              ((i >= n) ? n-1:i)

pad(1,1,clamp, [a,b,c,d]) =
    [a,a,b,c,d,d]
```

pad-constant.lift

```
constant(i, n) = C


pad(1,1,constant, [a,b,c,d]) =
    [C,a,b,c,d,C]
```

NEIGHBORHOOD CREATION USING SLIDE

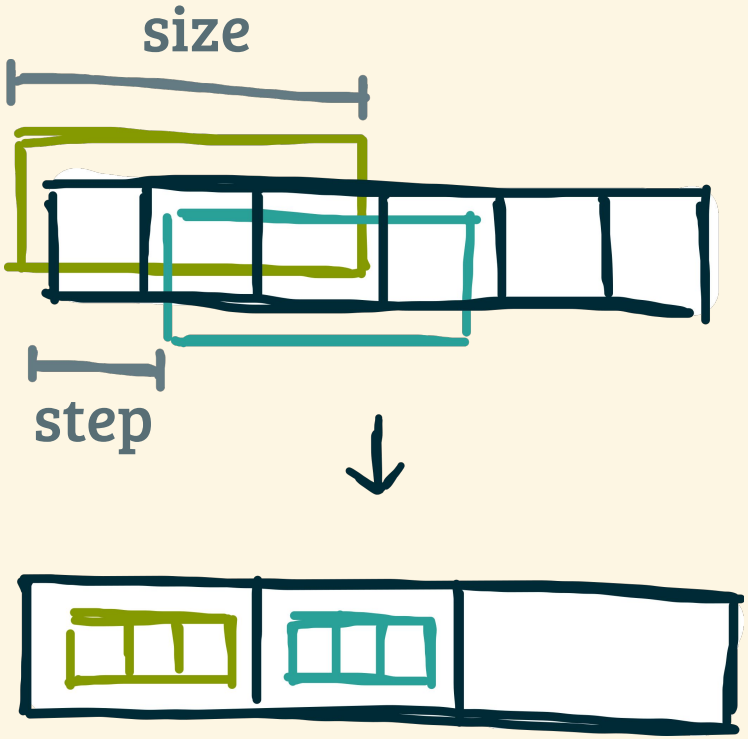

slide-example.lift

```
slide(3,1,[a,b,c,d,e]) =
    [[a,b,c],[b,c,d],[c,d,e]]
```

Applying stencil function using map

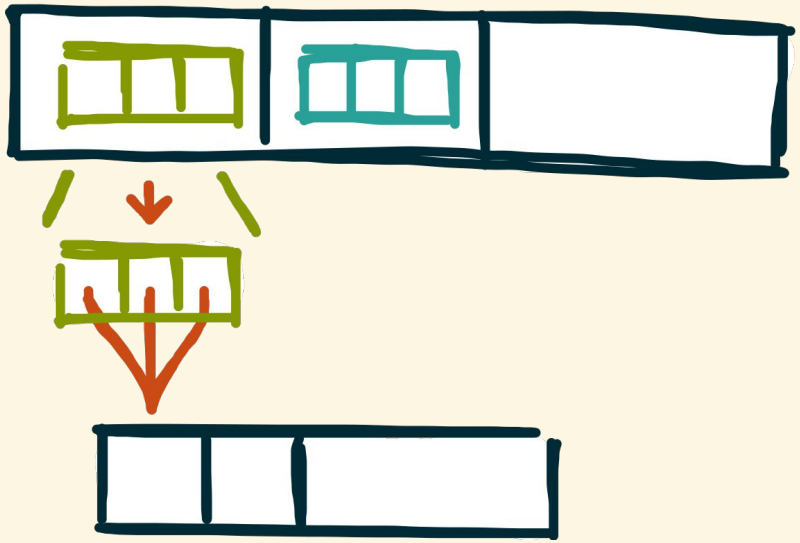

sum-neighborhoods.lift

```
map(nbh =>
  reduce(add, 0.0f, nbh))
```

PUTTING IT TOGETHER

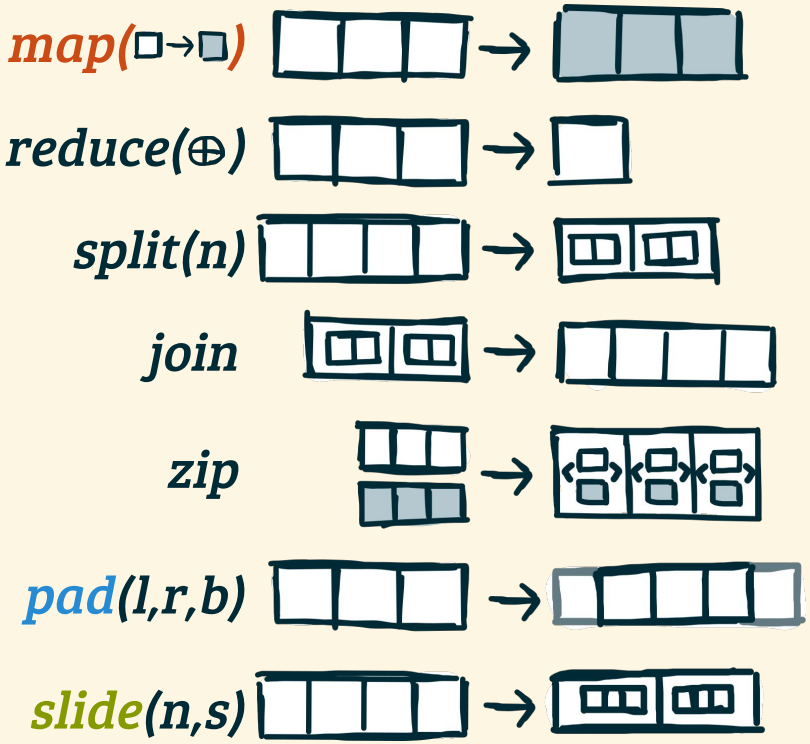

stencil1D.lift

```
def stencil1D =
 fun(A =>
  map(reduce(add, 0.0f),
   slide(3,1,
    pad(1,1,clamp,A))))
```

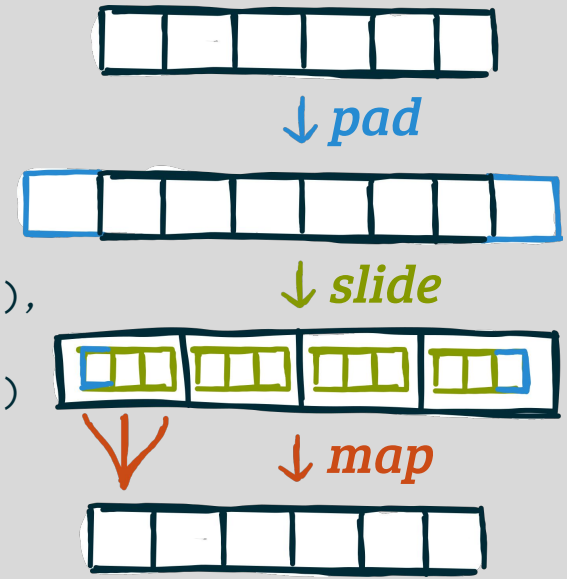

MULTIDIMENSIONAL STENCIL COMPUTATIONS

are expressed as compositions of intuitive, generic 1D primitives

Decompose to Re-Compose

MULTIDIMENSIONAL STENCIL COMPUTATIONS

are expressed as compositions of intuitive, generic 1D primitives

Decompose to Re-Compose

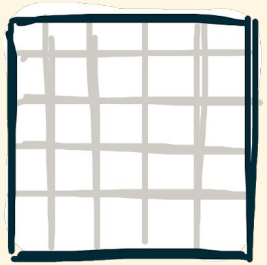

MULTIDIMENSIONAL STENCIL COMPUTATIONS

are expressed as compositions of intuitive, generic 1D primitives

Decompose to Re-Compose

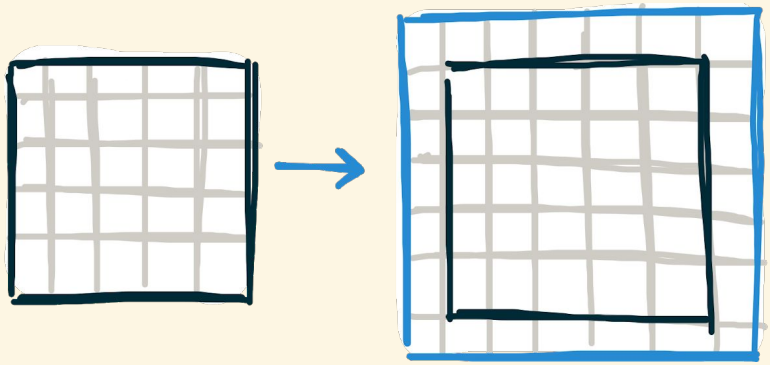

$pad_2(1,1,clamp,\text{input})$

MULTIDIMENSIONAL BOUNDARY HANDLING USING pad_2

input

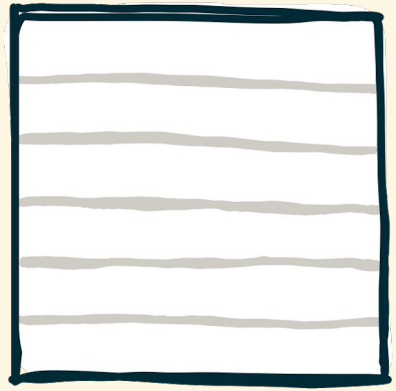

pad_2 =

MULTIDIMENSIONAL BOUNDARY HANDLING USING pad_2

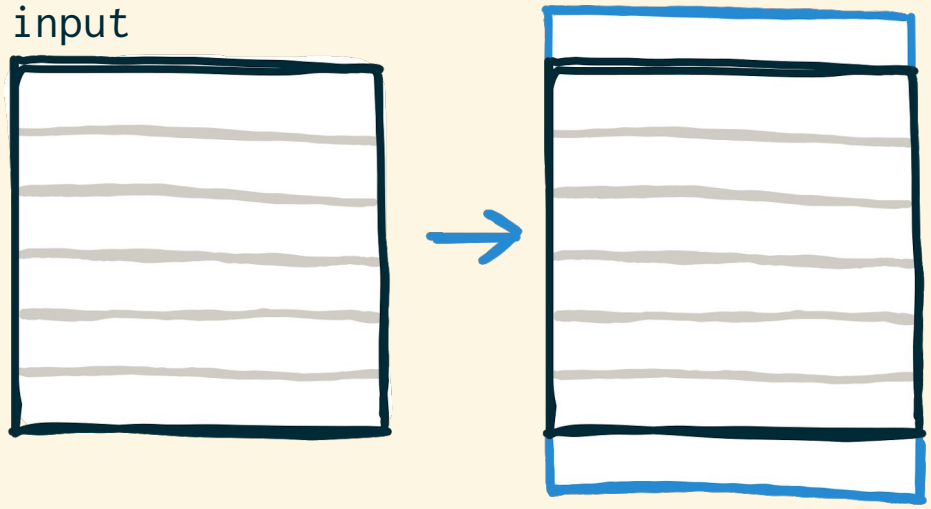

pad_2 = ***pad***(l,r,b,input)

MULTIDIMENSIONAL BOUNDARY HANDLING USING pad_2

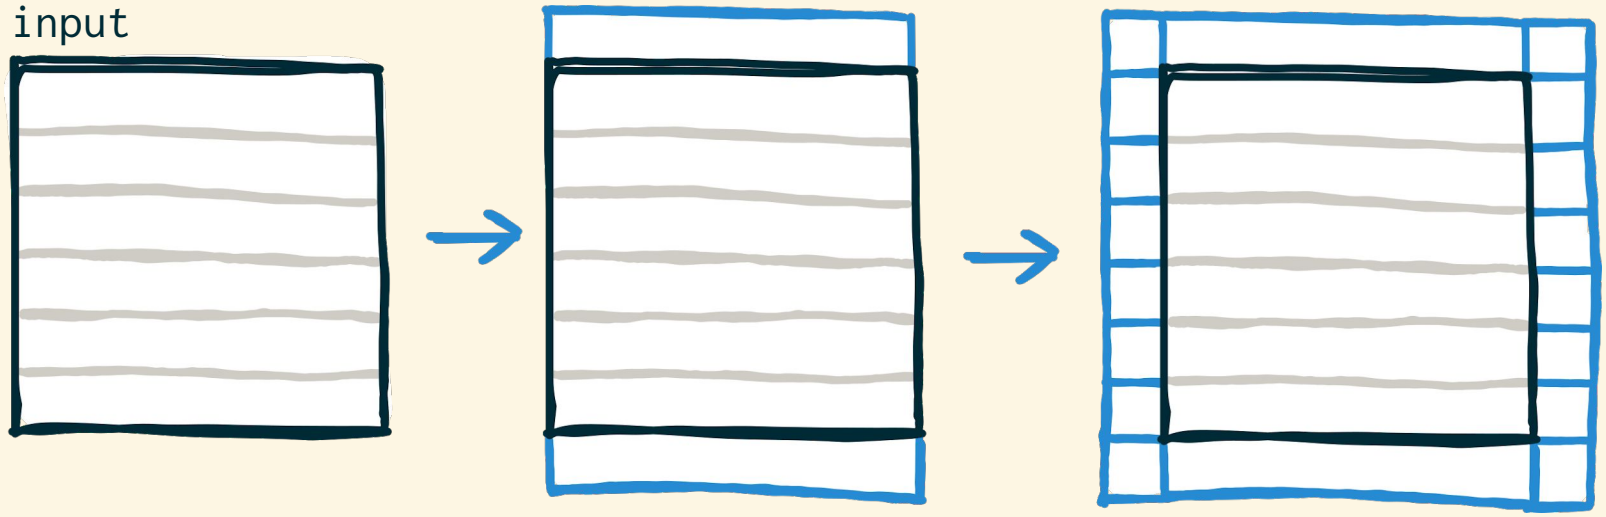

$$pad_2 = map(pad(l,r,b,pad(l,r,b,input)))$$

MULTIDIMENSIONAL STENCIL COMPUTATIONS

are expressed as compositions of intuitive, generic 1D primitives

Decompose to Re-Compose

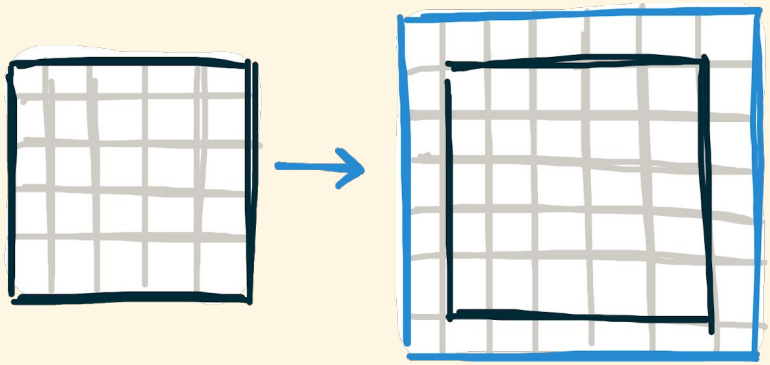

$pad_2(1, 1, clamp, input)$

MULTIDIMENSIONAL STENCIL COMPUTATIONS

are expressed as compositions of intuitive, generic 1D primitives

Decompose to Re-Compose

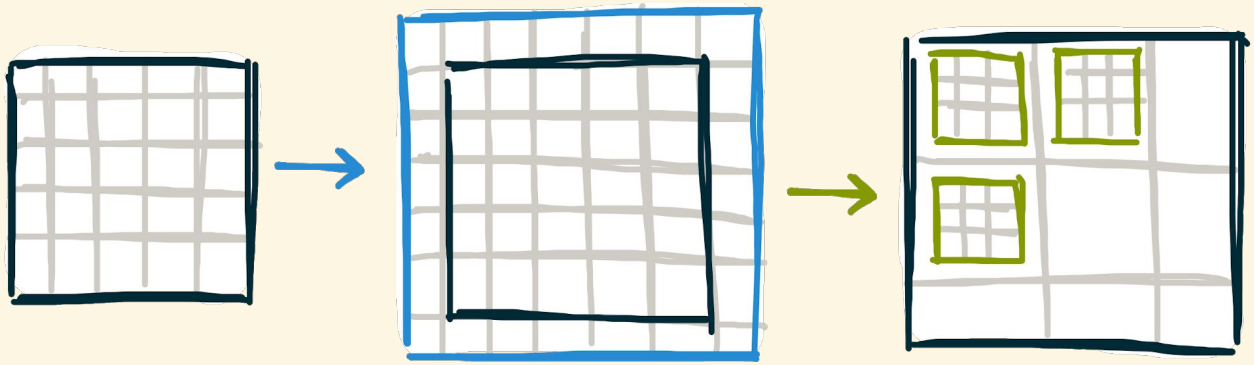

$$slide_2(3,1,\ pad_2(1,1,clamp,input))$$

MULTIDIMENSIONAL STENCIL COMPUTATIONS

are expressed as compositions of intuitive, generic 1D primitives

Decompose to Re-Compose

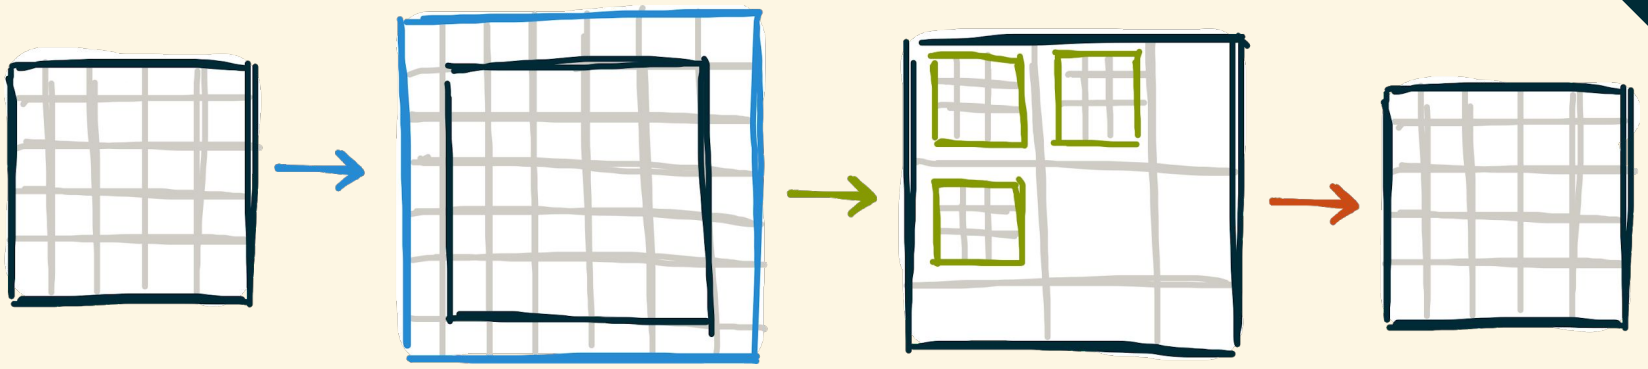

$$map_2(sum,\ slide_2(3,1,\ pad_2(1,1,clamp,input)))$$

MULTIDIMENSIONAL STENCIL COMPUTATIONS

are expressed as compositions of intuitive, generic 1D primitives

Decompose to Re-Compose

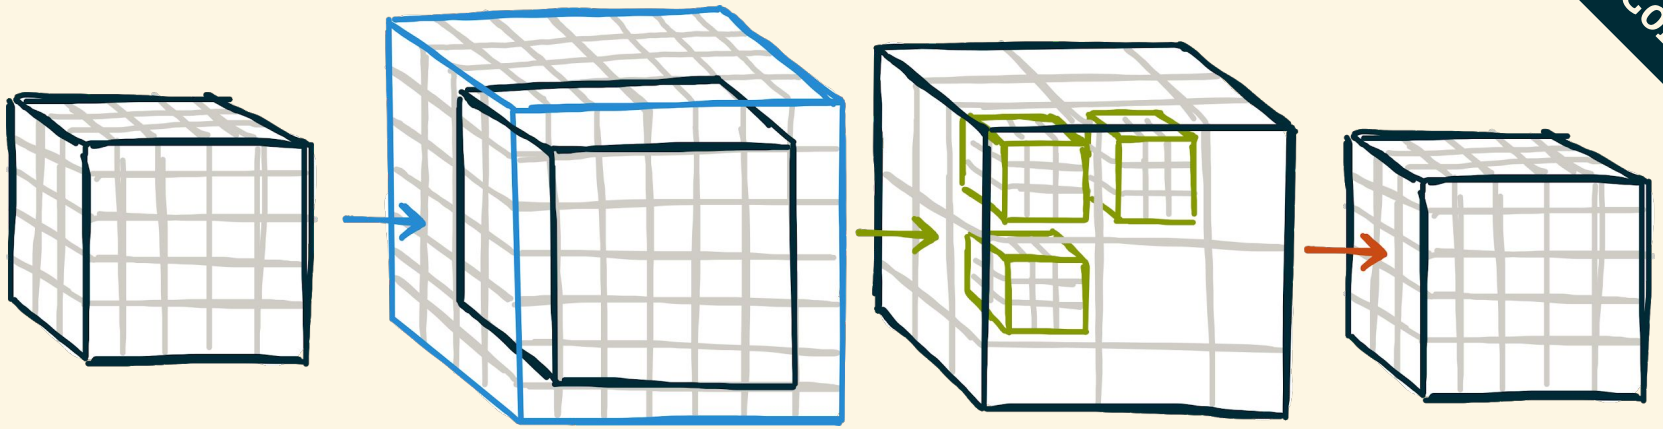

$$map_3(sum,\ slide_3(3,1,\ pad_3(1,1,clamp,input)))$$

MULTIDIMENSIONAL STENCIL COMPUTATIONS

are expressed as compositions of intuitive, generic 1D primitives

Decompose to Re-Compose

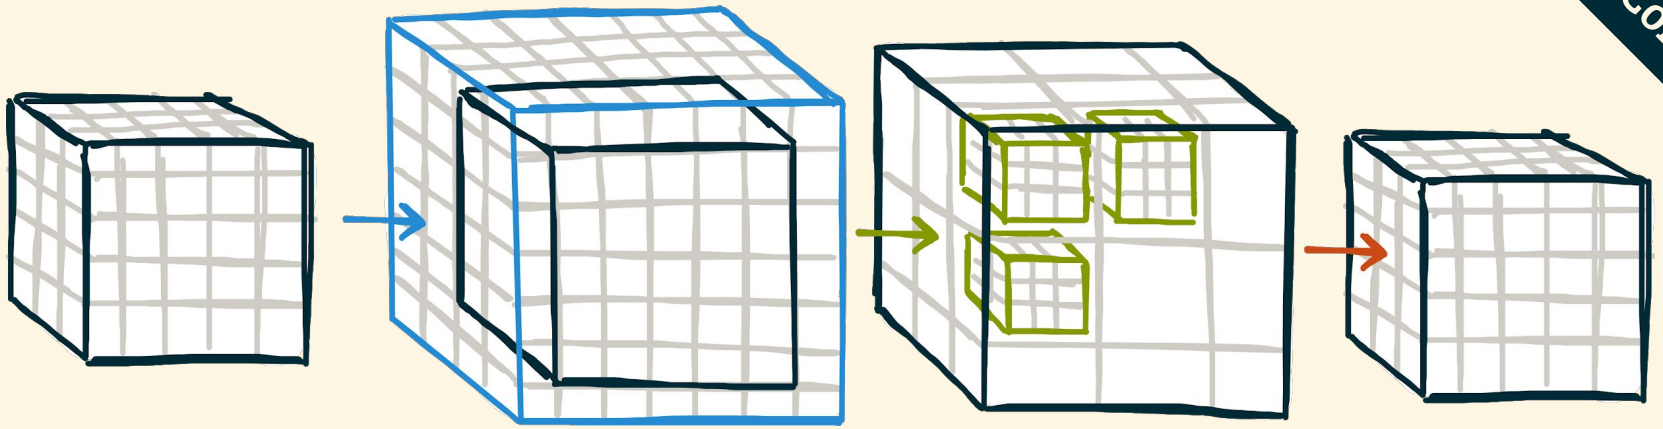

$$map_3(sum,\ slide_3(3,1,\ pad_3(1,1,clamp,input)))$$

Advantages: ✓ Compact Language ✓ Reuse Rewrites ✓ Simple Compilation

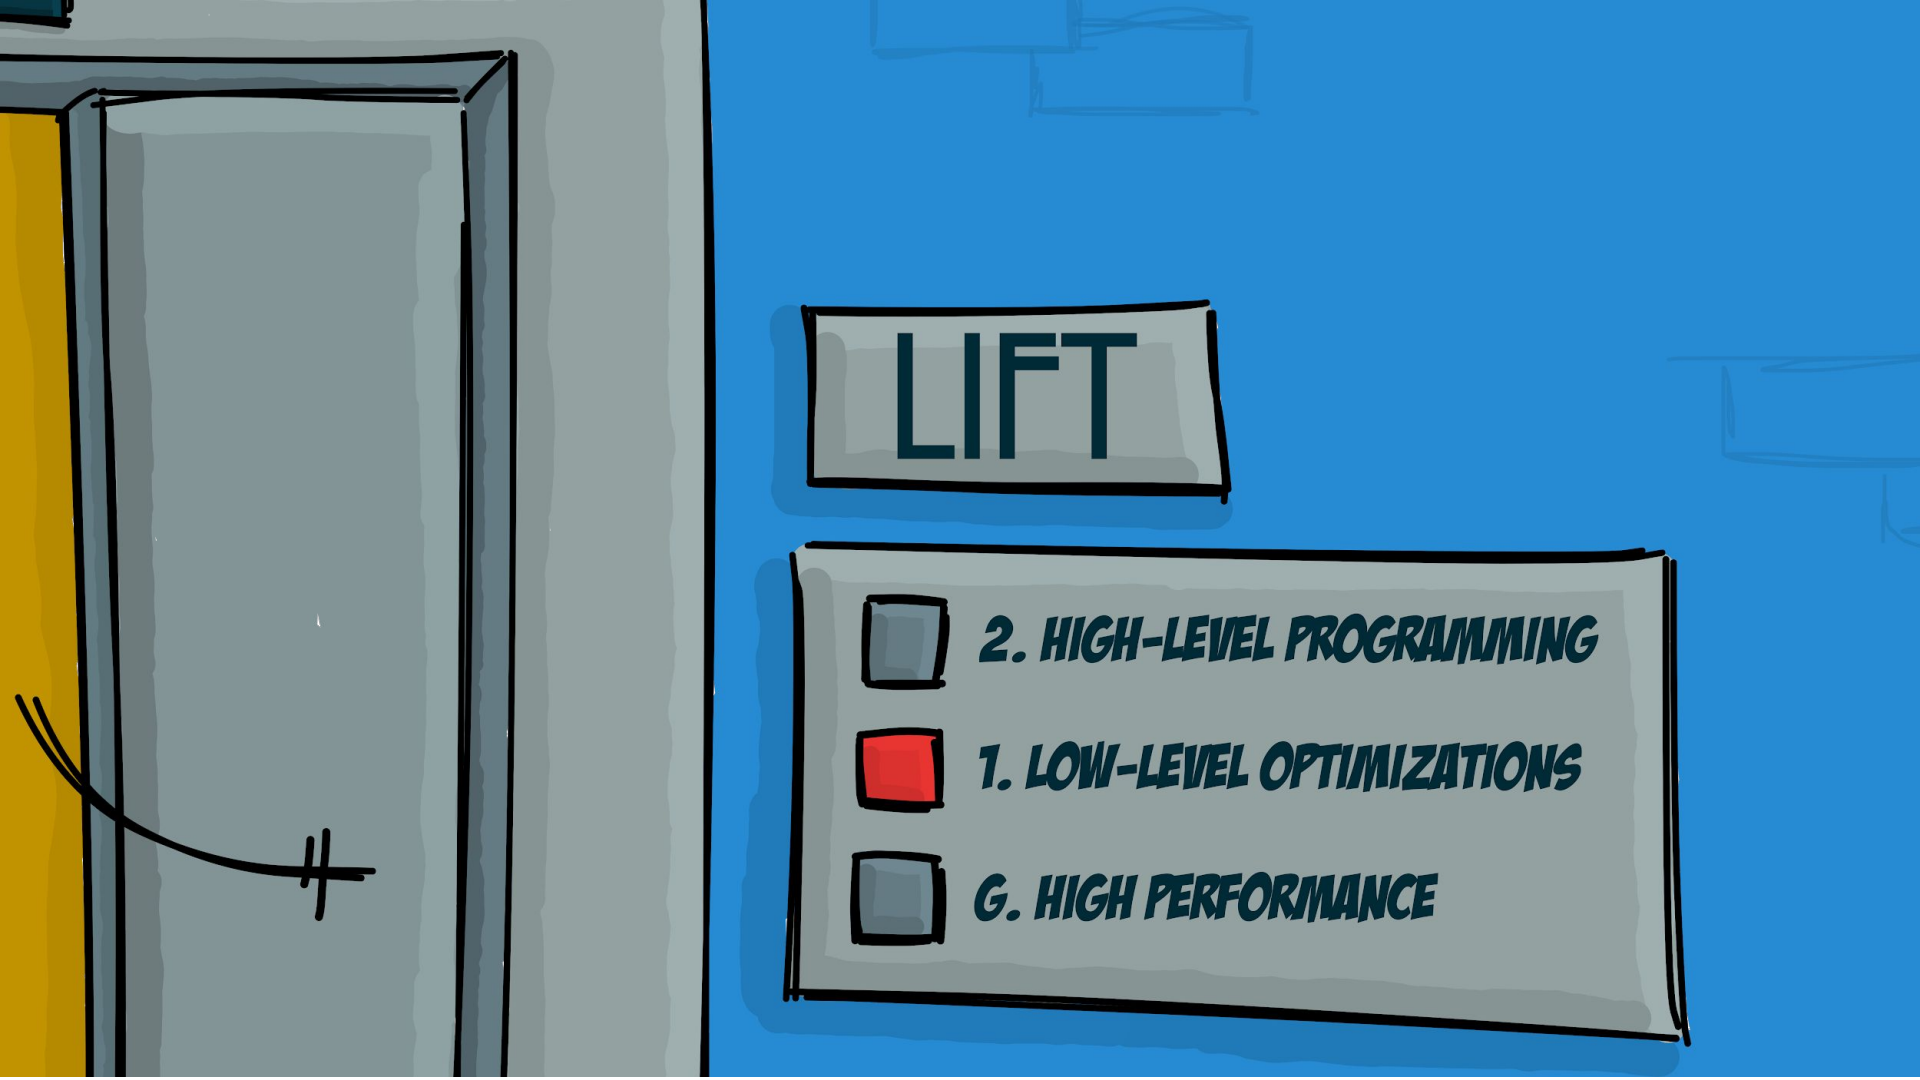
LIFT
2. HIGH-LEVEL PROGRAMMING
1. LOW-LEVEL OPTIMIZATIONS
G. HIGH PERFORMANCE

REUSING EXISTING REWRITE RULES

Divide & Conquer

map(f, A)

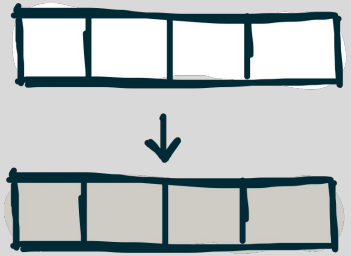

REUSING EXISTING REWRITE RULES

Divide & Conquer

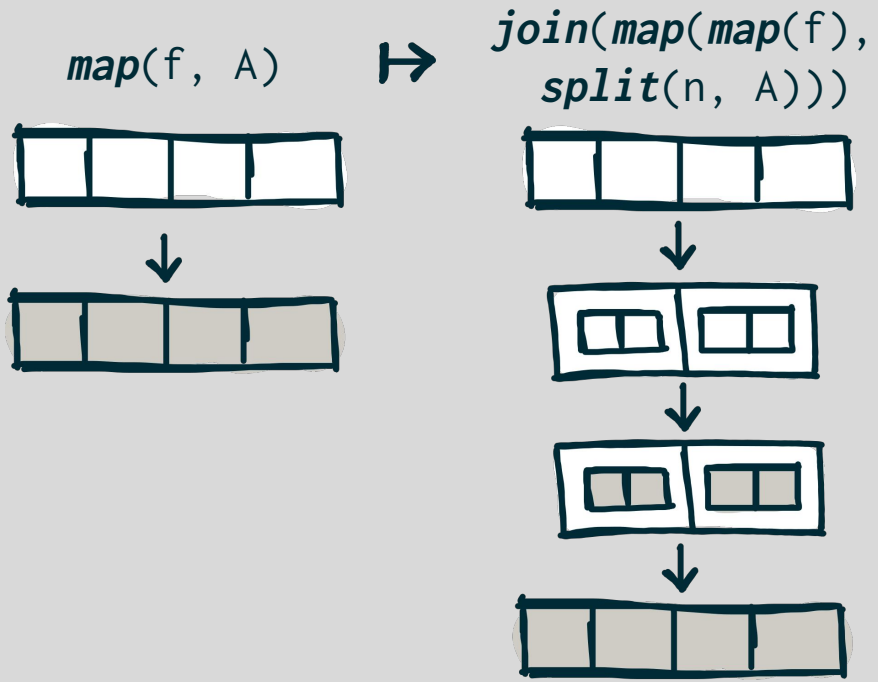

Optimization: Overlapped Tiling

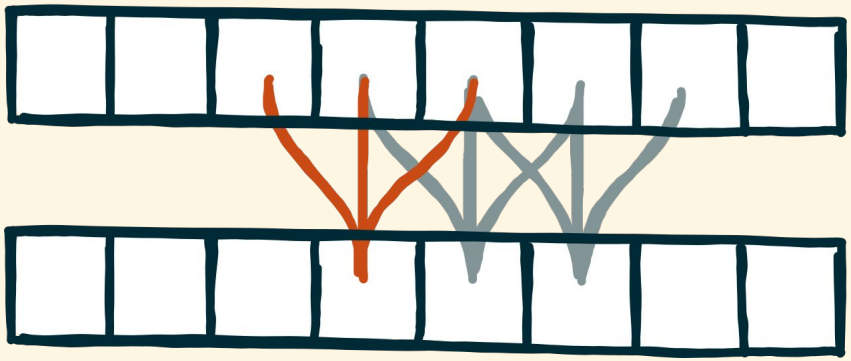

- **Exploit Locality**

  Close neighborhoods share elements that can be grouped in tiles

- **Shared Memory**

  Fast memory can be used to cache tiles

Optimization: Overlapped Tiling

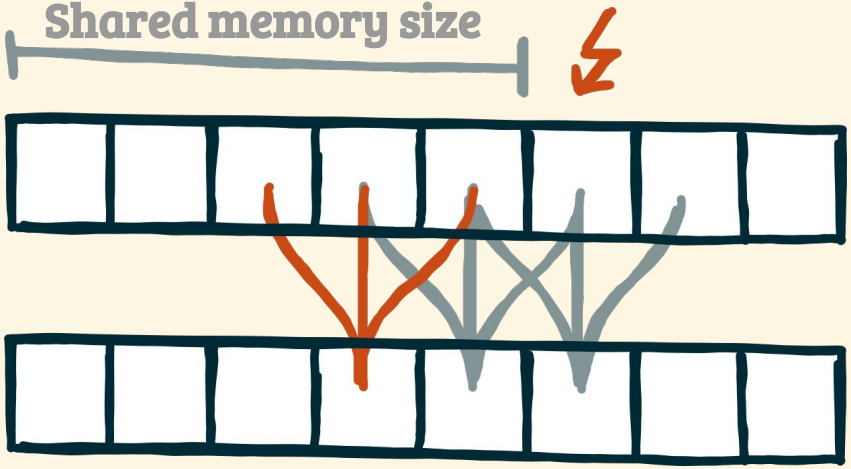

Exploit Locality

Close neighborhoods share elements that can be grouped in tiles

Shared Memory

Fast memory can be used to cache tiles

Optimization: Overlapped Tiling

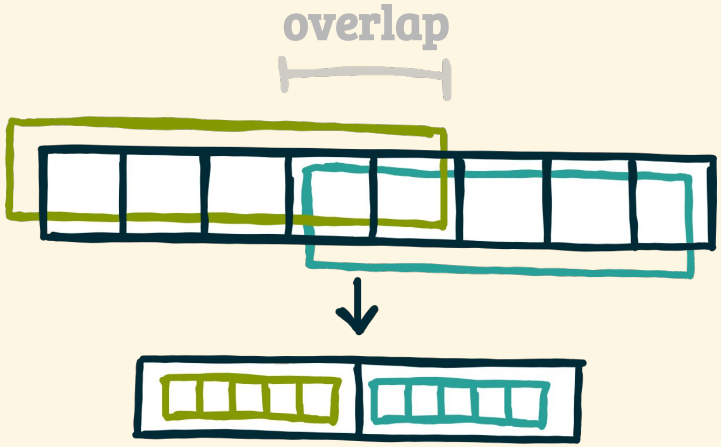

- **Exploit Locality**

  Close neighborhoods share elements that can be grouped in tiles

- **Shared Memory**

  Fast memory can be used to cache tiles

Optimization: Overlapped Tiling

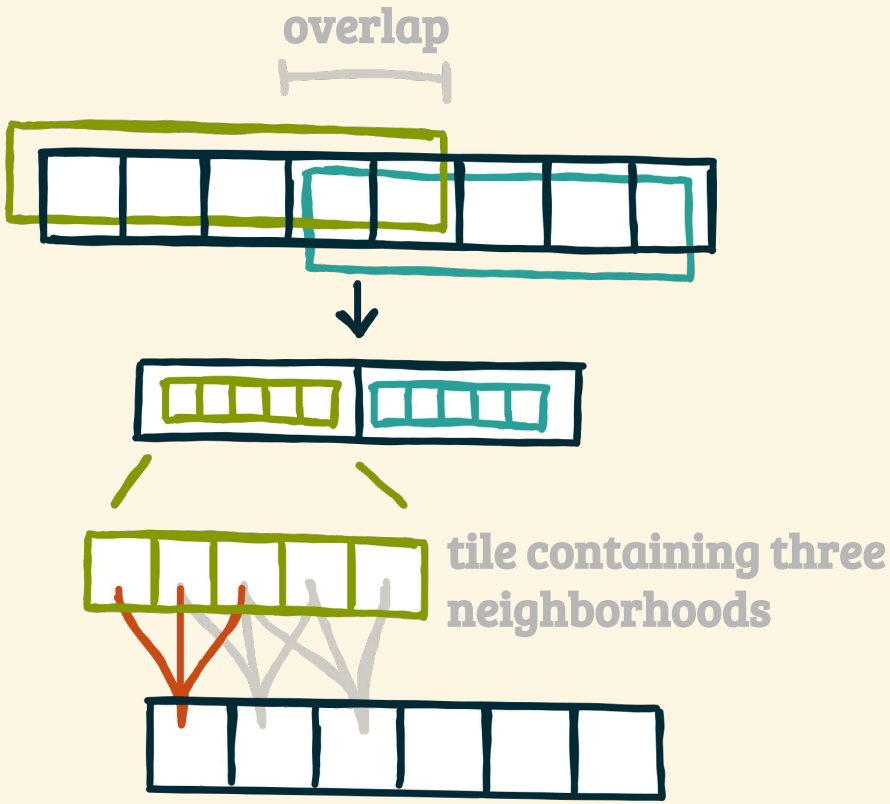

Exploit Locality

Close neighborhoods share elements that can be grouped in tiles

Shared Memory

Fast memory can be used to cache tiles

OVERLAPPED TILING AS A REWRITE RULE

overlapped tiling rule

```
map(f, slide(3,1,input))
```

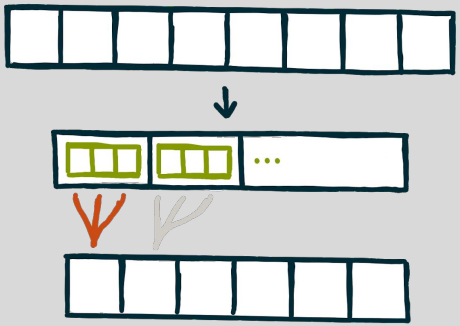

OVERLAPPED TILING AS A REWRITE RULE

overlapped tiling rule

map(f, ***slide***(3,1,input)) ↦ ***slide***(u,v,input)

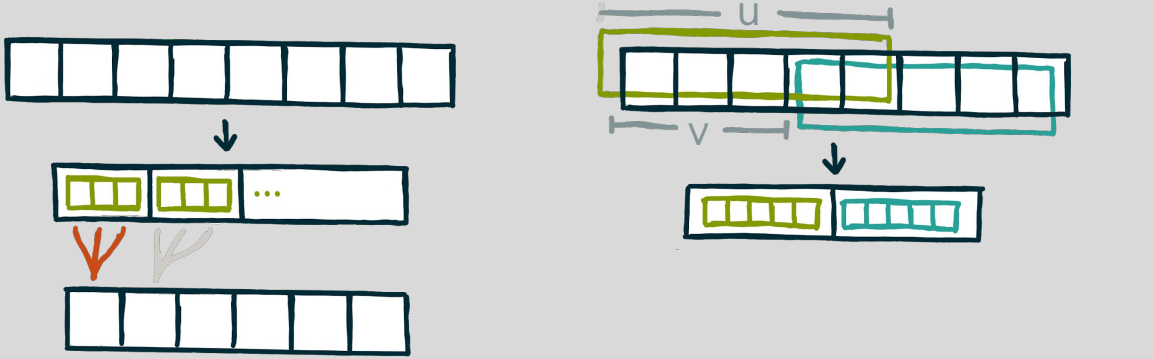

OVERLAPPED TILING AS A REWRITE RULE

overlapped tiling rule

$$\texttt{map}(f,\ \texttt{slide}(3,1,input)) \mapsto \texttt{join}(\texttt{map}(tile \Rightarrow \texttt{map}(f,\ \texttt{slide}(3,1,tile)),\ \texttt{slide}(u,v,input)))$$

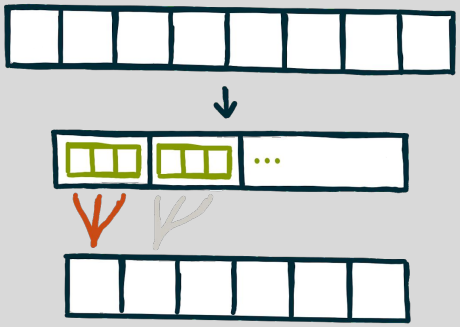

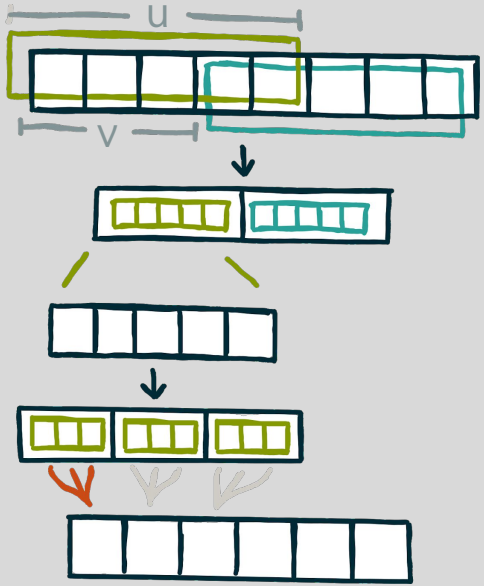

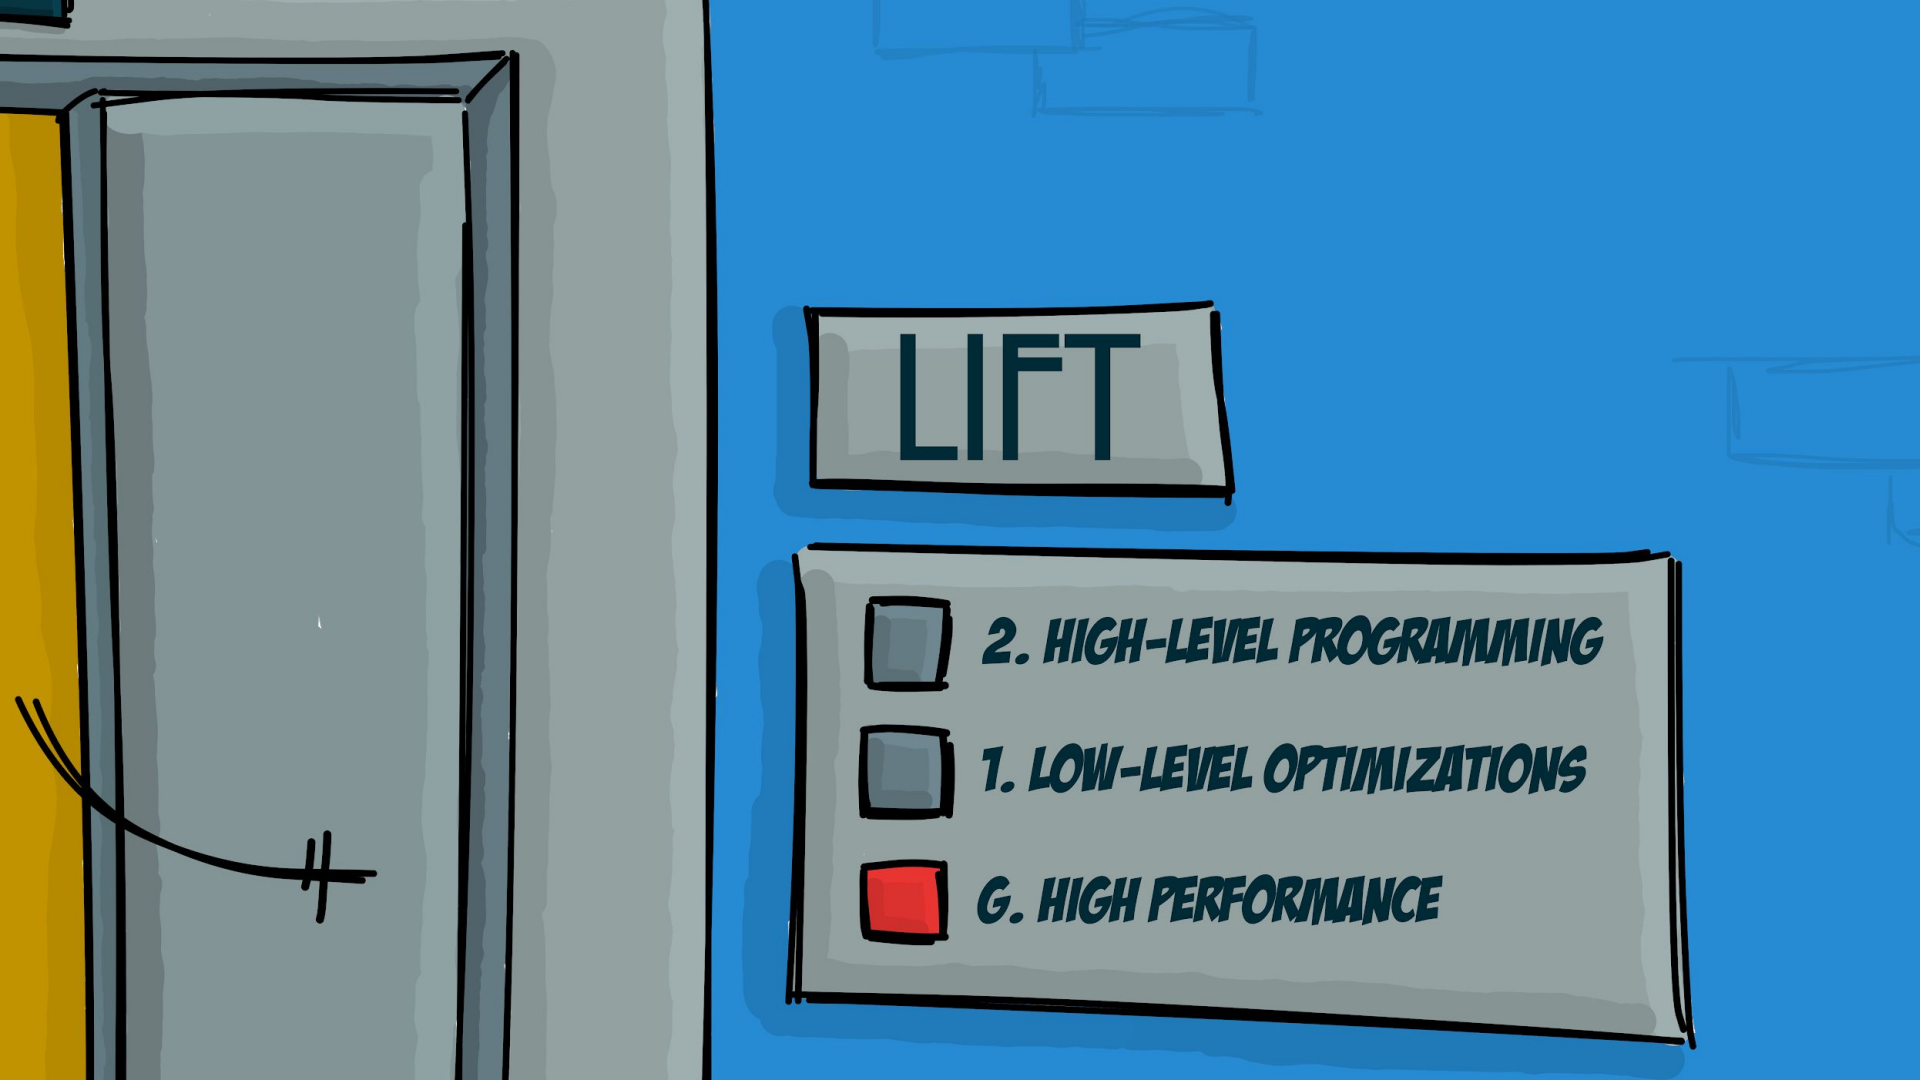
LIFT
2. HIGH-LEVEL PROGRAMMING
1. LOW-LEVEL OPTIMIZATIONS
G. HIGH PERFORMANCE

EXPERIMENTAL EVALUATION

14 Benchmarks

6 hand-optimized
8 polyhedral compilation

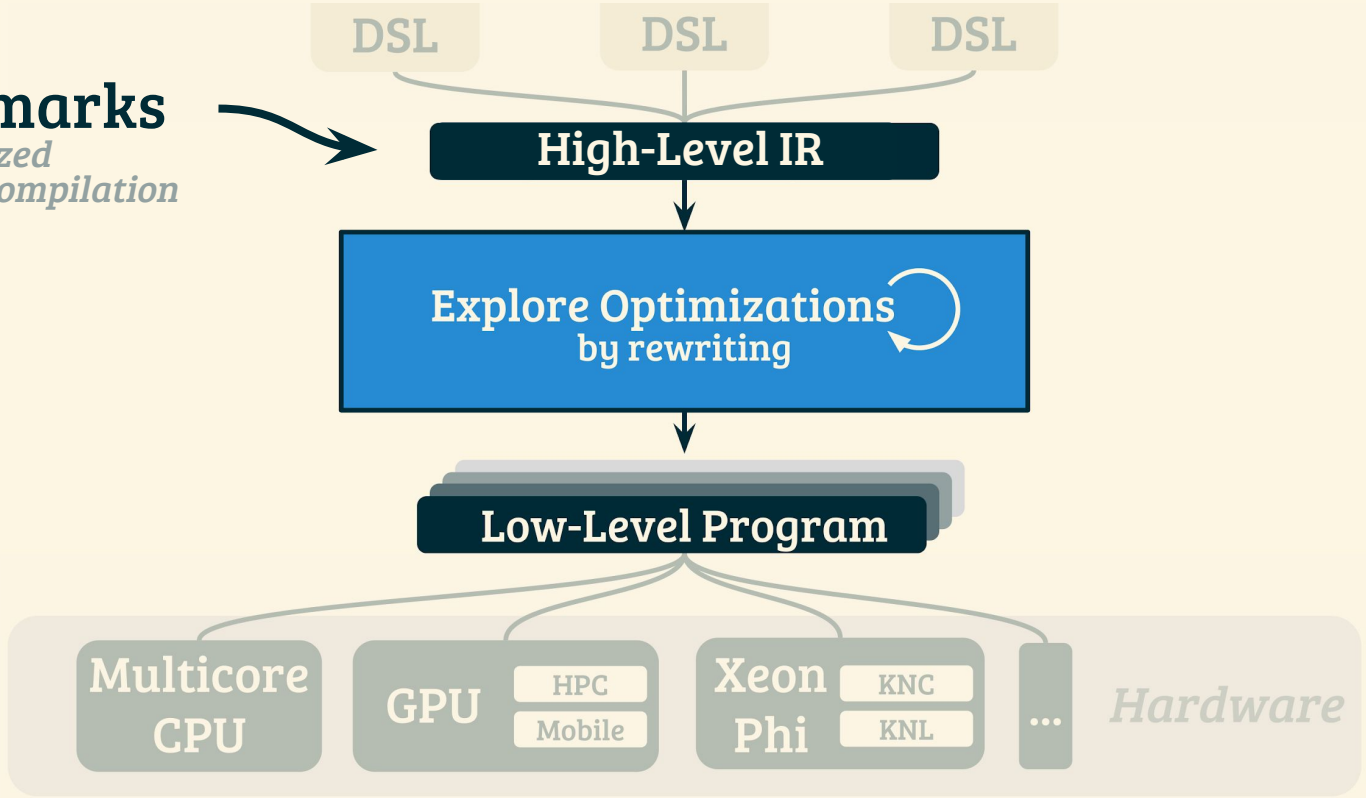

EXPERIMENTAL EVALUATION

14 Benchmarks
6 hand-optimized
8 polyhedral compilation

< 3h Exploration
per benchmark

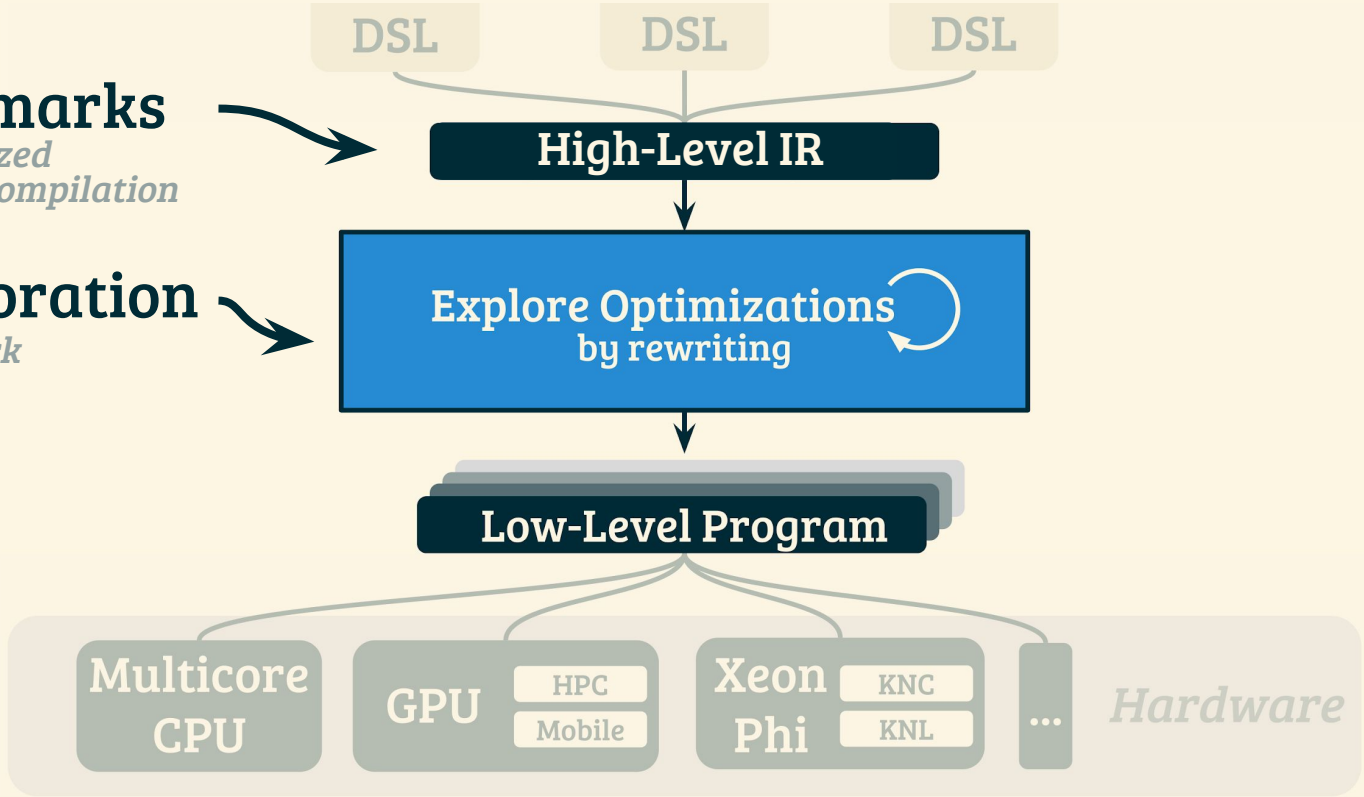

EXPERIMENTAL EVALUATION

14 Benchmarks
6 hand-optimized
8 polyhedral compilation

< 3h Exploration
per benchmark

up to 20 algorithmically different variants
+ auto-tuning of numerical parameters

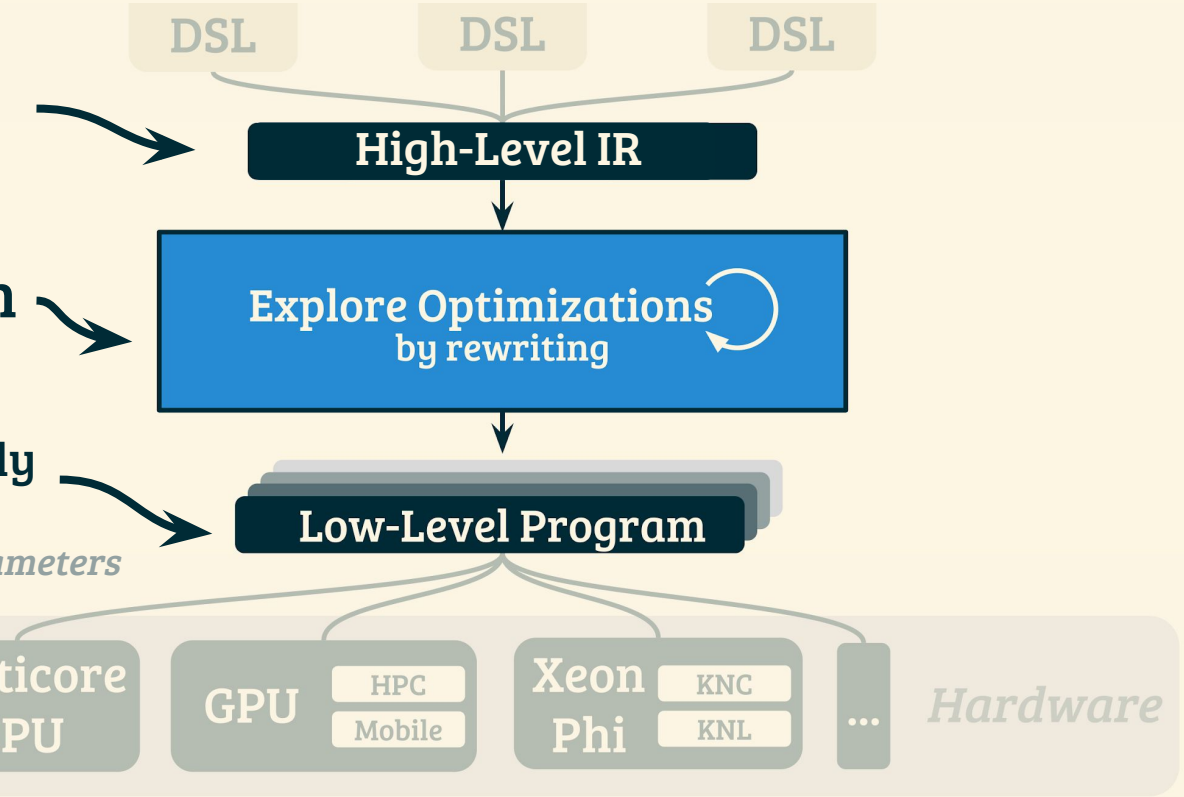

EXPERIMENTAL EVALUATION

14 Benchmarks

6 hand-optimized
8 polyhedral compilation

< 3h Exploration

per benchmark

up to 20 algorithmically different variants

+ auto-tuning of numerical parameters

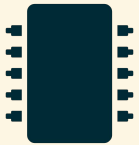

3 GPU Architectures

2 Desktop GPUs
1 Mobile GPU

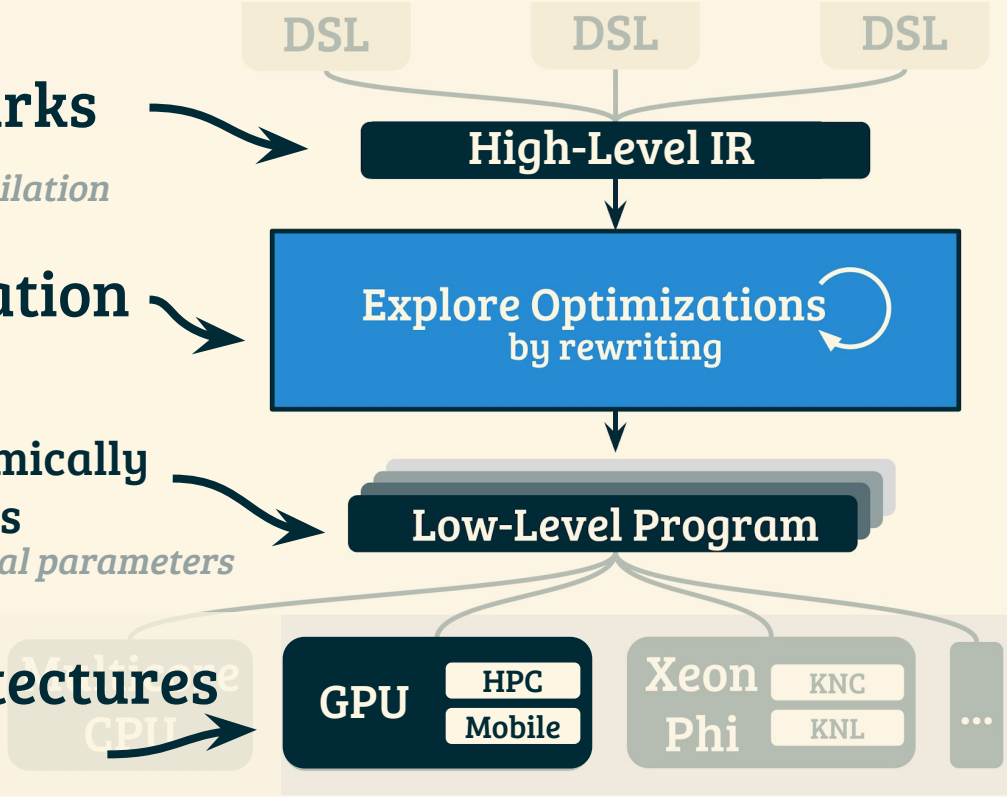

COMPARISON WITH HAND-OPTIMIZED CODES

COMPARISON WITH HAND-OPTIMIZED CODES

higher is better

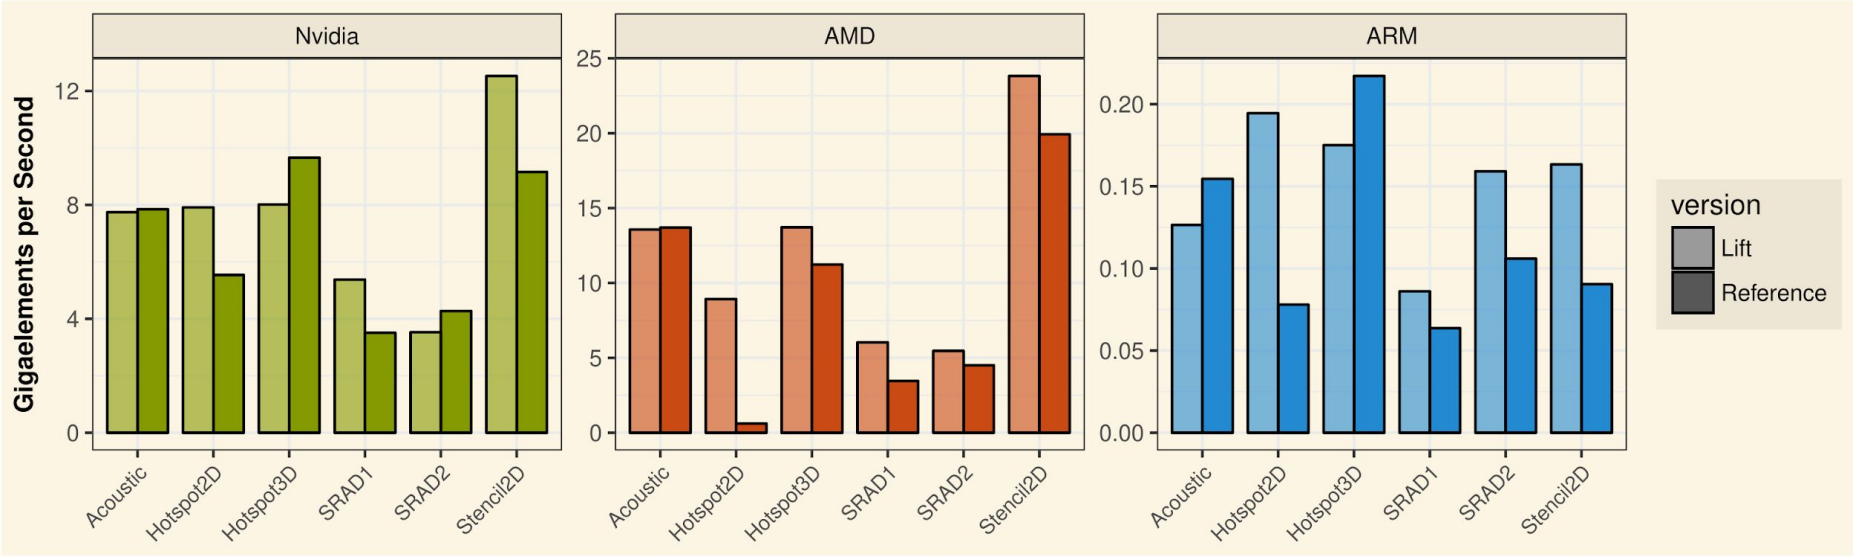

Lift achieves the same performance as hand optimized code

COMPARISON WITH HAND-OPTIMIZED CODES

higher is better

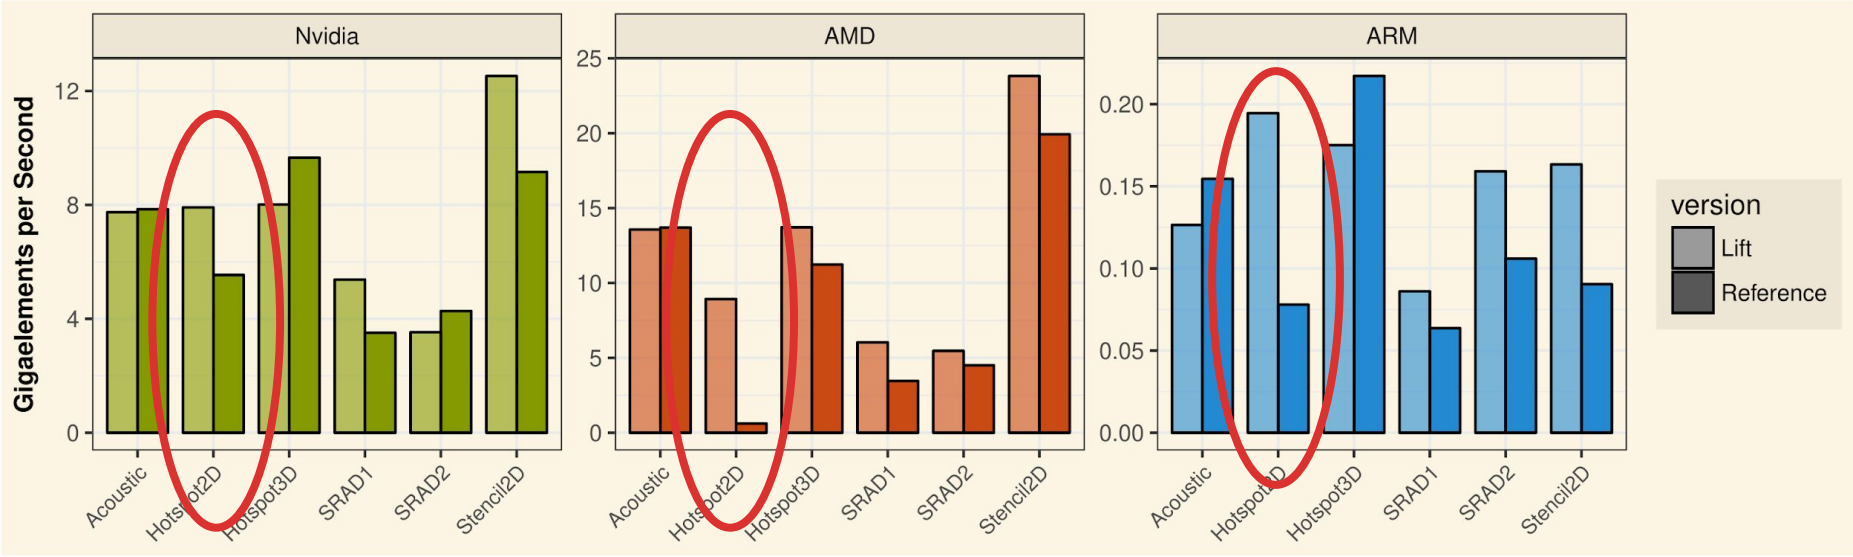

Lift achieves the same performance
as hand optimized code

COMPARISON WITH POLYHEDRAL COMPILATION

higher is better

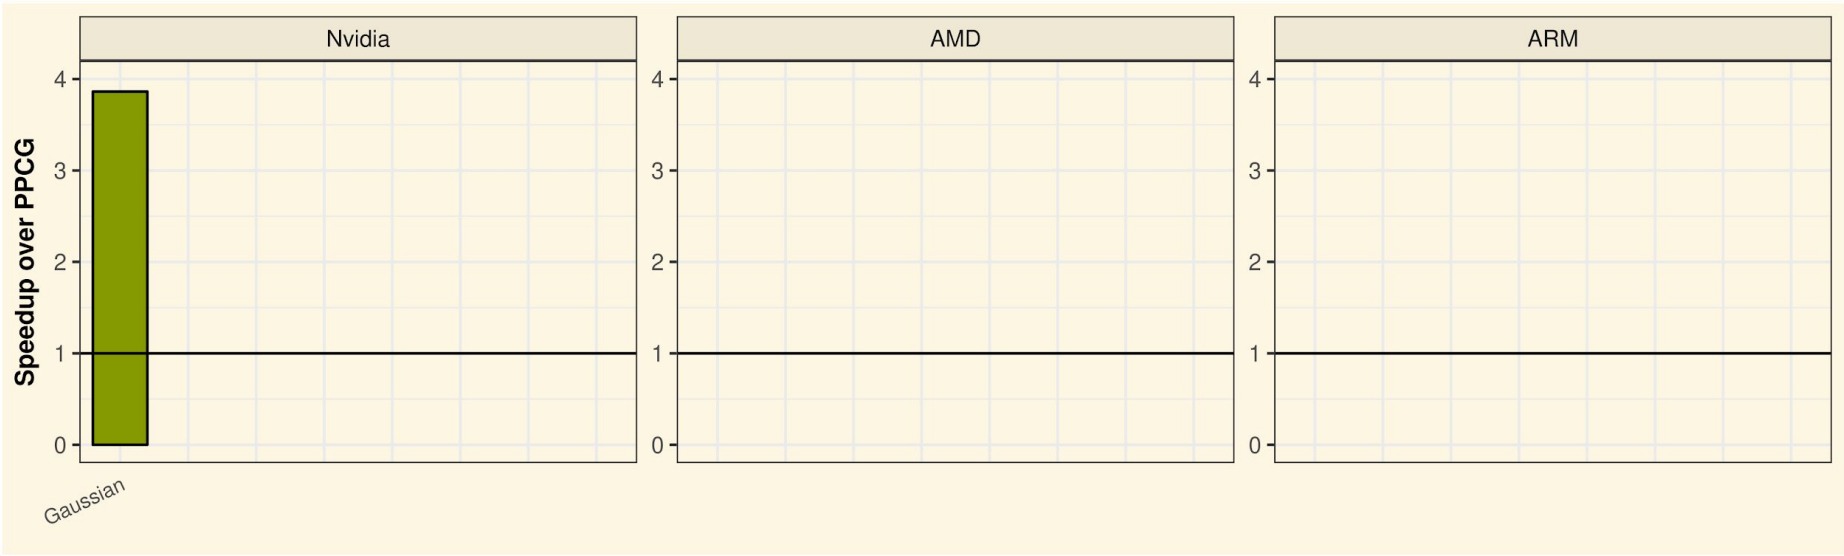

COMPARISON WITH POLYHEDRAL COMPILATION

higher is better

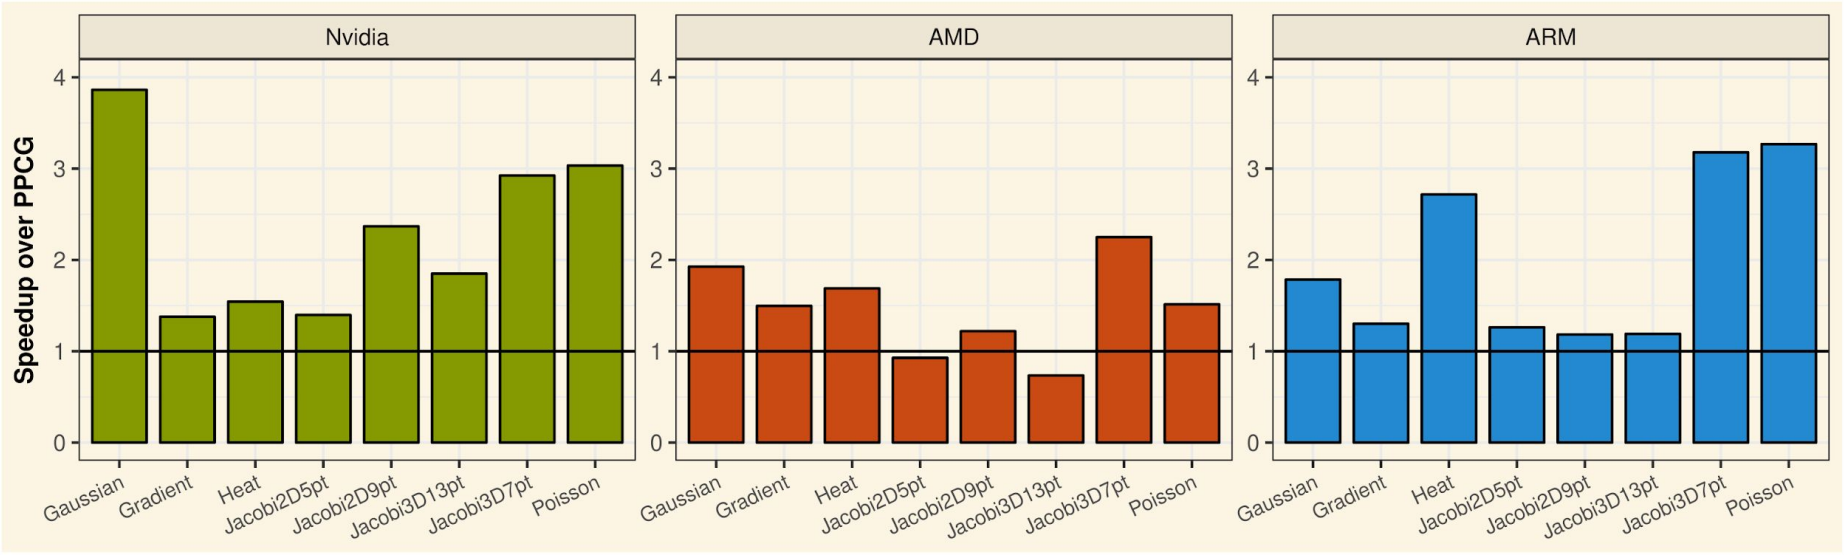

Lift outperforms state-of-the-art optimizing compilers

STENCIL COMPUTATIONS IN LIFT

We added:

2 Primitives
pad, *slide*

1 Rewrite Rule
overlapped tiling

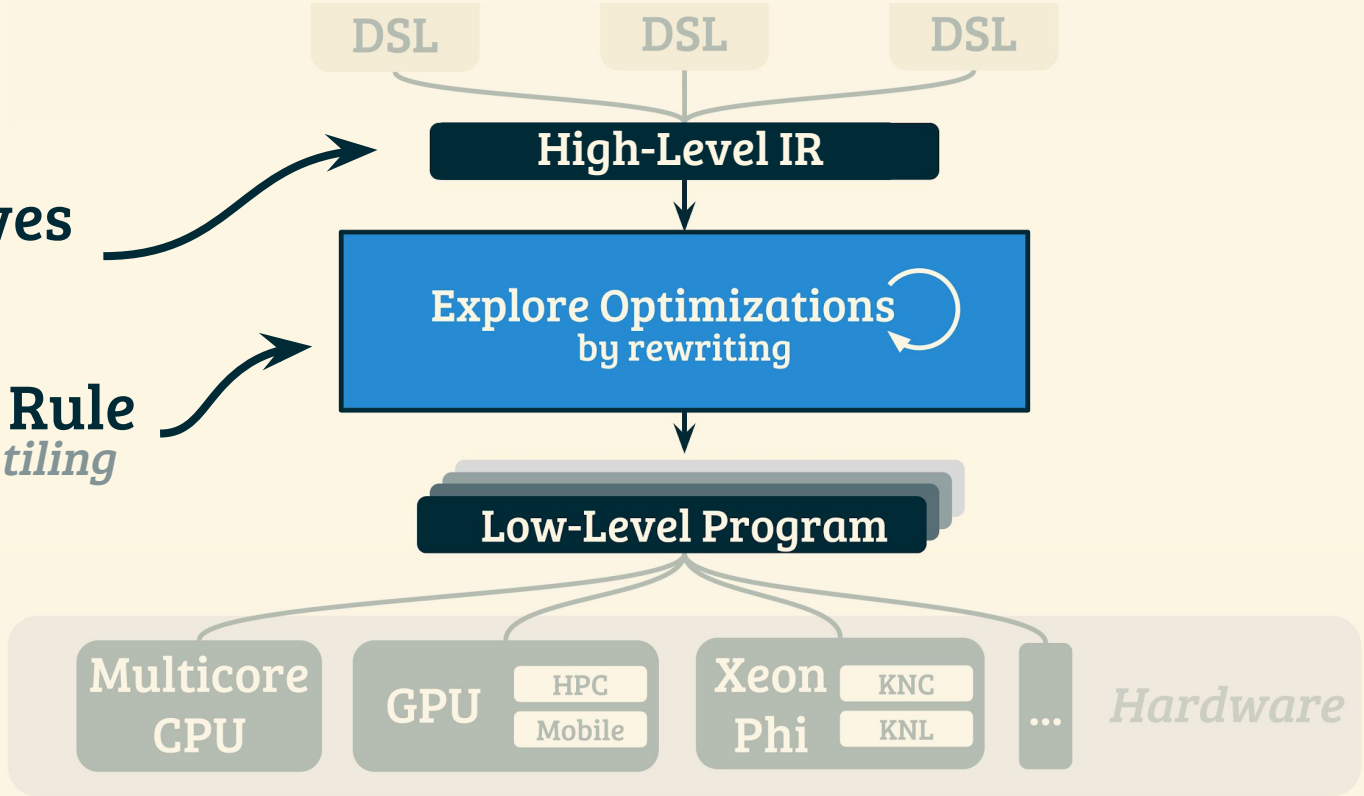

A STEP FURTHER

towards multi-kernel stencil computations

conv

```
def conv =
 fun(A =>
  map2D(multWeights,
   slide2D(3,1,
    pad2D(1,1,clamp,A))))
```

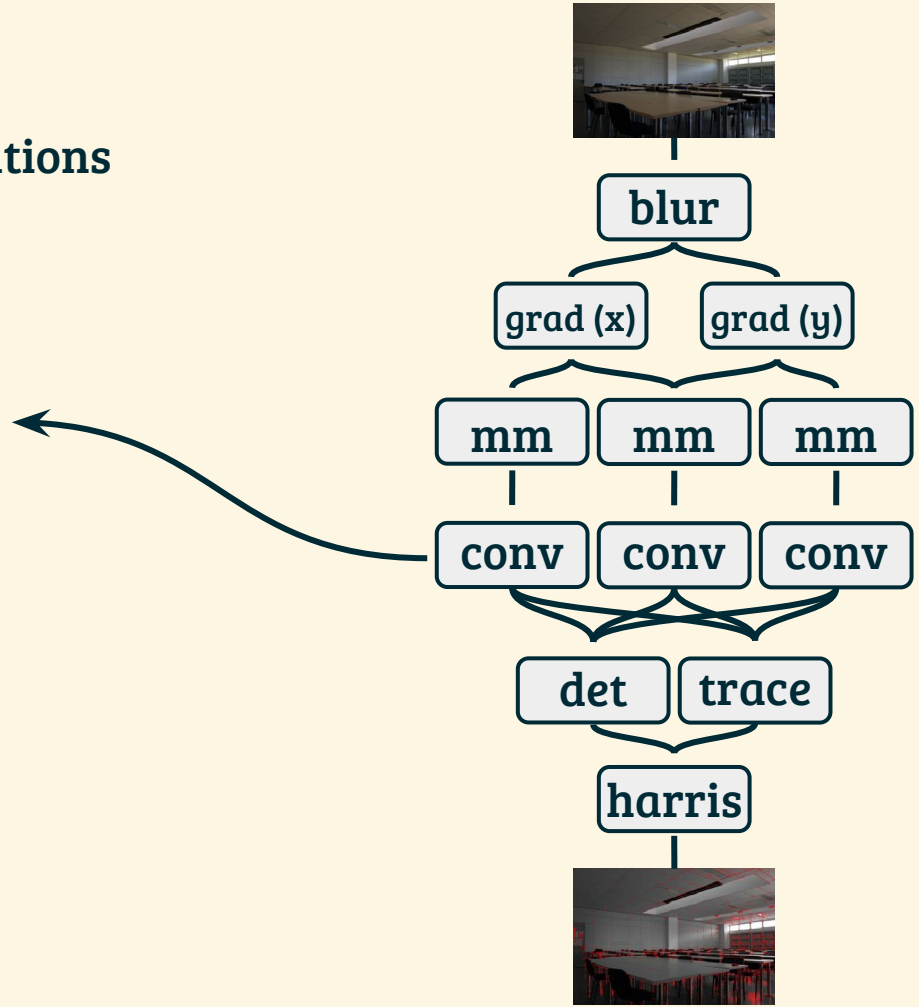

A STEP FURTHER

towards multi-kernel stencil computations

conv

```
def conv =
 fun(A =>
  map2D(multWeights,
   slide2D(3,1,
    pad2D(1,1,clamp,A))))
```

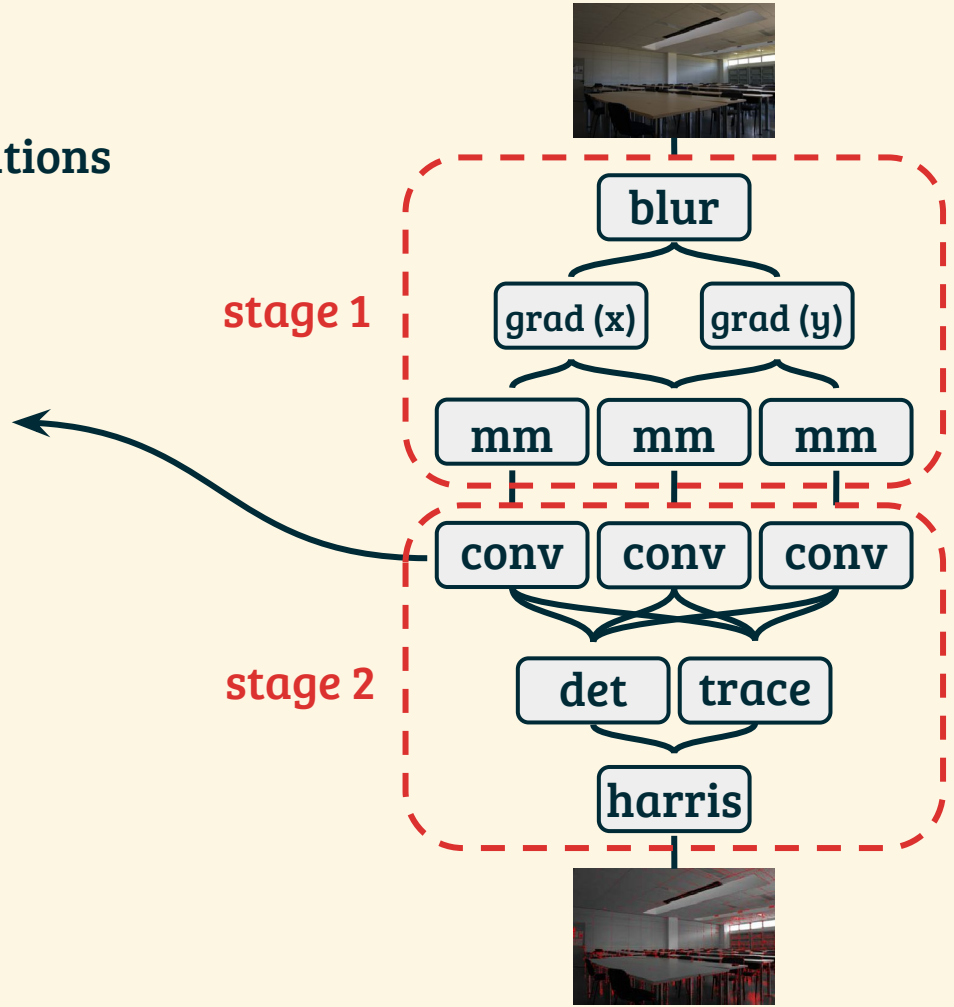

A STEP FURTHER

towards multi-kernel stencil computations

conv

```
def conv =
 fun(A =>
  map2D(multWeights,
   slide2D(3,1,
    pad2D(1,1,clamp,A))))
```

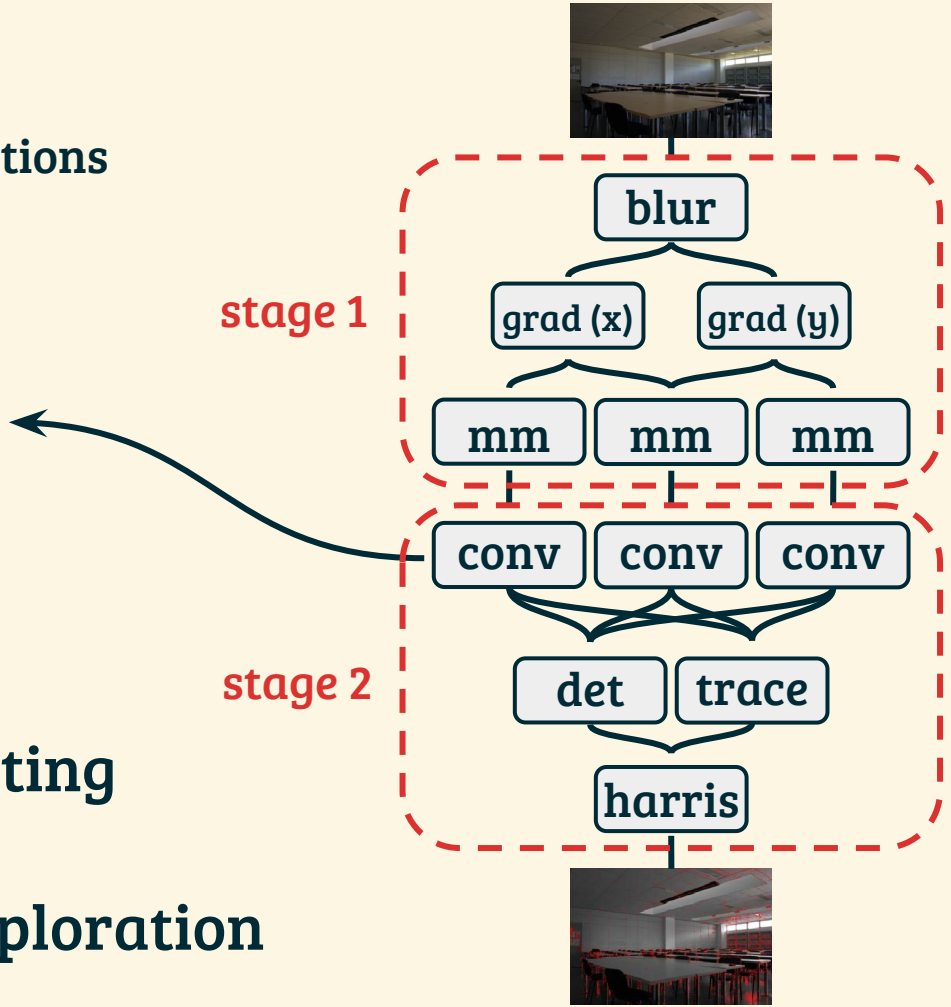

Kernel Fusion via Rewriting

Guided Search Space Exploration

LIFT IS OPEN SOURCE!

more info at:

lift-project.org

Paper 

CGO Artifact  Source Code

 Best Paper Award (CGO'18)

Bastian Hagedorn: b.hagedorn@wwu.de