

# Two Things I Did:

## Parallel Differentiation and Rank Polymorphism

**Robert Schenck**

March 18th, 2025


# Overview

- This talk is about two things I did during my PhD studies:



# Overview


- This talk is about two things I did during my PhD studies:
  - ▶ Thing #1: parallel automatic differentiation:



# Overview

- This talk is about two things I did during my PhD studies:

- ▶ Thing #1: parallel automatic differentiation:




- ▶ Thing #2: (pseudo-)rank polymorphism in a statically-typed language:



# Overview

- This talk is about two things I did during my PhD studies:

- ▶ Thing #1: parallel automatic differentiation:



- ▶ Thing #2: (pseudo-)rank polymorphism in a statically-typed language:




- How are these things related?

# Overview

- This talk is about two things I did during my PhD studies:

- ▶ Thing #1: parallel automatic differentiation:



- ▶ Thing #2: (pseudo-)rank polymorphism in a statically-typed language:




- How are these things related?

- ▶ Well, they are and aren't—both features belong in scientific programming languages.

# Overview

- This talk is about two things I did during my PhD studies:

- ▶ Thing #1: parallel automatic differentiation:



- ▶ Thing #2: (pseudo-)rank polymorphism in a statically-typed language:



- How are these things related?

- ▶ Well, they are and aren't—both features belong in scientific programming languages.
- ▶ Scientific programming is about adopting mathematics for (efficient) computation with a computer.

# Scientific Programming

- Desirables for a scientific programming language:





# Scientific Programming

- Desirables for a scientific programming language:
  - ▶ **High-level of abstraction:**
    - ▶ Should support the writing of mathematics as programs easily.
    - ▶ Program transformations/optimizations should be simple to express.



# Scientific Programming

- Desirables for a scientific programming language:
  - ▶ **High-level of abstraction:**
    - ▶ Should support the writing of mathematics as programs easily.
    - ▶ Program transformations/optimizations should be simple to express.
  - ▶ **Principled:**
    - ▶ Programming model should follow simple rules and be unsurprising.



# Scientific Programming

- Desirables for a scientific programming language:
  - ▶ **High-level of abstraction:**
    - ▶ Should support the writing of mathematics as programs easily.
    - ▶ Program transformations/optimizations should be simple to express.
  - ▶ **Principled:**
    - ▶ Programming model should follow simple rules and be unsurprising.
  - ▶ **Fast:**
    - ▶ Gotta go fast!!



# Scientific Programming

- Desirables for a scientific programming language:
  - ▶ **High-level of abstraction:**
    - ▶ Should support the writing of mathematics as programs easily.
    - ▶ Program transformations/optimizations should be simple to express.
  - ▶ **Principled:**
    - ▶ Programming model should follow simple rules and be unsurprising.
  - ▶ **Fast:**
    - ▶ Gotta go fast!!

	FORTRAN/C	APL	NumPy	?
High-level	X	✓	✓	✓
Principled	X	✓	X	✓
Fast	✓	?	?	✓



# Futhark

- Both things (i.e., Thing #1 and Thing #2) were implemented as extensions to the Futhark programming language.



# Futhark

- Both things (i.e., Thing #1 and Thing #2) were implemented as extensions to the Futhark programming language.
- Futhark is...



# Futhark

- Both things (i.e., Thing #1 and Thing #2) were implemented as extensions to the Futhark programming language.
- Futhark is...
  - ▶ Statically typed (with parametric polymorphism):

$$id : \alpha \rightarrow \alpha = \lambda x. x$$



# Futhark

- Both things (i.e., Thing #1 and Thing #2) were implemented as extensions to the Futhark programming language.
- Futhark is...
  - ▶ Statically typed (with parametric polymorphism):

$$id : \alpha \rightarrow \alpha = \lambda x. x$$

- ▶ A data-parallel functional array language.
    - ▶ Uses a library of parallel operators to build parallel-by-construction programs: **map**, **reduce**, **scan**, **hist**, ...





# Futhark

- Both things (i.e., Thing #1 and Thing #2) were implemented as extensions to the Futhark programming language.
- Futhark is...
  - ▶ Statically typed (with parametric polymorphism):

$$id : \alpha \rightarrow \alpha = \lambda x.x$$

- ▶ A data-parallel functional array language.
  - ▶ Uses a library of parallel operators to build parallel-by-construction programs: **map**, **reduce**, **scan**, **hist**, ...


## Example

```
def dotprod x y = reduce (+) 0 (map (*) x y)
```




# Thing #1: Parallel Automatic Differentiation

Robert Schenck, Ola Rønning, Troels Henriksen, Cosmin E. Oancea




- **Automatic differentiation (AD)** is a program transformation for differentiation.



# Overview

- **Automatic differentiation (AD)** is a program transformation for differentiation.




- Considering AD for a **functional**, **high-level**, and **nested-parallel** array language.



# Overview

- **Automatic differentiation (AD)** is a program transformation for differentiation.



- Considering AD for a **functional, high-level, and nested-parallel** array language.
- All parallelism is made explicit via **parallel combinators**—`map`, `reduce`, `scan`, etc.



# High-level AD

Key idea #1: High-level AD

Parallel constructs are differentiated at a **high-level**.



# High-level AD

## Key idea #1: High-level AD

Parallel constructs are differentiated at a **high-level**.

- Parallel combinators are differentiated with **specialized rewrite rules**.

$$\text{map} \xRightarrow[\text{AD}]{} \text{reduce} \circ \text{map},$$

$$\text{reduce} \xRightarrow[\text{AD}]{} \text{map} \circ \text{scan}$$



## Key idea #1: High-level AD

Parallel constructs are differentiated at a **high-level**.

- Parallel combinators are differentiated with **specialized rewrite rules**.

$$\text{map} \xrightarrow[\text{AD}]{} \text{reduce} \circ \text{map},$$

$$\text{reduce} \xrightarrow[\text{AD}]{} \text{map} \circ \text{scan}$$

- Differentiated programs benefit from entire optimization pipeline in the compiler.





## Key idea #1: High-level AD

Parallel constructs are differentiated at a **high-level**.

- Parallel combinators are differentiated with **specialized rewrite rules**.

$$\text{map} \xrightarrow[\text{AD}]{} \text{reduce} \circ \text{map},$$

$$\text{reduce} \xrightarrow[\text{AD}]{} \text{map} \circ \text{scan}$$

- Differentiated programs benefit from entire optimization pipeline in the compiler.
- Differentiation occurs **before** parallelism is mapped to hardware.



# The Tape

- Variables of the original program appear in the differentiated program.




# The Tape

- Variables of the original program appear in the differentiated program.
- All intermediate variables in the original program must be accessible in the differentiated program.



# The Tape

- Variables of the original program appear in the differentiated program.
- All intermediate variables in the original program must be accessible in the differentiated program.
- In classic AD, these variables are stored on a dynamically allocated **tape**.



# AD by Re-execution

- In our **nested-parallel** context, the tape is **complex/irregular** and must be passed across **deeply nested scopes**. Challenging to implement efficiently.



# AD by Re-execution

- In our **nested-parallel** context, the tape is **complex/irregular** and must be passed across **deeply nested scopes**. Challenging to implement efficiently.

## Key idea #2: Re-execution

Instead of storing intermediate variables, re-compute them by **re-execution**.

- A classic **space-time tradeoff**.



# AD by Re-execution

- In our **nested-parallel** context, the tape is **complex/irregular** and must be passed across **deeply nested scopes**. Challenging to implement efficiently.

## Key idea #2: Re-execution

Instead of storing intermediate variables, re-compute them by **re-execution**.

- A classic **space-time tradeoff**.
- **Asymptotics-preserving**: re-execution overhead is a constant for non-recursive programs.



# AD by Re-execution

- In our **nested-parallel** context, the tape is **complex/irregular** and must be passed across **deeply nested scopes**. Challenging to implement efficiently.

## Key idea #2: Re-execution

Instead of storing intermediate variables, re-compute them by **re-execution**.

- A classic **space-time tradeoff**.
- **Asymptotics-preserving**: re-execution overhead is a constant for non-recursive programs.
- Pretty fast in practice!





## Related Parallel AD Work

- **PyTorch, JAX, etc:** Restricted parallel DSLs; AD on fixed set of array primitives.



## Related Parallel AD Work

- **PyTorch, JAX, etc:** Restricted parallel DSLs; AD on fixed set of array primitives.
- **Enzyme:** LLVM compiler plugin that does AD on a post-optimization, low-level representation.



## Related Parallel AD Work

- **PyTorch, JAX, etc:** Restricted parallel DSLs; AD on fixed set of array primitives.
- **Enzyme:** LLVM compiler plugin that does AD on a post-optimization, low-level representation.
- **Dex:** High-level AD that uses multiple tapes; hard to implement efficiently.



# Related Parallel AD Work

- **PyTorch, JAX, etc:** Restricted parallel DSLs; AD on fixed set of array primitives.
- **Enzyme:** LLVM compiler plugin that does AD on a post-optimization, low-level representation.
- **Dex:** High-level AD that uses multiple tapes; hard to implement efficiently.

	PyTorch, JAX, etc.	Enzyme	Dex	?
High-level	✓	✗	✓	✓
Principled	✗	✗	✓	✓
Fast	?	✓	?	✓



# A Very Short Introduction to AD



# Introduction to AD

- **Goal:** compute the sensitivity of the output  $y$  to its inputs  $x_0, x_1$ .

$P(x_0, x_1) :$   
 $t_0 = \sin(x_0)$   
 $t_1 = x_1 \cdot t_0$   
 $y = x_0 + t_1$   
**return**  $y$

$\implies$

$P'(x_0, x_1, \bar{y}) :$   
?





# Introduction to AD

```
P(x0, x1) :  
  t0 = sin(x0)  
  t1 = x1 · t0  
  y = x0 + t1  
  return y
```

⇒

```
P'(x0, x1,  $\bar{y}$ ) :  
  ?
```

```
return  $\bar{x}_0$ ,  $\bar{x}_1$ 
```

- **Goal:** compute  $\bar{x}_0$  and  $\bar{x}_1$

## Adjoint of a variable

$$\bar{v} \equiv \frac{\partial y}{\partial v}$$

The **sensitivity** of the output  $y$  to  $v$ .





# Introduction to AD

$P(x_0, x_1) :$   
 $t_0 = \sin(x_0)$   
 $t_1 = x_1 \cdot t_0$   
 $y = x_0 + t_1$   
**return**  $y$

$\implies$

$P'(x_0, x_1, \bar{y}) :$   
 $t_0 = \sin(x_0)$   
 $t_1 = x_1 \cdot t_0$   
 $y = x_0 + t_1$

**return**  $\bar{x}_0, \bar{x}_1$

- **Goal:** compute  $\bar{x}_0$  and  $\bar{x}_1$

## Adjoint of a variable

$$\bar{v} \equiv \frac{\partial y}{\partial v}$$

The **sensitivity** of the output  $y$  to  $v$ .

- Adjoints depend on primal values. **Add the statements of the original program.**



# Introduction to AD

```
P(x0, x1) :  
  t0 = sin(x0)  
  t1 = x1 · t0  
  y = x0 + t1  
  return y
```

⇒

```
P'(x0, x1,  $\bar{y}$ ) :  
  t0 = sin(x0)  
  t1 = x1 · t0  
  y = x0 + t1
```

return  $\bar{x}_0$ ,  $\bar{x}_1$

- **Goal:** compute  $\bar{x}_0$  and  $\bar{x}_1$

## Adjoint of a variable

$$\bar{v} \equiv \frac{\partial y}{\partial v}$$

The **sensitivity** of the output  $y$  to  $v$ .

- Adjoint values depend on primal values. Add the statements of the original program.
- Compute  $\bar{x}_1$  by the chain rule:

$$\bar{x}_1 \equiv \frac{\partial y}{\partial x_1} = \frac{\partial y}{\partial t_1} \frac{\partial t_1}{\partial x_1} = \bar{t}_1 \frac{\partial t_1}{\partial x_1}$$



# Introduction to AD

$P(x_0, x_1)$  :  
 $t_0 = \sin(x_0)$   
 $t_1 = x_1 \cdot t_0$   
 $y = x_0 + t_1$   
**return**  $y$

$\implies$

$P'(x_0, x_1, \bar{y})$  :  
 $t_0 = \sin(x_0)$   
 $t_1 = x_1 \cdot t_0$   
 $y = x_0 + t_1$

$$\bar{t}_1 = \bar{y}$$
$$\bar{x}_1 = t_0 \cdot \bar{t}_1$$

**return**  $\bar{x}_0, \bar{x}_1$

- **Goal:** compute  $\bar{x}_0$  and  $\bar{x}_1$

## Adjoint of a variable

$$\bar{v} \equiv \frac{\partial y}{\partial v}$$

The **sensitivity** of the output  $y$  to  $v$ .

- Adjoints depend on primal values. Add the statements of the original program.
- Compute  $\bar{x}_1$  by the chain rule:

$$\bar{x}_1 \equiv \frac{\partial y}{\partial x_1} = \frac{\partial y}{\partial t_1} \frac{\partial t_1}{\partial x_1} = \bar{t}_1 \frac{\partial t_1}{\partial x_1}$$



# Introduction to AD

$P(x_0, x_1) :$   
 $t_0 = \sin(x_0)$   
 $t_1 = x_1 \cdot t_0$   
 $y = x_0 + t_1$   
**return**  $y$

$\implies$

$P'(x_0, x_1, \bar{y}) :$   
 $t_0 = \sin(x_0)$   
 $t_1 = x_1 \cdot t_0$   
 $y = x_0 + t_1$   
 $\bar{x}_0 = \bar{y}$   
 $\bar{t}_1 = \bar{y}$   
 $\bar{x}_1 = t_0 \cdot \bar{t}_1$   
 $\bar{t}_0 = x_1 \cdot \bar{t}_1$   
 $\bar{x}_0 += \cos(x_0) \cdot \bar{t}_0$   
**return**  $\bar{x}_0, \bar{x}_1$

- **Goal:** compute  $\bar{x}_0$  and  $\bar{x}_1$

## Adjoint of a variable

$$\bar{v} \equiv \frac{\partial y}{\partial v}$$

The **sensitivity** of the output  $y$  to  $v$ .

- Adjoints depend on primal values. Add the statements of the original program.
- Compute  $\bar{x}_1$  by the chain rule.
- Do the same for  $\bar{x}_0$ .



# Introduction to AD

$P(x_0, x_1) :$   
 $t_0 = \sin(x_0)$   
 $t_1 = x_1 \cdot t_0$   
 $y = x_0 + t_1$   
**return**  $y$

$\implies$

$P'(x_0, x_1, \bar{y}) :$   
 $t_0 = \sin(x_0)$   
 $t_1 = x_1 \cdot t_0$   
 $y = x_0 + t_1$   
 $\bar{x}_0 = \bar{y}$   
 $\bar{t}_1 = \bar{y}$   
 $\bar{x}_1 = t_0 \cdot \bar{t}_1$   
 $\bar{t}_0 = x_1 \cdot \bar{t}_1$   
 $\bar{x}_0 += \cos(x_0) \cdot \bar{t}_0$   
**return**  $\bar{x}_0, \bar{x}_1$

- **Goal:** compute  $\bar{x}_0$  and  $\bar{x}_1$

## Adjoint of a variable

$$\bar{v} \equiv \frac{\partial y}{\partial v}$$

The **sensitivity** of the output  $y$  to  $v$ .

- Adjoints depend on primal values. Add the statements of the original program.
- Compute  $\bar{x}_1$  by the chain rule.
- Do the same for  $\bar{x}_0$ .
- Since  $\bar{x}_0$  is read **twice**, its adjoint gets **two** contributions.



# Introduction to AD

$P(x_0, x_1) :$   
 $t_0 = \sin(x_0)$   
 $t_1 = x_1 \cdot t_0$   
 $y = x_0 + t_1$   
**return**  $y$

$\implies$

$P'(x_0, x_1, \bar{y}) :$   
 $t_0 = \sin(x_0)$   
 $t_1 = x_1 \cdot t_0$   
 $y = x_0 + t_1$   
 $\bar{x}_0 = \bar{y}$   
 $\bar{t}_1 = \bar{y}$   
 $\bar{x}_1 = t_0 \cdot \bar{t}_1$   
 $\bar{t}_0 = x_1 \cdot \bar{t}_1$   
 $\bar{x}_0 += \cos(x_0) \cdot \bar{t}_0$   
**return**  $\bar{x}_0, \bar{x}_1$



- **Goal:** compute  $\bar{x}_0$  and  $\bar{x}_1$

## Adjoint of a variable

$$\bar{v} \equiv \frac{\partial y}{\partial v}$$

The **sensitivity** of the output  $y$  to  $v$ .

- Adjoint values depend on primal values. Add the statements of the original program.
- Compute  $\bar{x}_1$  by the chain rule.
- Do the same for  $\bar{x}_0$ .
- Since  $\bar{x}_0$  is read twice, its adjoint gets two contributions.
- Adjoint values appear in **reverse** program-order.

# Introduction to AD

$$\begin{array}{l} P(x_0, x_1) : \\ t_0 = \sin(x_0) \\ t_1 = x_1 \cdot t_0 \\ y = x_0 + t_1 \\ \mathbf{return} \ y \end{array} \quad \Longrightarrow \quad \begin{array}{l} P'(x_0, x_1, \bar{y}) : \\ t_0 = \sin(x_0) \\ t_1 = x_1 \cdot t_0 \\ y = x_0 + t_1 \\ \bar{x}_0 = \bar{y} \\ \bar{t}_1 = \bar{y} \\ \bar{x}_1 = t_0 \cdot \bar{t}_1 \\ \bar{t}_0 = x_1 \cdot \bar{t}_1 \\ \bar{x}_0 += \cos(x_0) \cdot \bar{t}_0 \\ \mathbf{return} \ \bar{x}_0, \bar{x}_1 \end{array}$$

- Can express AD as a rewrite rule:

## AD rewrite rule

$$\begin{array}{l} v = f(u, w) \\ \vdots \\ v = f(u, w) \Longrightarrow \begin{array}{l} \bar{u} += \frac{\partial f(u, w)}{\partial u} \bar{v} \\ \bar{w} += \frac{\partial f(u, w)}{\partial w} \bar{v} \end{array} \end{array}$$



# AD Transformation





# AD Transformation

- Programs are built from **bodies**; i.e., a list of statements concluding in a result.


```
let  $x = a + b$  } stm ) )  
let  $res = x * c$  } stms  
in  $res$  } body
```





# AD Transformation

- Programs are built from **bodies**; i.e., a list of statements concluding in a result.



- To differentiate:
  - Execute the statements of the original body;  $\overrightarrow{\text{stms}}$  is the **forward sweep**.
  - Compute the adjoint contributions;  $\overleftarrow{\text{stms}}$  is the **reverse sweep**.





# AD by Re-execution

```
let zS = map ( $\lambda a bs \rightarrow$   
  let z = reduce ( $\lambda x y \rightarrow$   
    let t = sin(x)  
    let red_res = t · y  
    in red_res) 0 bs  
  let map_res = z · a  
  in map_res) as bss  
in zS
```


}  
*stms<sub>0</sub>*  
*stms<sub>1</sub>*  
*stms<sub>2</sub>*




# AD by Re-execution



# AD by Re-execution



# AD by Re-execution



- The amount of re-execution is proportional to the depth of the deepest scope.





# Re-execution in Perfect Scope Nests

$stms_0$   
 $stms_1$   
 $stms_2$

**General case**

$map\ f$   
 $map\ g$   
 $stms_2$

**Perfect nest**

$\Rightarrow$

$\overrightarrow{stms_0}$   
 $\overrightarrow{stms_1}$   
 $\overrightarrow{stms_2}$   
 $\overleftarrow{stms_0}$   
 $\overleftarrow{stms_1}$   
 $\overleftarrow{stms_2}$   
 $\overrightarrow{stms_2}$   
 $\overleftarrow{stms_1}$   
 $\overleftarrow{stms_2}$

**General case**

$\overrightarrow{map\ f}$   
 $\overrightarrow{map\ g}$   
 $\overrightarrow{stms_2}$   
 $\overleftarrow{map\ f}$   
 $\overleftarrow{map\ g}$   
 $\overrightarrow{stms_2}$   
 $\overleftarrow{stms_2}$   
 $\overleftarrow{stms_2}$

**Perfect nest**

- In perfect scope nests, only the **outermost** and **innermost** scopes are re-executed.



# Differentiating Parallel Constructs



# reduce

- **reduce** combines all elements of an array with a binary associative operator  $\odot$ :

$$\mathbf{let\ } y = \mathbf{reduce\ } \odot\ e_{\odot}\ [a_0, a_1, \dots, a_{n-1}]$$

$\equiv$

$$\mathbf{let\ } y = a_0 \odot a_1 \odot \dots \odot a_{n-1}$$



# reduce

- **reduce** combines all elements of an array with a binary associative operator  $\odot$ :

$$\mathbf{let\ } y = \mathbf{reduce\ } \odot\ e_{\odot}\ [a_0, a_1, \dots, a_{n-1}]$$

$\equiv$

$$\mathbf{let\ } y = a_0 \odot a_1 \odot \dots \odot a_{n-1}$$

- For each  $a_i$  in the array, we can group the terms of the reduce as

$$\underbrace{a_0 \odot \dots \odot a_{i-1}}_{l_i} \odot a_i \odot \underbrace{a_{i+1} \odot \dots \odot a_{n-1}}_{r_i}$$



# reduce

- **reduce** combines all elements of an array with a binary associative operator  $\odot$ :

$$\begin{aligned}\mathbf{let } y &= \mathbf{reduce } \odot e_{\odot} [a_0, a_1, \dots, a_{n-1}] \\ &\equiv \\ \mathbf{let } y &= a_0 \odot a_1 \odot \dots \odot a_{n-1}\end{aligned}$$

- For each  $a_i$  in the array, we can group the terms of the reduce as

$$\underbrace{a_0 \odot \dots \odot a_{i-1}}_{l_i} \odot a_i \odot \underbrace{a_{i+1} \odot \dots \odot a_{n-1}}_{r_i}$$

And then directly apply the AD rewrite rule

$$\bar{a}_i \bar{\vdash} = \frac{\partial(l_i \odot a_i \odot r_i)}{\partial a_i} \bar{y}$$



## Computing $l_i$ and $r_i$

- For each  $i \in \{0, \dots, n-1\}$ , need to compute  $l_i$  and  $r_i$

$$\underbrace{a_0 \odot \cdots \odot a_{i-1}}_{l_i} \odot a_i \odot \underbrace{a_{i+1} \odot \cdots \odot a_{n-1}}_{r_i}$$



## Computing $l_i$ and $r_i$

- For each  $i \in \{0, \dots, n-1\}$ , need to compute  $l_i$  and  $r_i$

$$\underbrace{a_0 \odot \dots \odot a_{i-1}}_{l_i} \odot a_i \odot \underbrace{a_{i+1} \odot \dots \odot a_{n-1}}_{r_i}$$

- For the  $l_i$ s, do a parallel scan

$$\mathbf{let} \ ls = \mathbf{scan} \odot e_\odot [a_0, a_1, \dots, a_{n-1}] \equiv \left[ \underbrace{e_\odot}_{l_0}, \underbrace{a_0}_{l_1}, \underbrace{a_0 \odot a_1}_{l_2}, \dots, \underbrace{a_0 \odot \dots \odot a_{n-2}}_{l_{n-1}} \right]$$



# Computing $l_i$ and $r_i$

- For each  $i \in \{0, \dots, n-1\}$ , need to compute  $l_i$  and  $r_i$

$$\underbrace{a_0 \odot \dots \odot a_{i-1}}_{l_i} \odot a_i \odot \underbrace{a_{i+1} \odot \dots \odot a_{n-1}}_{r_i}$$

- For the  $l_i$ s, do a parallel scan

$$\text{let } ls = \text{scan } \odot e_{\odot} [a_0, a_1, \dots, a_{n-1}] \equiv [ \underbrace{e_{\odot}}_{l_0}, \underbrace{a_0}_{l_1}, \underbrace{a_0 \odot a_1}_{l_2}, \dots, \underbrace{a_0 \odot \dots \odot a_{n-2}}_{l_{n-1}} ]$$

- For the  $r_i$ s, the array must be reversed

$$\text{let } rs = \text{reverse } as \triangleright \text{scan } (\lambda x y \rightarrow y \odot x) e_{\odot} [a_0, a_1, \dots, a_{n-1}] \triangleright \text{reverse} \\ \equiv [ \underbrace{a_0 \odot \dots \odot a_{n-2}}_{r_0}, \dots, \underbrace{a_{n-2} \odot a_{n-1}}_{r_{n-3}}, \underbrace{a_{n-1}}_{r_{n-2}}, \underbrace{e_{\odot}}_{r_{n-1}} ]$$





# reduce

- The differentiation of reduce results in the following statements:

```
let  $y = \text{reduce } \odot e_{\odot} [a_0, a_1, \dots, a_{n-1}]$  } Forward sweep
    ⋮
let  $ls = \text{scan } \odot e_{\odot} as$ 
let  $rs = \text{reverse } as \triangleright \text{scan } (\lambda x y \rightarrow y \odot x) e_{\odot} \triangleright \text{reverse}$  } Reverse sweep
let  $\overline{as} \overline{r} = \text{map } \left( \lambda l_i a_i r_i \rightarrow \frac{\partial (l_i \odot a_i \odot r_i)}{\partial a_i} \overline{y} \right) ls as rs$ 
```



# reduce

- The differentiation of reduce results in the following statements:

```
let y = reduce ⊙ e ⊙ [a0, a1, ..., an-1] } Forward sweep
      ⋮
let ls = scan ⊙ e ⊙ as
let rs = reverse as ▷ scan (λx y → y ⊙ x) e ⊙ ▷ reverse } Reverse sweep
let  $\overline{as}$   $\overline{r}$  = map  $\left( \lambda l_i a_i r_i \rightarrow \frac{\partial(l_i \odot a_i \odot r_i)}{\partial a_i} \overline{y} \right)$  ls as rs
```

- The rule is **asymptotics-preserving**: scan has the same asymptotics as reduce.



# reduce

- The differentiation of reduce results in the following statements:

$\mathbf{let} \ y = \mathbf{reduce} \ \odot \ e_{\odot} \ [a_0, a_1, \dots, a_{n-1}]$	}	Forward sweep
$\vdots$		
$\mathbf{let} \ ls = \mathbf{scan} \ \odot \ e_{\odot} \ as$	}	Reverse sweep
$\mathbf{let} \ rs = \mathbf{reverse} \ as \triangleright \ \mathbf{scan} \ (\lambda x \ y \rightarrow y \ \odot \ x) \ e_{\odot} \ \triangleright \ \mathbf{reverse}$		
$\mathbf{let} \ \overline{as} \ \overline{r} = \mathbf{map} \ \left( \lambda l_i \ a_i \ r_i \rightarrow \frac{\partial(l_i \odot a_i \odot r_i)}{\partial a_i} \ \overline{y} \right) \ ls \ as \ rs$		

- The rule is **asymptotics-preserving**: scan has the same asymptotics as reduce.
- **Specialized rules** for other operators (+, min, max, ·) admit even more efficient implementations.



- Consider the following **map** :

```
let xs = map ( $\lambda a\ b \rightarrow$  let res = a · b in res) as bs
```



- Consider the following **map** :

$$\mathbf{let} \ xs = \mathbf{map} \ (\lambda a \ b \rightarrow \mathbf{let} \ res = a \cdot b \ \mathbf{in} \ res) \ as \ bs$$

- Differentiating **map** is straightforward: just differentiate the lambda and pass in the necessary adjoints as well:

$$\begin{aligned} \mathbf{let} \ \overline{as}, \overline{bs} = \mathbf{map} \ (\lambda a \ b \ \overline{x} \ \overline{a_0} \ \overline{b_0} \rightarrow \\ \mathbf{let} \ res = a \cdot b \\ \mathbf{let} \ \overline{a} = b \cdot \overline{x} + \overline{a_0} \\ \mathbf{let} \ \overline{b} = a \cdot \overline{x} + \overline{b_0} \\ \mathbf{in} \ \overline{a}, \overline{b}) \ as \ bs \ \overline{xs} \ \overline{as_0} \ \overline{bs_0} \end{aligned}$$


# map with Free Variables

- maps involving **free variables** are more complicated to differentiate

`let xs = map ( $\lambda a \rightarrow a \cdot b$ ) as`



# map with Free Variables

- maps involving **free variables** are more complicated to differentiate

```
let xs = map ( $\lambda a \rightarrow a \cdot b$ ) as
```

- Naive approach:** turn free variables into bound variables.

```
let xs = map ( $\lambda a \ b' \rightarrow a \cdot b'$ ) as (replicate n b)
```



# map with Free Variables

- maps involving **free variables** are more complicated to differentiate

`let xs = map ( $\lambda a \rightarrow a \cdot b$ ) as`

- Naive approach:** turn free variables into bound variables.

`let xs = map ( $\lambda a b' \rightarrow a \cdot b'$ ) as (replicate n b)`

- Problem:** asymptotically **inefficient** for partially used free arrays.

`map ( $\lambda(i, as') \rightarrow as'[i]$ ) is (replicate n as),`





# Efficient `map`s with Free Variables

- In an impure language, asymptotics-preserving adjoint updates for free array variables can be implemented as a **generalized reduction**.



# Efficient maps with Free Variables

- In an impure language, asymptotics-preserving adjoint updates for free array variables can be implemented as a **generalized reduction**.
  - ▶ The adjoint of a free array variable  $as[i]$  can be updated with an operation  $\overline{as}[i] += v$ .



# Efficient maps with Free Variables

- In an impure language, asymptotics-preserving adjoint updates for free array variables can be implemented as a **generalized reduction**.
  - ▶ The adjoint of a free array variable  $as[i]$  can be updated with an operation  $\overline{as}[i] += v$ .
- In our pure setting, we introduce **accumulators**.
  - ▶ **Write-only** view of an array.
  - ▶ Preserves purely functional reasoning in the compiler.
  - ▶ Preserves asymptotics by operationally doing in-place updates.



# Loops



# Loops

- **Loops** in Futhark are sugar for tail-recursive functions.
- **Loop parameters** are **variables** which are variant through the loop and are returned as the result of the loop.

```
loop y = 2 for  $i = 0 \dots n - 1$  do  
  let  $y' = y * y$   
  in  $y'$ 
```

```
y = 2  
for  $i = 0 \dots n - 1$  do  
   $y = y * y$ 
```

(Imperative analog)



# Loops

- **Loops** in Futhark are sugar for tail-recursive functions.
- **Loop parameters** are **variables** which are variant through the loop and are returned as the result of the loop.

```
loop  $y = 2$  for  $i = 0 \dots n - 1$  do  
  let  $y' = y * y$   
  in  $y'$ 
```

```
 $y = 2$   
for  $i = 0 \dots n - 1$  do  
   $y = y * y$ 
```

(Imperative analog)

- Since the adjoints of the **loop body** are computed in **reverse order**, the loop parameter  $y$  needs to be **saved** for each iteration.



# Differentiating Loops

```
let  $y'' =$   
  loop  $y = y_0$  for  $i = 0 \dots n - 1$  do  
     $stms_{loop}$   
  in  $y'$ 
```

1. Execute the original loop, save the value of  $y$  in each iteration in  $ys$ .

```
2 let  $ys_0 = scratch(n,$   
3      $sizeof(y_0))$   
4 let  $(y'', ys) =$   
5 loop  $(y, ys) = (y_0, ys_0)$   
6 for  $i = 0 \dots n - 1$  do  
7   let  $ys[i] = y$   
8    $stms_{loop}$   
9   in  $(y', ys)$   
12 let  $(\overline{y''''}, \overline{fvs_l}) =$   
13 loop  $(\overline{y}, \overline{fvs_l}) = (\overline{y''}, \overline{fvs_{l_0}})$   
14 for  $i = n - 1 \dots 0$  do  
15   let  $y = ys[i]$   
16    $\overrightarrow{stms_{loop}}$   
17    $\overleftarrow{stms_{loop}}$   
18   in  $(\overline{y'}, \overline{fvs'_l})$   
19 let  $\overline{y_0} += \overline{y''''}$ 
```

Forward sweep

Reverse sweep



# Differentiating Loops

```
let  $y'' =$   
  loop  $y = y_0$  for  $i = 0 \dots n - 1$  do  
     $stms_{loop}$   
  in  $y'$ 
```

1. Execute the original loop, save the value of  $y$  in each iteration in  $ys$ .
2. Compute the adjoint contributions of the loop.

```
2 let  $ys_0 = scratch(n,$   
3      $sizeof(y_0))$   
4 let  $(y'', ys) =$   
5 loop  $(y, ys) = (y_0, ys_0)$   
6 for  $i = 0 \dots n - 1$  do  
7   let  $ys[i] = y$   
8    $stms_{loop}$   
9   in  $(y', ys)$   
12 let  $(\overline{y'''}, \overline{fvs}_l) =$   
13 loop  $(\overline{y}, \overline{fvs}_l) = (\overline{y''}, \overline{fvs}_{l_0})$   
14 for  $i = n - 1 \dots 0$  do  
15   let  $y = ys[i]$   
16    $\overrightarrow{stms_{loop}}$   
17    $\overleftarrow{stms_{loop}}$   
18   in  $(y', \overline{fvs}'_l)$   
19 let  $\overline{y_0} += \overline{y'''}$ 
```

Forward sweep

Reverse sweep





# Differentiating Loops

```
let  $y'' =$   
  loop  $y = y_0$  for  $i = 0 \dots n - 1$  do  
     $stms_{loop}$   
  in  $y'$ 
```

1. Execute the original loop, save the value of  $y$  in each iteration in  $ys$ .
2. Compute the adjoint contributions of the loop.
  - ▶ Run the loop backwards

```
2 let  $ys_0 = scratch(n,$   
3      $sizeof(y_0))$   
4 let  $(y'', ys) =$   
5 loop  $(y, ys) = (y_0, ys_0)$   
6 for  $i = 0 \dots n - 1$  do  
7   let  $ys[i] = y$   
8    $stms_{loop}$   
9   in  $(y', ys)$   
12 let  $(\overline{y''''}, \overline{fvs'_l}) =$   
13 loop  $(\overline{y}, \overline{fvs'_l}) = (\overline{y''}, \overline{fvs'_{l_0}})$   
14 for  $i = n - 1 \dots 0$  do  
15   let  $y = ys[i]$   
16    $\overrightarrow{stms_{loop}}$   
17    $\overleftarrow{stms_{loop}}$   
18   in  $(\overline{y'}, \overline{fvs'_l})$   
19 let  $\overline{y_0} += \overline{y''''}$ 
```

Forward sweep

Reverse sweep



# Differentiating Loops

```
let  $y'' =$   
  loop  $y = y_0$  for  $i = 0 \dots n - 1$  do  
     $stms_{loop}$   
  in  $y'$ 
```

1. Execute the original loop, save the value of  $y$  in each iteration in  $ys$ .
2. Compute the adjoint contributions of the loop.
  - ▶ Run the loop backwards
  - ▶ Restore the value of  $y$  from  $ys$

```
2 let  $ys_0 = scratch(n,$   
3      $sizeof(y_0))$   
4 let  $(y'', ys) =$   
5 loop  $(y, ys) = (y_0, ys_0)$   
6 for  $i = 0 \dots n - 1$  do  
7   let  $ys[i] = y$   
8    $stms_{loop}$   
9   in  $(y', ys)$   
12 let  $(\overline{y''''}, \overline{fvs}_l) =$   
13 loop  $(\overline{y}, \overline{fvs}_l) = (\overline{y''}, \overline{fvs}_{l_0})$   
14 for  $i = n - 1 \dots 0$  do  
15   let  $y = ys[i]$   
16    $\overrightarrow{stms_{loop}}$   
17    $\overleftarrow{stms_{loop}}$   
18   in  $(\overline{y'}, \overline{fvs}'_l)$   
19 let  $\overline{y_0} += \overline{y''''}$ 
```

Forward sweep

Reverse sweep



# Differentiating Loops

```
let  $y'' =$   
  loop  $y = y_0$  for  $i = 0 \dots n - 1$  do  
     $stms_{loop}$   
  in  $y'$ 
```

1. Execute the original loop, save the value of  $y$  in each iteration in  $ys$ .
2. Compute the adjoint contributions of the loop.
  - ▶ Run the loop backwards
  - ▶ Restore the value of  $y$  from  $ys$
  - ▶ Re-execute the body of the original loop

```
2 let  $ys_0 = scratch(n,$   
3      $sizeof(y_0))$   
4 let  $(y'', ys) =$   
5 loop  $(y, ys) = (y_0, ys_0)$   
6 for  $i = 0 \dots n - 1$  do  
7   let  $ys[i] = y$   
8    $stms_{loop}$   
9   in  $(y', ys)$   
12 let  $(\overline{y'''}, \overline{fvs}_l) =$   
13 loop  $(\overline{y}, \overline{fvs}_l) = (\overline{y''}, \overline{fvs}_{l_0})$   
14 for  $i = n - 1 \dots 0$  do  
15   let  $y = ys[i]$   
16    $\overrightarrow{stms_{loop}}$   
17    $\overleftarrow{stms_{loop}}$   
18   in  $(\overline{y'}, \overline{fvs}'_l)$   
19 let  $\overline{y}_0 += \overline{y'''}$ 
```

Forward sweep

Reverse sweep



# Differentiating Loops

```
let  $y'' =$   
  loop  $y = y_0$  for  $i = 0 \dots n - 1$  do  
     $stms_{loop}$   
  in  $y'$ 
```

1. Execute the original loop, save the value of  $y$  in each iteration in  $ys$ .
2. Compute the adjoint contributions of the loop.
  - ▶ Run the loop backwards
  - ▶ Restore the value of  $y$  from  $ys$
  - ▶ Re-execute the body of the original loop
  - ▶ Compute the adjoints of the body



```
2 let  $ys_0 = scratch(n,$   
3      $sizeof(y_0))$   
4 let  $(y'', ys) =$   
5 loop  $(y, ys) = (y_0, ys_0)$   
6 for  $i = 0 \dots n - 1$  do  
7   let  $ys[i] = y$   
8    $stms_{loop}$   
9   in  $(y', ys)$   
12 let  $(\overline{y'''}, \overline{fvs_l}) =$   
13 loop  $(\overline{y}, \overline{fvs_l}) = (\overline{y''}, \overline{fvs_{l_0}})$   
14 for  $i = n - 1 \dots 0$  do  
15   let  $y = ys[i]$   
16    $\overrightarrow{stms_{loop}}$   
17    $\overleftarrow{stms_{loop}}$   
18   in  $(\overline{y'}, \overline{fvs'_l})$   
19 let  $\overline{y_0} += \overline{y'''}$ 
```

Forward sweep

Reverse sweep

# Loop Strip-mining

- **Loop strip-mining** partitions a loop into a loop nest

```
loop  $y = y_0$  for  $i = 0 \dots n^3 - 1$  do  
  stms
```

⇒

```
loop  $y_j = y_0$  for  $j = 0 \dots n - 1$  do  
  loop  $y_k = y_j$  for  $k = 0 \dots n - 1$  do  
    loop  $y_m = y_k$  for  $m = 0 \dots n - 1$  do  
      let  $i = j * n^{2/3} + k * n^{1/3} + m$   
      stms
```



# Loop Strip-mining

- **Loop strip-mining** partitions a loop into a loop nest

```
loop  $y = y_0$  for  $i = 0 \dots n^3 - 1$  do  
  stms
```

$\implies$

```
loop  $y_j = y_0$  for  $j = 0 \dots n - 1$  do  
  loop  $y_k = y_j$  for  $k = 0 \dots n - 1$  do  
    loop  $y_m = y_k$  for  $m = 0 \dots n - 1$  do  
      let  $i = j * n^{2/3} + k * n^{1/3} + m$   
      stms
```

- For the original loop, we save  $n^3$  versions of  $y$  on the tape.



# Loop Strip-mining

- **Loop strip-mining** partitions a loop into a loop nest

```
loop  $y = y_0$  for  $i = 0 \dots n^3 - 1$  do  
  stms  $\implies$  loop  $y_j = y_0$  for  $j = 0 \dots n - 1$  do  
    loop  $y_k = y_j$  for  $k = 0 \dots n - 1$  do  
      loop  $y_m = y_k$  for  $m = 0 \dots n - 1$  do  
        let  $i = j * n^{2/3} + k * n^{1/3} + m$   
        stms
```

- For the original loop, we save  $n^3$  versions of  $y$  on the tape.
- For the strip-mined loop, only  $3n$  versions are saved. (With an increased re-execution overhead factor of 3.)




# Benchmarks





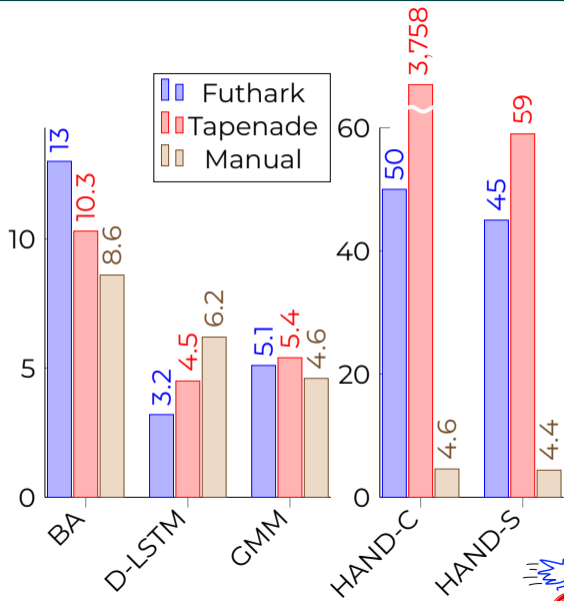
# CPU Benchmarks - ADBench



- ADBench: a collection of AD benchmarks for comparing sequential AD tools.




# CPU Benchmarks - ADBench



- ADBench: a collection of AD benchmarks for comparing sequential AD tools.
- Benchmarked Futhark using its C backend.



# CPU Benchmarks - ADBench




- ADBench: a collection of AD benchmarks for comparing sequential AD tools.
- Benchmarked Futhark using its C backend.
- Performance measured in **AD overhead**:

$$\frac{\text{differentiated runtime}}{\text{original runtime}}$$



# GPU Benchmarks - vs. Enzyme




- Performance measured in **AD overhead**:

$$\frac{\text{differentiated runtime}}{\text{original runtime}}$$



# GPU Benchmarks - vs. Enzyme



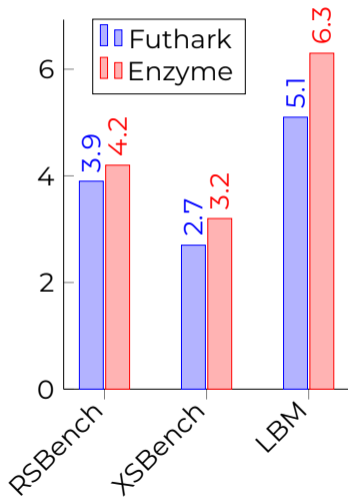
- Performance measured in **AD overhead**:

$$\frac{\text{differentiated runtime}}{\text{original runtime}}$$

- Enzyme is LLVM compiler plugin that performs AD on a low-level imperative IR.



# GPU Benchmarks - vs. Enzyme




- Performance measured in **AD overhead**:

$$\frac{\text{differentiated runtime}}{\text{original runtime}}$$

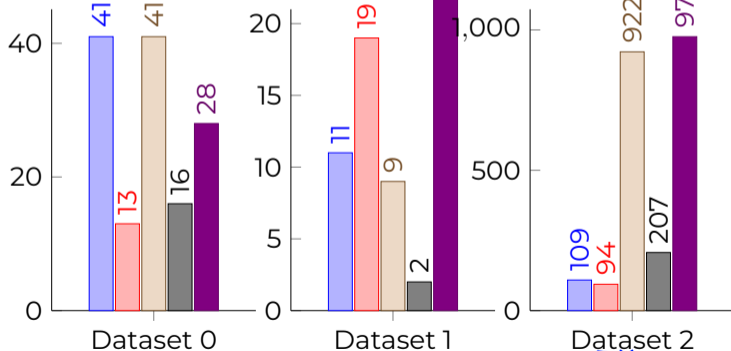
- Enzyme is LLVM compiler plugin that performs AD on a low-level imperative IR.
- RS Bench and XS Bench are comprised of a large parallel loop with inner sequential loops and branches.
- LBM consists of a large sequential loop containing a parallel loop.



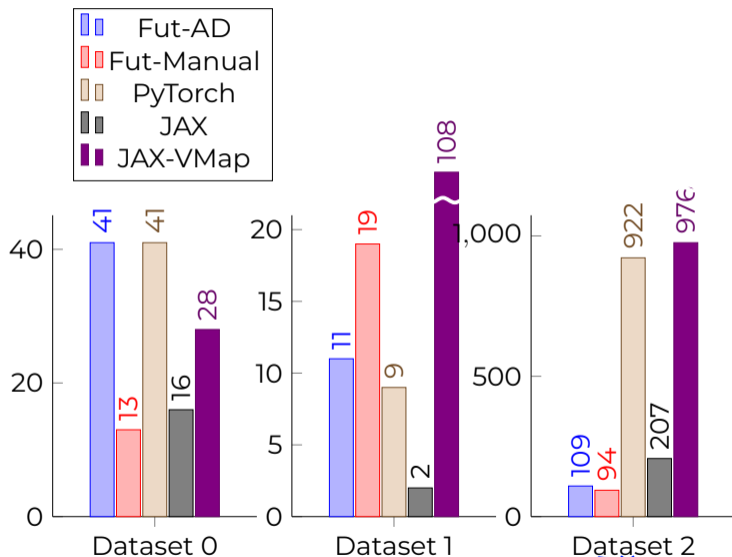
# GPU Benchmarks - $k$ -means



- Performance measured in **milliseconds**.
- $k$ -means clustering using AD-based Newton's method to find cluster centers.



# GPU Benchmarks - $k$ -means




- Performance measured in **milliseconds**.
- $k$ -means clustering using AD-based Newton's method to find cluster centers.
- PyTorch and JAX use hand-tuned matrix primitives; JAX(vmap) instead uses JAX's vectorizing map operation for these operations, in analog with Futhark.





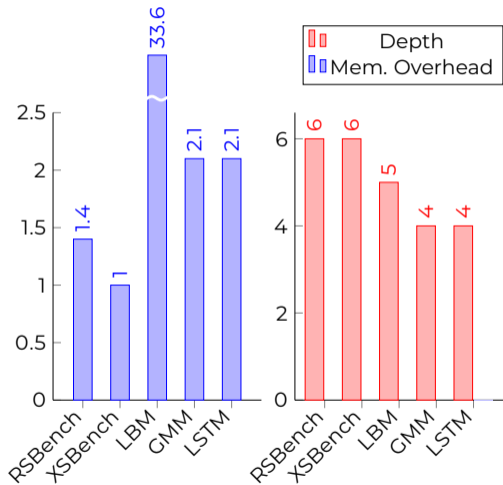
# GPU Benchmarks - Sparse $k$ -means



- Performance measured in **seconds**.
- PyTorch and JAX use hand-tuned matrix primitives and sparse libraries.
- Futhark just uses a standard CSR implementation.



# GPU Benchmarks - Depth and Memory Consumption



- **AD Memory overhead:**

$$\frac{\text{differentiated mem. consumption}}{\text{original mem. consumption}}$$

- With **loop strip-mining**, LBM's memory overhead is **reduced to 8.7**, with only a 1.3× increase in runtime.
- **Strong performance** on programs with **non-trivial depth** demonstrates the viability of a recomputation-based approach to AD.



## Conclusions



# Conclusions

- AD in a **nested-parallel, high-level** and **hardware-neutral** functional language.



# Conclusions

- AD in a **nested-parallel, high-level** and **hardware-neutral** functional language.
- **Key idea:** high-level differentiation using specialized rules for parallel combinators.



# Conclusions

- AD in a **nested-parallel, high-level** and **hardware-neutral** functional language.
- **Key idea:** high-level differentiation using specialized rules for parallel combinators.
- **Key idea:** re-computation instead of a tape (except for loops!).



# Conclusions

- AD in a **nested-parallel, high-level** and **hardware-neutral** functional language.
- **Key idea:** high-level differentiation using specialized rules for parallel combinators.
- **Key idea:** re-computation instead of a tape (except for loops!).
- Strong performance against state-of-the-art AD competitors.



# Conclusions


- AD in a **nested-parallel, high-level** and **hardware-neutral** functional language.
- **Key idea:** high-level differentiation using specialized rules for parallel combinators.
- **Key idea:** re-computation instead of a tape (except for loops!).
- Strong performance against state-of-the-art AD competitors.
- The implementation is now mature and available in the Futhark compiler.





# Thing #2: AUTOMAP

Robert Schenck, Nikolaj Hey Hinnerskov, Troels Henriksen,  
Magnus Madsen, Martin Elsmann



# Rank polymorphism

- $1 + 2 \Rightarrow 3$



# Rank polymorphism

- $1 + 2 \Rightarrow 3$

- $[1, 2, 3] + [4, 5, 6] \Rightarrow [5, 7, 9]$



# Rank polymorphism

- $1 + 2 \Rightarrow 3$
- $[1, 2, 3] + [4, 5, 6] \Rightarrow [5, 7, 9]$
- $[1, 2, 3] + 4 \Rightarrow [5, 6, 7]$



# Rank polymorphism

- $1 + 2 \Rightarrow 3$
- $[1, 2, 3] + [4, 5, 6] \Rightarrow [5, 7, 9]$
- $[1, 2, 3] + 4 \Rightarrow [5, 6, 7]$
- $\text{sqrt } [[1, 4, 9], [16, 25, 36]] \Rightarrow [[1, 2, 3], [4, 5, 6]]$



# Rank polymorphism

## Rank polymorphism

The ability to apply functions to arguments with different ranks than the function expects.



# Rank polymorphism

## Rank polymorphism

The ability to apply functions to arguments with different ranks than the function expects.

- Makes code easier to read, more enjoyable to write, and closer to math:

`map (+) [1, 2, 3] [4, 5, 6]` vs. `[1, 2, 3] + [4, 5, 6]`



# Rank polymorphism

## Rank polymorphism

The ability to apply functions to arguments with different ranks than the function expects.

- Makes code easier to read, more enjoyable to write, and closer to math:

`map (+) [1, 2, 3] [4, 5, 6]` vs. `[1, 2, 3] + [4, 5, 6]`

- This work: how do we get **rank polymorphic applications** in a statically-typed language with parametric polymorphism?





- `map f xs` applies `f` to each element of `xs`:

`map f [x_0, x_1, ..., x_n] = [f x_0, f x_1, ..., f x_n]`



# map and rep

- `map f xs` applies `f` to each element of `xs`:

`map f [x_0, x_1, ..., x_n] = [f x_0, f x_1, ..., f x_n]`

- You can `map` functions that take multiple arguments too:

`map (+) [x_0, ..., x_n] [y_0, ..., y_n]`  
`= [x_0 + y_0, ..., x_n + y_n]`



## map and rep

- **map**  $f$   $x_s$  applies  $f$  to each element of  $x_s$ :

**map**  $f$   $[x_0, x_1, \dots, x_n] = [f\ x_0, f\ x_1, \dots, f\ x_n]$

- You can **map** functions that take multiple arguments too:

**map**  $(+)$   $[x_0, \dots, x_n]$   $[y_0, \dots, y_n]$   
 $= [x_0 + y_0, \dots, x_n + y_n]$

- **rep**  $x$  makes an array of unspecified length whose elements are all  $x$ :

**rep**  $x = [x, x, \dots, x]$

- ▶ We'll ignore the question of how many elements are needed.



# An example

`[[1,2],[3,4]] + 1`



# An example

`[[1,2],[3,4]] + 1`

elaborates to

`[[1,2],[3,4]] + rep (rep 1)`



# An example

`[[1,2],[3,4]] + 1`

elaborates to

`[[1,2],[3,4]] + rep (rep 1)`

which further elaborates to

```
map (map (+)) [[1,2],[3,4]]
  (rep (rep 1))
```



# An example

```
[[1,2],[3,4]] + 1
```

elaborates to

```
[[1,2],[3,4]] + rep (rep 1)
```

which further elaborates to

```
map (map (+)) [[1,2],[3,4]]  
  (rep (rep 1))
```

```
xss : [][]int  
yss : [][]int  
f: []int → [][]int → int
```

---

```
f xss yss
```



# An example

```
[[1,2],[3,4]] + 1
```

elaborates to

```
[[1,2],[3,4]] + rep (rep 1)
```

which further elaborates to

```
map (map (+)) [[1,2],[3,4]]  
  (rep (rep 1))
```

```
xss : [][]int  
yss : [][]int  
f: []int → [][]int → int
```

---

```
f xss yss
```

First, we map `f` across both matrices:

```
map f xss yss
```





# An example

```
[[1,2],[3,4]] + 1
```

elaborates to

```
[[1,2],[3,4]] + rep (rep 1)
```

which further elaborates to

```
map (map (+)) [[1,2],[3,4]]  
  (rep (rep 1))
```

```
xss : [][]int  
yss : [][]int  
f: []int → [][]int → int
```

---

```
f xss yss
```

First, we map `f` across both matrices:

```
map f xss yss
```

Because of the `map`, `yss` must be replicated:

```
map f xss (rep yss)
```



# An example

```
[[1,2],[3,4]] + 1
```

elaborates to

```
[[1,2],[3,4]] + rep (rep 1)
```

which further elaborates to

```
map (map (+)) [[1,2],[3,4]]  
  (rep (rep 1))
```

```
xss : [][]int  
yss : [][]int  
f: []int → [][]int → int
```

---

```
f xss yss
```

First, we map `f` across both matrices:

```
map f xss yss
```

Because of the `map`, `yss` must be replicated:

```
map f xss (rep yss)
```

`rep`s can often be eliminated

```
map (λxs → f xs yss) xss
```



## Goal

For each function application, the compiler should automatically insert **maps** or **reps** to make the application **rank-correct**.



## Goal

For each function application, the compiler should automatically insert **maps** or **reps** to make the application **rank-correct**.

$$f\ x \implies \text{map} \left( \dots (\text{map } f) \dots \right) \left( \text{rep } \dots (\text{rep } x) \dots \right)$$


# Challenge: ambiguity

```
sum : []int → int  
length : []a → int  
xss : [][]int
```

-----  
sum (length xss)



# Challenge: ambiguity

```
sum : []int → int  
length : []a → int  
xss : [][]int
```

-----  
sum (length xss)

Many rank-correct elaborations:

1. sum (**rep** (length xss))



# Challenge: ambiguity

```
sum : []int → int  
length : []a → int  
xss : [][]int
```

-----  
sum (length xss)

Many rank-correct elaborations:

1. sum (**rep** (length xss))
2. sum (**map** length xss)



# Challenge: ambiguity

```
sum : []int → int  
length : []a → int  
xss : [][]int
```

-----  
sum (length xss)

Many rank-correct elaborations:

1. sum (**rep** (length xss))
2. sum (**map** length xss)
3. **map** sum (**map** (**map** length) (**rep** xss))
4. ...





# The Strategy

## Rule 1

An application can be **mapped** or **repped** (or neither) but **never both**.



# The Strategy

## Rule 1

An application can be **mapped** or **repped** (or neither) but **never both**.

- **OK:**

- ▶ `map f x`



## Rule 1

An application can be **mapped** or **repped** (or neither) but **never both**.

- **OK:**

- ▶ `map f x`

- ▶ `g (rep (rep x))`



## Rule 1

An application can be **mapped** or **repped** (or neither) but **never both**.

- **OK:**

- ▶ `map f x`
- ▶ `g (rep (rep x))`
- ▶ `(map (map h) x) (rep y)`



## Rule 1

An application can be **mapped** or **repped** (or neither) but **never both**.

### ▪ OK:

- ▶ `map f x`
- ▶ `g (rep (rep x))`
- ▶ `(map (map h) x) (rep y)`

### BAD:

- ▶ `map f (rep x)`
- ▶ `(map (map g)) (rep x)`



## Rule 1

An application can be **mapped** or **repped** (or neither) but **never both**.

- **OK:**

- ▶ `map f x`
- ▶ `g (rep (rep x))`
- ▶ `(map (map h) x) (rep y)`

- **BAD:**

- ▶ `map f (rep x)`
- ▶ `(map (map g)) (rep x)`

- Never necessary to **map** and **rep** in the same application to obtain a rank-correct program.



# The Strategy

## Rule 2

**Minimize** the number of inserted **maps** and **reps**.



# The Strategy

## Rule 2

**Minimize** the number of inserted **maps** and **reps**.

- Generally aligns with programmer's intent; makes for a simple mental model.





## Rule 2

**Minimize** the number of inserted **maps** and **reps**.

- Generally aligns with programmer's intent; makes for a simple mental model.
- The minimization is over **all** the applications of a top-level definition.
  - ▶ Only have to choose from the set of minimal solutions.



## Rule 2

**Minimize** the number of inserted **maps** and **reps**.

- Generally aligns with programmer's intent; makes for a simple mental model.
- The minimization is over **all** the applications of a top-level definition.
  - ▶ Only have to choose from the set of minimal solutions.

`sum (length xss)` can be elaborated to:

1. `sum (rep (length xss))`



## Rule 2

**Minimize** the number of inserted **maps** and **reps**.

- Generally aligns with programmer's intent; makes for a simple mental model.
- The minimization is over **all** the applications of a top-level definition.
  - ▶ Only have to choose from the set of minimal solutions.

`sum (length xss)` can be elaborated to:

1. `sum (rep (length xss))`
2. `sum (map length xss)`



## Rule 2

**Minimize** the number of inserted **maps** and **reps**.

- Generally aligns with programmer's intent; makes for a simple mental model.
- The minimization is over **all** the applications of a top-level definition.
  - ▶ Only have to choose from the set of minimal solutions.

`sum (length xss)` can be elaborated to:

1. `sum (rep (length xss))`
2. `sum (map length xss)`
3. `map sum (map (map length) (rep xss))`



## Rule 2

**Minimize** the number of inserted **maps** and **reps**.

- Generally aligns with programmer's intent; makes for a simple mental model.
- The minimization is over **all** the applications of a top-level definition.
  - ▶ Only have to choose from the set of minimal solutions.

`sum (length xss)` can be elaborated to:

1. `sum (rep (length xss))`
2. `sum (map length xss)`
3. `map sum (map (map length) (rep xss))`
4. ...



## Rule 2

**Minimize** the number of inserted **maps** and **reps**.

- Generally aligns with programmer's intent; makes for a simple mental model.
- The minimization is over **all** the applications of a top-level definition.
  - ▶ Only have to choose from the set of minimal solutions.

`sum (length xss)` can be elaborated to:

1. `sum (rep (length xss))`
2. `sum (map length xss)`
3. ~~`map sum (map (map length) (rep xss))`~~
4. ...



# Challenge: elaboration is global

- `sum (length xss)` can be elaborated to:
  1. `sum (rep (length xss))`
    - ▶ **Local** reasoning: in the application of `sum` to `length xss`, the argument is underdimensioned, so a `rep` is inserted.



# Challenge: elaboration is global

- `sum (length xss)` can be elaborated to:
  1. `sum (rep (length xss))`
    - ▶ **Local** reasoning: in the application of `sum` to `length xss`, the argument is underdimensioned, so a `rep` is inserted.
  2. `sum (map length xss)`
    - ▶ **Global** reasoning: `length xss` is rank-correct as-is, but a `map` is inserted because of the outer `sum` application.





# Challenge: elaboration is global

- `sum (length xss)` can be elaborated to:
  1. `sum (rep (length xss))`
    - ▶ **Local** reasoning: in the application of `sum` to `length xss`, the argument is underdimensioned, so a **rep** is inserted.
  2. `sum (map length xss)`
    - ▶ **Global** reasoning: `length xss` is rank-correct as-is, but a **map** is inserted because of the outer `sum` application.
- Elaborations of inner applications affect outer applications.
  - ▶ To find all minimal elaborations, must consider all applications **simultaneously**.



# Challenge: type variables

- Futhark has parametric polymorphism:

```
id : a → a
```

```
length : []a → int
```



# Challenge: type variables

- Futhark has parametric polymorphism:

`id : a → a`

`length : []a → int`

- A type variable can have any rank!



# Challenge: type variables

- Futhark has parametric polymorphism:

`id : a → a`

`length : []a → int`

- A type variable can have any rank!
- How do we statically insert **maps** and **reps** in the presence of type variables, whose ranks aren't known?



# Constraints

- Suppose

$f : p \rightarrow b$

$x : a$



# Constraints

- Suppose

$$f : p \rightarrow b$$

$$x : a$$

- The application  $f \ x$  has constraint

$$p = a$$



# Constraints

- Suppose

$$f : p \rightarrow b$$

$$x : a$$

- The application  $f \ x$  has constraint

$$p = a$$

- We only care about rank, so relax to

$$|p| = |a|$$

where  $|p|$  is the **rank** of  $p$ .

For example,  $|[\text{[]}]_{\text{int}}| = 2$  and  $|\text{int}| = 0$ .



## Constraints - map

Rank polymorphisms means rank differences are allowed.





## Constraints - map

Rank polymorphisms means rank differences are allowed.

- Case  $|p| < |a|$ :
  - ▶ Introduce a **rank variable**  $M$  to account for the difference:

$$M + |p| = |a|$$



# Constraints - map

Rank polymorphisms means rank differences are allowed.

- Case  $|p| < |a|$ :

- ▶ Introduce a **rank variable**  $M$  to account for the difference:

$$M + |p| = |a|$$

- ▶ `sqrt : int → int`  
`[1,2,3] : []int`

Application `sqrt [1,2,3]` gives the constraint

$$M + \underbrace{|int|}_0 = \underbrace{|[]int|}_1 \implies M = 1$$



# Constraints - map

Rank polymorphisms means rank differences are allowed.

- Case  $|p| < |a|$ :

- ▶ Introduce a **rank variable**  $M$  to account for the difference:

$$M + |p| = |a|$$

- ▶ `sqrt : int → int`  
`[1,2,3] : []int`

Application `sqrt [1,2,3]` gives the constraint

$$M + \underbrace{|int|}_0 = \underbrace{|[]int|}_1 \implies M = 1$$

- ▶  $M$  is equal to the number of **maps** required:

`map sqrt [1,2,3]`



# Constraints - rep

- Case  $|p| > |a|$ :
  - ▶ Introduce a rank variable  $R$  to account for the difference:

$$|p| = R + |a|$$



# Constraints - rep

- Case  $|p| > |a|$ :
  - ▶ Introduce a rank variable  $R$  to account for the difference:

$$|p| = R + |a|$$

- ▶ Example: `length : []b → int` The application `length 3` gives the constraint

$$|[]b| = R + |\text{int}|$$

$$1 + |b| = R \implies R = 1, |b| = 0$$



# Constraints - rep

- Case  $|p| > |a|$ :
  - ▶ Introduce a rank variable  $R$  to account for the difference:

$$|p| = R + |a|$$

- ▶ Example: `length : []b → int` The application `length 3` gives the constraint

$$|[]b| = R + |\text{int}|$$

$$1 + |b| = R \implies R = 1, |b| = 0$$

$$R = 2, |b| = 1$$

...



# Constraints - rep

- Case  $|p| > |a|$ :
  - ▶ Introduce a rank variable  $R$  to account for the difference:

$$|p| = R + |a|$$

- ▶ Example: `length : []b → int` The application `length 3` gives the constraint

$$|[]b| = R + |\text{int}|$$

$$1 + |b| = R \implies R = 1, |b| = 0$$

$$R = 2, |b| = 1$$

...

- ▶  $R$  is equal to the number of **reps** required:
  - ▶ `length (rep 3)`
  - ▶ `length (rep (rep 3))`
  - ▶ ...



# Constraints - Summary

- Each application of a function  $f : p \rightarrow c$  to an argument  $x : a$  generates a constraint

$$M + |p| = R + |a|$$





# Constraints - Summary

- Each application of a function  $f : p \rightarrow c$  to an argument  $x : a$  generates a constraint

$$M + |p| = R + |a|$$

- Rule 1: can either **map** or **rep** but not both

$$M = 0 \text{ or } R = 0$$



# Constraints to ILPs

- Collect the constraints for each function application.



# Constraints to ILPs

- Collect the constraints for each function application.
- Example: `sum (length xss)`



# Constraints to ILPs

- Collect the constraints for each function application.
- Example: `sum (length xss)`

$$\left. \begin{array}{l} M_1 + 1 + |a| = R_1 + 2 \\ M_1 = 0 \text{ or } R_1 = 0 \end{array} \right\} \text{length}$$



# Constraints to ILPs

- Collect the constraints for each function application.
- Example: `sum (length xss)`

$$\begin{array}{l} M_1 + 1 + |a| = R_1 + 2 \\ M_1 = 0 \text{ or } R_1 = 0 \end{array} \left. \vphantom{\begin{array}{l} M_1 + 1 + |a| = R_1 + 2 \\ M_1 = 0 \text{ or } R_1 = 0 \end{array}} \right\} \text{length}$$
$$\begin{array}{l} M_2 + 1 = R_2 + M_1 \\ M_2 = 0 \text{ or } R_2 = 0 \end{array} \left. \vphantom{\begin{array}{l} M_2 + 1 = R_2 + M_1 \\ M_2 = 0 \text{ or } R_2 = 0 \end{array}} \right\} \text{sum}$$



# Constraints to ILPs

- Collect the constraints for each function application.
- Example: `sum (length xss)`

$$\left. \begin{array}{l} M_1 + 1 + |a| = R_1 + 2 \\ M_1 = 0 \text{ or } R_1 = 0 \end{array} \right\} \text{length}$$
$$\left. \begin{array}{l} M_2 + 1 = R_2 + M_1 \\ M_2 = 0 \text{ or } R_2 = 0 \end{array} \right\} \text{sum}$$

- Rule 2: Minimize the number of **maps** and **reps**



# Constraints to ILPs

- Collect the constraints for each function application.
- Example: `sum (length xss)`

minimize

$$M_1 + R_1 + M_2 + R_2$$

subject to

$$\left. \begin{array}{l} M_1 + 1 + |a| = R_1 + 2 \\ M_1 = 0 \text{ or } R_1 = 0 \end{array} \right\} \text{length}$$

$$\left. \begin{array}{l} M_2 + 1 = R_2 + M_1 \\ M_2 = 0 \text{ or } R_2 = 0 \end{array} \right\} \text{sum}$$

- Rule 2: Minimize the number of **maps** and **reps**



# Constraints to ILPs

- Collect the constraints for each function application.
- Example: `sum (length xss)`

minimize

$$M_1 + R_1 + M_2 + R_2$$

subject to

$$\left. \begin{array}{l} M_1 + 1 + |a| = R_1 + 2 \\ M_1 = 0 \text{ or } R_1 = 0 \end{array} \right\} \text{length}$$

$$\left. \begin{array}{l} M_2 + 1 = R_2 + M_1 \\ M_2 = 0 \text{ or } R_2 = 0 \end{array} \right\} \text{sum}$$

- Rule 2: Minimize the number of **maps** and **reps**
- The or-constraints can be linearized to obtain an integer linear program (ILP).





# AUTOMAP TL;DR

1. For each application generate rank equality and Rule 1 (**map** or **rep** but not both) constraints.



# AUTOMAP TL;DR

1. For each application generate rank equality and Rule 1 (**map** or **rep** but not both) constraints.
2. Transform the constraint set into an ILP and solve.



# AUTOMAP TL;DR

1. For each application generate rank equality and Rule 1 (**map** or **rep** but not both) constraints.
2. Transform the constraint set into an ILP and solve.
3. Use the ILP solution to elaborate. E.g., if the  $i$ -th application  $f\ x$  has  $M_i = 3$  and  $R_i = 0$ :

$$f\ x \quad \Longrightarrow \quad \mathbf{map}\ (\mathbf{map}\ (\mathbf{map}\ f))\ x$$



# AUTOMAP TL;DR

1. For each application generate rank equality and Rule 1 (**map** or **rep** but not both) constraints.
2. Transform the constraint set into an ILP and solve.
3. Use the ILP solution to elaborate. E.g., if the  $i$ -th application  $f\ x$  has  $M_i = 3$  and  $R_i = 0$ :

$$f\ x \quad \Longrightarrow \quad \mathbf{map}\ (\mathbf{map}\ (\mathbf{map}\ f))\ x$$

4. Type check elaborated program and continue with compilation as usual.



# User experience

- **map** and **rep** are normal source-level functions.
  - ▶ Programmer free to use AUTOMAP to whatever extent they wish.



# User experience

- **map** and **rep** are normal source-level functions.
  - ▶ Programmer free to use AUTOMAP to whatever extent they wish.

## Ambiguity feedback:

```
Error: sum (length xss) has multiple elaborations:  
  1. sum (rep (length xss))  
  2. sum (map length xss)
```

- ▶ Nice error messages.
- ▶ Disambiguation is easy: just insert a **map** or **rep** into the source.



# User experience

- **map** and **rep** are normal source-level functions.
  - ▶ Programmer free to use AUTOMAP to whatever extent they wish.

## Ambiguity feedback:

```
Error: sum (length xss) has multiple elaborations:
```

1. sum (**rep** (length xss))
2. sum (**map** length xss)

- ▶ Nice error messages.
  - ▶ Disambiguation is easy: just insert a **map** or **rep** into the source.
- 
- Fully transparent: the compiler can always elaborate any implicit **maps** or **reps**.



# User experience

- **map** and **rep** are normal source-level functions.
  - ▶ Programmer free to use AUTOMAP to whatever extent they wish.

## Ambiguity feedback:

```
Error: sum (length xss) has multiple elaborations:
```

1. sum (**rep** (length xss))
2. sum (**map** length xss)

- ▶ Nice error messages.
  - ▶ Disambiguation is easy: just insert a **map** or **rep** into the source.
- 
- Fully transparent: the compiler can always elaborate any implicit **maps** or **reps**.





## Practical impact

- We implemented AUTOMAP in **Futhark**, a functional array language that supports parametric polymorphism and top-level let-polymorphism.



## Practical impact

- We implemented AUTOMAP in **Futhark**, a functional array language that supports parametric polymorphism and top-level let-polymorphism.
- Difficult to quantify value of feature that is glorified syntax sugar.



## Practical impact

- We implemented AUTOMAP in **Futhark**, a functional array language that supports parametric polymorphism and top-level let-polymorphism.
- Difficult to quantify value of feature that is glorified syntax sugar.
- We (manually!) rewrote programs to take advantage of AUTOMAP when we judged it improved readability.



## Practical impact: before

```
def main [nK] [nX]
  (kx: [nK]f32) (ky: [nK]f32) (kz: [nK]f32)
  (x: [nX]f32) (y: [nX]f32) (z: [nX]f32)
  (phiR: [nK]f32) (phiI: [nK]f32)
  : ([nX]f32, [nX]f32) =
  let phiM = map (\r i → r*r + i*i) phiR phiI
  let as = map (\x_e y_e z_e →
    map (2*pi*)
      (map (\kx_e ky_e kz_e →
        kx_e*x_e + ky_e*y_e + kz_e*z_e)
        kx ky kz))
    x y z
  let qr = map (\a → sum(map2 (*) phiM (map cos a))) as
  let qi = map (\a → sum(map2 (*) phiM (map sin a))) as
  in (qr, qi)
```




## Practical impact: after

```
def main [nK] [nX]
    (kx: [nK]f32) (ky: [nK]f32) (kz: [nK]f32)
    (x: [nX]f32) (y: [nX]f32) (z: [nX]f32)
    (phiR: [nK]f32) (phiI: [nK]f32)
    : ([nX]f32, [nX]f32) =
let phiM = phiR*phiR + phiI*phiI
let as = 2*pi*(kx*transpose (rep x)
    + ky*transpose (rep y)
    + kz*transpose (rep z))
let qr = sum (cos as * phiM)
let qi = sum (sin as * phiM)
in (qr, qi)
```



# Metrics from changing a benchmark suite

Proportion of ILP problems that have less than some given number of constraints.



Number of programs: 67

Lines of code: 8621  $\Rightarrow$  8515

Change in maps: 467  $\Rightarrow$  213

Largest ILP size: 28104 constraints

Median ILP size: 16 constraints

Mean ILP size: 116 constraints

Mean type checking slowdown: 2.50 $\times$



# Related work

- **Typed Remora:**

- ▶ Very general/powerful; binds shape variables in types:

$$\text{sum} : \forall S.S \text{ int} \rightarrow \text{int}$$

- ▶ Inference is very difficult.



# Related work

- **Typed Remora:**

- ▶ Very general/powerful; binds shape variables in types:

$$\text{sum} : \forall S.S \text{ int} \rightarrow \text{int}$$

- ▶ Inference is very difficult.

- **Naperian Functors** (Jeremy Gibbons):

- ▶ Cool rank polymorphism encoding in Haskell.
- ▶ Complicated function types (and potentially error messages).





# Related work

- **Typed Remora:**

- ▶ Very general/powerful; binds shape variables in types:

$$\text{sum} : \forall S.S \text{ int} \rightarrow \text{int}$$

- ▶ Inference is very difficult.

- **Naperian Functors** (Jeremy Gibbons):

- ▶ Cool rank polymorphism encoding in Haskell.
- ▶ Complicated function types (and potentially error messages).

- **Single-assignment C:**

- ▶ Has *rank specialization* where functions have specialized definitions depending on the rank of the input.
- ▶ No parametric polymorphism or higher-order functions.



# Summary

- AUTOMAP is a conservative extension of/compatible with a Hindley-Milner-style type system for array programming.



# Summary

- AUTOMAP is a conservative extension of/compatible with a Hindley-Milner-style type system for array programming.
- Anything inferred can also be inserted explicitly (much like classic type systems!)



# Summary

- AUTOMAP is a conservative extension of/compatible with a Hindley-Milner-style type system for array programming.
- Anything inferred can also be inserted explicitly (much like classic type systems!)
- Type checking based on some heavy machinery (ILP), but we suspect of a fairly simple kind.



# Summary

- AUTOMAP is a conservative extension of/compatible with a Hindley-Milner-style type system for array programming.
- Anything inferred can also be inserted explicitly (much like classic type systems!)
- Type checking based on some heavy machinery (ILP), but we suspect of a fairly simple kind.
- Implemented in Futhark, but not really production ready yet.
  - ▶ TODO: quality of type errors, type checking speed, better ambiguity checking.



# Final Things


- Check out Futhark: <https://futhark-lang.org>
  - ▶ There's a blog post on AUTOMAP that covers the AUTOMAP-portion of this talk in more detail.
  - ▶ The papers for each thing can also be found there, along with my PhD thesis.
- These slides and more about me at <https://rschenck.com>.




# Final Thanks

- Thanks for coming to my defense!
- Thanks to my advisors:
  - ▶ Fritz Henglein, Cosmin E. Oancea, and Troels Henriksen.
  - ▶ And everyone else at the PLTC/DIKU!
- Thanks to my committee:
  - ▶ Michael Kirkedal for being my committee chair.
  - ▶ Sven-Bodo Scholz and Paul Kelly (and for making the trip all the way here to Copenhagen).
- And thanks to all the hedgehog artists:
  - ▶ Nikolaj Hey Hinnerskov, Fillippa Biil, Lea Henriksen, Lys Sanz Moreta, the internet, and more.

Gotta go  
FAST



# That's all!



FUTHARK

The problem of being faster than light is that you can only live in darkness.



GO FAST.