



WESTFÄLISCHE
WILHELMS-UNIVERSITÄT
MÜNSTER

> Performance Portable Stencil Code Generation in LIFT

Bastian Hagedorn | Larisa Stoltzfus |

Michel Steuer | Sergei Gorlatch | Christophe Dubach




Introduction

Bastian Hagedorn:


- PhD Student:
University of Münster
- Group: Parallel and
Distributed Systems
- Work: Stencil Computations -
LIFT Project



Stencil Computations are all around...




Stencil Computations are all around...



...and are executed on a wide range of hardware...




Stencil Computations are all around...



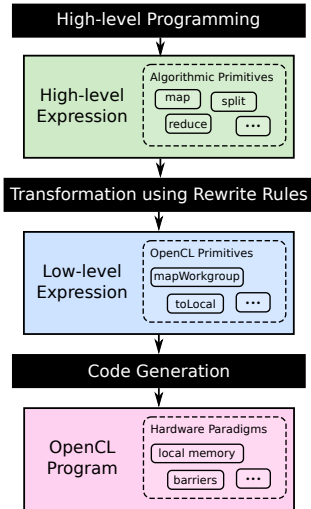
...and are executed on a wide range of hardware...



...which requires experts to achieve high-performance!



The LIFT Project



High-Level LIFT provides high-level interface of composable functional *primitives*


Optimization Automatic rewriting to optimize high-level expressions

Performance High-performance OpenCL Code Generation

So far mainly used for linear algebra applications

Existing Primitives in LIFT


map $\square \rightarrow \square$




reduce \oplus




zip



split n



join



at i




transpose



Existing Primitives in LIFT


map $\square \rightarrow \square$




reduce \oplus




zip



split n



join



at i




transpose



Existing Primitives in LIFT


map $\square \rightarrow \square$




reduce \oplus




zip



split n



join




at i



transpose




**set of primitives is
→ too restricted for
stencil computations**



Decomposing Stencil Computations

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;  
}
```



Decomposing Stencil Computations

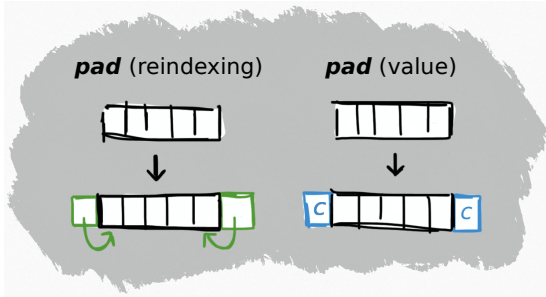
stencil.c

```
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) **Neighborhood:** accessing neighboring elements according to stencil shape
- (b) **Boundary Handling:** what happens at the border of the input array?
- (c) **Stencil Function:** compute single output element for a given neighborhood

(b) Boundary Handling using Pad



reindexing.example

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

$\text{pad}(1, 1, \text{clamp}, [a, b, c, d]) =$
 $[a, a, b, c, d, d]$


value.example

```
constant(i, n) = C
```

$\text{pad}(1, 1, \text{constant}, [a, b, c, d]) =$
 $[C, a, b, c, d, C]$

(a) Create Neighborhoods using Slide

slide (size, step, in)




slide.example

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

$slide : (size : Int, step : Int, in : [T]_n) \rightarrow [[T]_{size}]_{\frac{n-size+step}{step}}$

(c) Apply Stencil Function using Map



stencil-fun.example

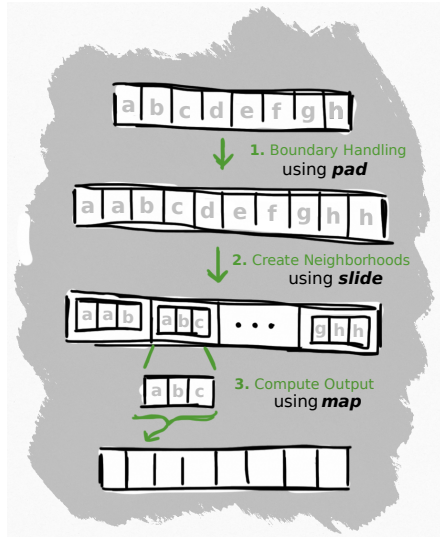
```
map(nbh ⇒  
  reduce(add, 0.0f, nbh),  
  [[0, 1, 2], [1, 2, 3]]) =  
[[3], [6]]
```

Expressing Stencil Computations in LIFT

```
stencil.lift
```

```
val sumNbh = fun(nbh =>
  reduce(add, 0.0f, nbh))


val stencil =
  fun( A: Array(Float, N) =>
    map(sumNbh, // 3.
      slide(3, 1, // 2.
        pad(1, 1, clamp, A)))) // 1.
```



Multidimensional Stencil Computations

Idea: Express complex computations as compositions of simple primitives

$map_n(f, slide_n(size, step, pad_n(l, r, h, input)))$



$map_2(f, slide_2(size, step, pad_2(l, r, h, in)))$

Multidimensional Boundary Handling

$$\mathbf{pad}_1(l, r, h, \mathit{input}) = \mathbf{pad}(l, r, h, \mathit{input})$$




$$\mathbf{pad}_n(l, r, h, \mathit{input}) = \mathbf{map}_{n-1}(\mathbf{pad}(l, r, h), \\ \mathbf{pad}_{n-1}(l, r, h, \mathit{input}))$$

where \mathbf{map}_n are n nested \mathbf{map} s:

$$\mathbf{map}_1(f, \mathit{input}) = \mathbf{map}(f, \mathit{input})$$

$$\mathbf{map}_n(f, \mathit{input}) = \mathbf{map}_{n-1}(\mathbf{map}(f), \mathit{input})$$

$$\mathbf{pad}_2 = \\ \mathbf{map}(\mathbf{pad}^{(1,1h)}, \\ \mathbf{pad}^{(1,1,h,in)})$$



Multidimensional Neighborhood Creation

$$\mathit{slide}_2(2, 1, \begin{bmatrix} [a, b, c], \\ [d, e, f], \\ [g, h, i] \end{bmatrix}) =$$

$\mathit{map}(\mathit{transpose},$

$\mathit{slide}(2, 1,$

$\mathit{map}(\mathit{slide}(2, 1), [[a, b, c], [d, e, f], [g, h, i]]))) =$

Multidimensional Neighborhood Creation

$$\mathit{slide}_2(2, 1, \begin{bmatrix} [a, b, c], \\ [d, e, f], \\ [g, h, i] \end{bmatrix}) =$$

$\mathit{map}(\mathit{transpose},$

$\mathit{slide}(2, 1,$

$\mathit{map}(\mathit{slide}(2, 1), [[a, b, c], [d, e, f], [g, h, i]]))) =$

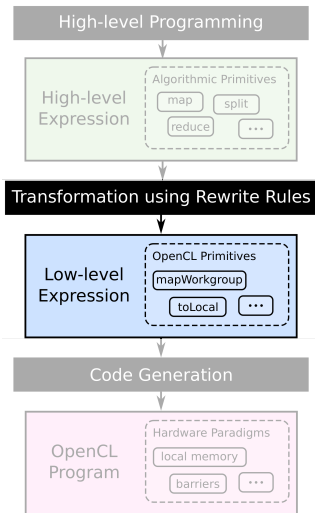
$\mathit{map}(\mathit{transpose}, \mathit{slide}(2, 1,$

$[[[a, b], [b, c]], [[d, e], [e, f]], [[g, h], [h, i]]]))) =$

Multidimensional Neighborhood Creation

$$\begin{aligned} & \mathit{slide}_2(2, 1, \begin{bmatrix} [a, b, c], \\ [d, e, f], \\ [g, h, i] \end{bmatrix}) = \\ & \mathit{map}(\mathit{transpose}, \\ & \quad \mathit{slide}(2, 1, \\ & \quad \quad \mathit{map}(\mathit{slide}(2, 1), \begin{bmatrix} [a, b, c], [d, e, f], [g, h, i] \end{bmatrix}))) = \\ & \mathit{map}(\mathit{transpose}, \mathit{slide}(2, 1, \\ & \quad \begin{bmatrix} [[a, b], [b, c]], [[d, e], [e, f]], [[g, h], [h, i]] \end{bmatrix}))) = \\ & \mathit{map}(\mathit{transpose}, \\ & \quad \begin{bmatrix} [[[a, b], [b, c]], [[d, e], [e, f]]], \\ [[[d, e], [e, f]], [[g, h], [h, i]]] \end{bmatrix}))) = \\ & \quad \begin{bmatrix} \begin{bmatrix} [a, b], \\ [d, e] \end{bmatrix}, \begin{bmatrix} [b, c], \\ [e, f] \end{bmatrix}, \\ \begin{bmatrix} [d, e], \\ [g, h] \end{bmatrix}, \begin{bmatrix} [e, f], \\ [h, i] \end{bmatrix} \end{bmatrix} \end{aligned}$$

Optimizing by Rewriting



Optimizations are encoded as Rewrite Rules:

map-fusion


$$\mathit{map}(f, \mathit{map}(g, in))$$
$$\mapsto$$
$$\mathit{map}((f \circ g), in)$$

divide-and-conquer

$$\mathit{map}(f, in)$$
$$\mapsto$$
$$\mathit{join}(\mathit{map}(\mathit{map}(f), \mathit{split}(n, in)))$$

Optimizing Stencil Computations

Exploiting Locality through Overlapped Tiling:



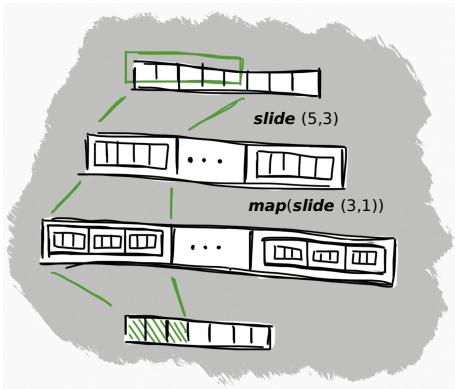
Locality Close neighborhoods share elements that can be grouped in tiles

Local Memory On GPUs, local memory can be used to cache tiles.

Overlap The shape of the stencil determines the overlap at the edges of tiles

Overlapped Tiling in LIFT

We reuse the *slide* primitive to represent overlapped tiles:



$\uparrow u = 5, v = 3, \text{size} = 3, \text{step} = 1$

Tiling as a rewrite rule:

overlapped-tiling

map(*f*, *slide*(*size*, *step*, *input*))


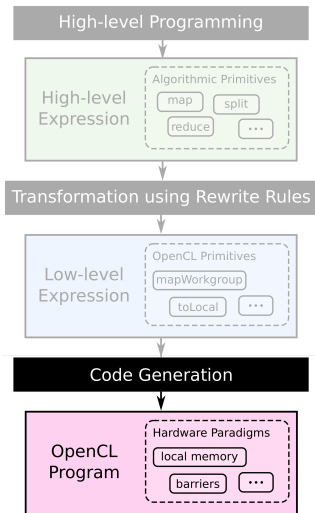
\mapsto

join(*map*(*tile* \Rightarrow

map(*f*, *slide*(*size*, *step*, *tile*)),

slide(*u*, *v*, *input*)))

Code Generation and Exploration

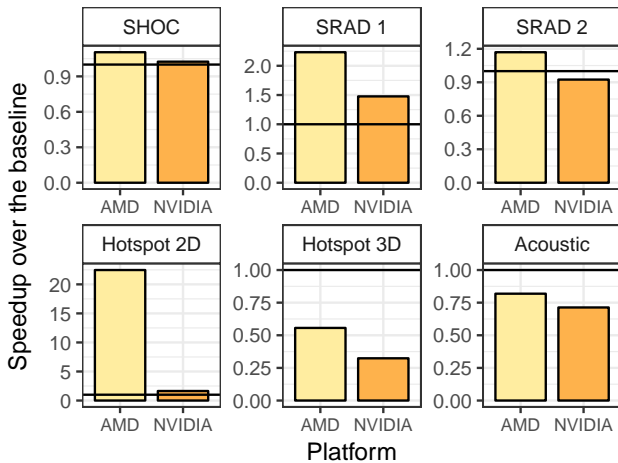


Benchmarks implemented in LIFT

Benchmark	Dim	Stencil	Input size	# Input grids
SHOC	2	9-point	8194^2	1
SRAD 1	2	5-point	504×458	1
SRAD 2	2	3-point	504×458	2
Hotspot2D	2	5-point	8192^2	2
Hotspot3D	3	7-point	$512^2 \times 8$	2
Acoustic	3	7-point	$512^2 \times 404$	2

Table: Benchmarks used in the evaluation.

Evaluation



A horizontal banner image featuring a blue sky with scattered white and light-colored clouds. The word "LIFT" is centered in the sky in a white, sans-serif font.

LIFT


Where High-Level Programming Meets Performance Portability

Questions?

b.hagedorn@wwu.de

<http://www.lift-project.org/>

Non-Rectengular Stencils



```
val T = at(1. at(0, nbh))  
val B = at(1. at(2, nbh))  
val C = at(1. at(1, nbh))  
val L = at(0. at(1, nbh))  
val R = at(2. at(1, nbh))  
return T+B+C+L+R
```

→ stencilFunction

```
map(stencilFunction, nbh)
```

```

// toGlobal(mapGlobal(id)) o
// mapGlobal(reduceSeq (+) 0) o slide 3 1 o pad 1 1 clamp input

float add(float x, float y){
    { return x+y; }
}
float id(float x){
    { return x; }
}
kernel void KERNEL(const global float* restrict v__9,
    global float* v__15, int v_N_0){

    /* Static local memory */
    /* Typed Value memory */
    float v__11;
    /* Private Memory */
    for (int v_gl_id_6 = get_global_id(0); v_gl_id_6 < v_N_0;
        v_gl_id_6 = (v_gl_id_6 + get_global_size(0))) {
        float v_tmp_20 = 0.0f;
        v__11 = v_tmp_20;
        /* reduce_seq */
        for (int v_i_7 = 0; v_i_7 < 3; v_i_7 = (1 + v_i_7)) {
            v__11 = add(v__11,
                v__9[ ((-1 + v_gl_id_6 + v_i_7) >= 0) ?
                    ((-1 + v_gl_id_6 + v_i_7) < v_N_0) ?
                    (-1 + v_gl_id_6 + v_i_7) : (-1 + v_N_0) ) : 0 ] );
        }
        /* end reduce_seq */
    }
    for (int v_gl_id_8 = get_global_id(0); v_gl_id_8 < v_N_0;
        v_gl_id_8 = (v_gl_id_8 + get_global_size(0))) {
        v__15[v_gl_id_8] = id(v__11);
    }
}

```