

Andrea Arteaga

Bit-Reproducibility in HPC Applications




About Me

BSc and MSc in Computational Science and Engineering at ETH (2008-2013)

About Me



BSc and MSc in Computational Science and Engineering at ETH (2008-2013)



About Me

BSc and MSc in Computational Science and Engineering at ETH (2008-2013)


Oliver Fuhrer, MeteoSwiss




About Me

BSc and MSc in Computational Science and Engineering at ETH (2008-2013)

Oliver Fuhrer, MeteoSwiss




Torsten Höfler, ETH




About Me

BSc and MSc in Computational Science and Engineering at ETH (2008-2013)

Oliver Fuhrer, MeteoSwiss



Torsten Höfler, ETH



Piz Daint, CSCS

Floating-Point Algebra

Floating-Point Algebra

Floating-Point Operations

$$a \oplus b \quad := \quad \text{rd}(a + b)$$

$$a \ominus b \quad := \quad \text{rd}(a - b)$$

$$a \otimes b \quad := \quad \text{rd}(a \times b)$$

$$a \oslash b \quad := \quad \text{rd}(a \div b)$$

$$\text{sqrt}(a) \quad := \quad \text{rd}(\sqrt{a})$$

Floating-Point Algebra

a := 0.11100010100

b := 0.11100111011

c := 0.011000000101

exact: 10.001010100011

rounded: **10.001010100**

Floating-Point Algebra

```

a := 0.11100010100
b := 0.11100111011
c := 0.011000000101

```

$(a \oplus b) \oplus c$

```

exact: 10.001010100011
rounded: 10.001010100

```

```

a + b := 1.11001001111
rounded: 1.1100101000

+ c := 10.001010100101
rounded: 10.001010101

```

Floating-Point Algebra

```

a := 0.11100010100
b := 0.11100111011
c := 0.011000000101
  
```

$(a \oplus b) \oplus c$

$a \oplus (b \oplus c)$

```

exact: 10.001010100011
rounded: 10.001010100
  
```

```

a + b := 1.11001001111
rounded: 1.1100101000

+ c := 10.001010100101
rounded: 10.001010101
  
```

```

b + c := 1.010001111011
rounded: 1.0100011111

+ a := 10.00101010010
rounded: 10.001010100
  
```

Questions


1. How to implement bit-reproducible scalar functions?
2. How to deterministically reduce large arrays of numbers?

Questions

1. How to implement bit-reproducible scalar functions?
2. How to deterministically reduce large arrays of numbers?

Stencil functions

$$b_{i,j} = -4 \cdot a_{i,j} + a_{i-1,j} + a_{i+1,j} + a_{i,j-1} + a_{i,j+1}$$



```
for i = 1..N-1
  for j = 1..N-1
    b[i,j] = f(a, i, j)
  endfor
endfor
```


Questions

1. How to implement bit-reproducible scalar functions?

2. How to deterministically reduce large arrays of numbers?

Stencil functions

$$b_{i,j} = -4 \cdot a_{i,j} + a_{i-1,j} + a_{i+1,j} + a_{i,j-1} + a_{i,j+1}$$



```
for i = 1..N-1
  for j = 1..N-1
    b[i,j] = f(a, i, j)
  endfor
endfor
```

Reproducible summation

$$b = \sum_{i=0}^{N-1} a_i$$

```
float a[N];
float b_local = sum(a, N), b_global;
MPI_Reduce(b_local, b_global, 1, ...)
```


Reproducible Scalar Functions

$$b_{i,j} = -4 \cdot a_{i,j} + a_{i-1,j} + a_{i+1,j} + a_{i,j-1} + a_{i,j+1}$$

$$\theta_{i,j} = T_{i,j} \cdot \exp(-rocvp \cdot \log(p_{i,j} \cdot 10^{-5}))$$

Reproducible Scalar Functions

$$b_{i,j} = -4 \cdot a_{i,j} + a_{i-1,j} + a_{i+1,j} + a_{i,j-1} + a_{i,j+1}$$



$$b_{i,j} = (((-4 \cdot a_{i,j} + a_{i-1,j}) + a_{i+1,j}) + a_{i,j-1}) + a_{i,j+1}$$

Impose a fixed
execution order
(e.g. bracketing)

$$\theta_{i,j} = T_{i,j} \cdot \exp(-rocvp \cdot \log(p_{i,j} \cdot 10^{-5}))$$

Reproducible Scalar Functions

$$b_{i,j} = -4 \cdot a_{i,j} + a_{i-1,j} + a_{i+1,j} + a_{i,j-1} + a_{i,j+1}$$


$$b_{i,j} = ((((-4 \cdot a_{i,j} + a_{i-1,j})) + a_{i+1,j}) + a_{i,j-1}) + a_{i,j+1}$$

Impose a fixed
execution order
(e.g. bracketing)

$$\theta_{i,j} = T_{i,j} \cdot \exp(-rocvp \cdot \log(p_{i,j} \cdot 10^{-5}))$$

Reproducible Scalar Functions

$$b_{i,j} = -4 \cdot a_{i,j} + a_{i-1,j} + a_{i+1,j} + a_{i,j-1} + a_{i,j+1}$$

$$b_{i,j} = ((((-4 \cdot a_{i,j} + a_{i-1,j}) + a_{i+1,j}) + a_{i,j-1}) + a_{i,j+1})$$

Impose a fixed execution order (e.g. bracketing)

Fix the instruction set (e.g. FMA)

$$\theta_{i,j} = T_{i,j} \cdot \exp(-rocvp \cdot \log(p_{i,j} \cdot 10^{-5}))$$

Reproducible Scalar Functions

$$b_{i,j} = -4 \cdot a_{i,j} + a_{i-1,j} + a_{i+1,j} + a_{i,j-1} + a_{i,j+1}$$

$$b_{i,j} = ((((-4 \cdot a_{i,j} + a_{i-1,j}) + a_{i+1,j}) + a_{i,j-1}) + a_{i,j+1})$$

$$\theta_{i,j} = T_{i,j} \cdot \exp(-rocvp \cdot \log(p_{i,j} \cdot 10^{-5}))$$

Impose a fixed execution order (e.g. bracketing)

Fix the instruction set (e.g. FMA)

Use reproducible transcendental functions

Reproducible Summation: Error-Free Transform

$a := 10.110110011$

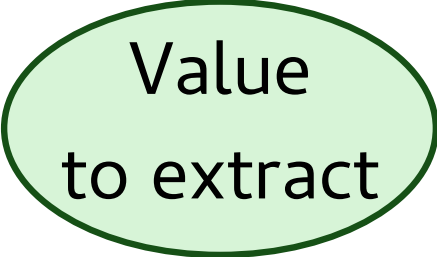
Value
to extract

Reproducible Summation: Error-Free Transform

A red oval containing the word "Extractor".

Extractor

$a := 10.110110011$
 $M := 1000.0000000$

A green oval containing the text "Value to extract".

Value
to extract

Reproducible Summation: Error-Free Transform

Extractor

$a := 10.110110011$
 $M := 1000.0000000$

Value
to extract

$a+M = 1010.110110011$
 $= 1010.1101101$

Reproducible Summation: Error-Free Transform

Extractor

$a := 10.110110011$

$M := 1000.0000000$

Value
to extract

$a+M = 1010.110110011$

$= 1010.1101101$

$(a+M)-M = 10.1101101$

Rounded
value

Reproducible Summation: Error-Free Transform

Extractor

$$\begin{aligned}
 a &:= 10.110110011 \\
 M &:= 1000.0000000
 \end{aligned}$$

Value
to extract

$$\begin{aligned}
 a+M &= 1010.110110011 \\
 &= 1010.1101101
 \end{aligned}$$

$$(a+M) - M = 10.1101101$$

Rounded
value

Remainder

$$a - ((a+M) - M) = -0.000000001$$

Reproducible Summation: Algorithm

a[0] 10.110110011

a[1] .10010010101

Reproducible Summation: Algorithm

M 10000.00000000


a[0] 10.110110011

a[1] .10010010101

Reproducible Summation: Algorithm

M	1 0000.00000000
a[0]	10.110110011
a[1]	.10010010101

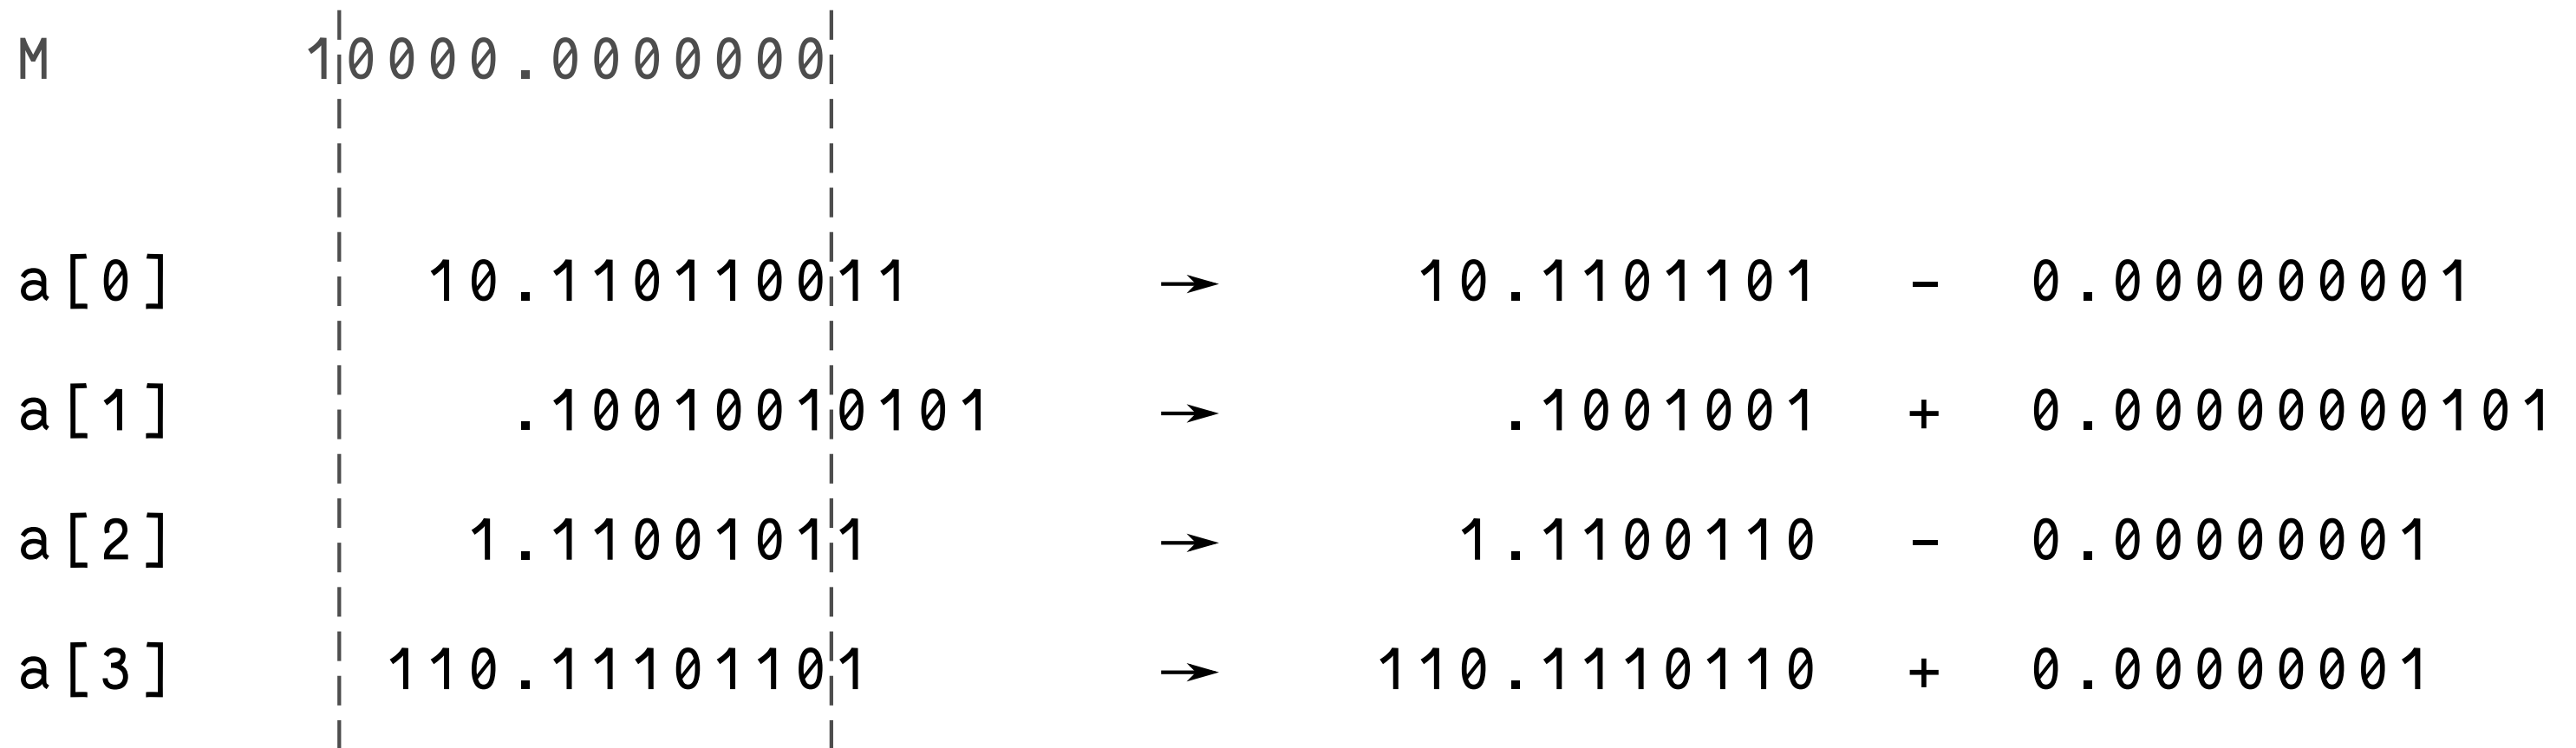
Reproducible Summation: Algorithm




Reproducible Summation: Algorithm

M	$1 0000.00000000 $				
a[0]	10.110110011	→	10.1101101	-	0.000000001
a[1]	$.10010010101$	→	$.1001001$	+	0.00000000101
a[2]	1.11001011	→	1.1100110	-	0.000000001

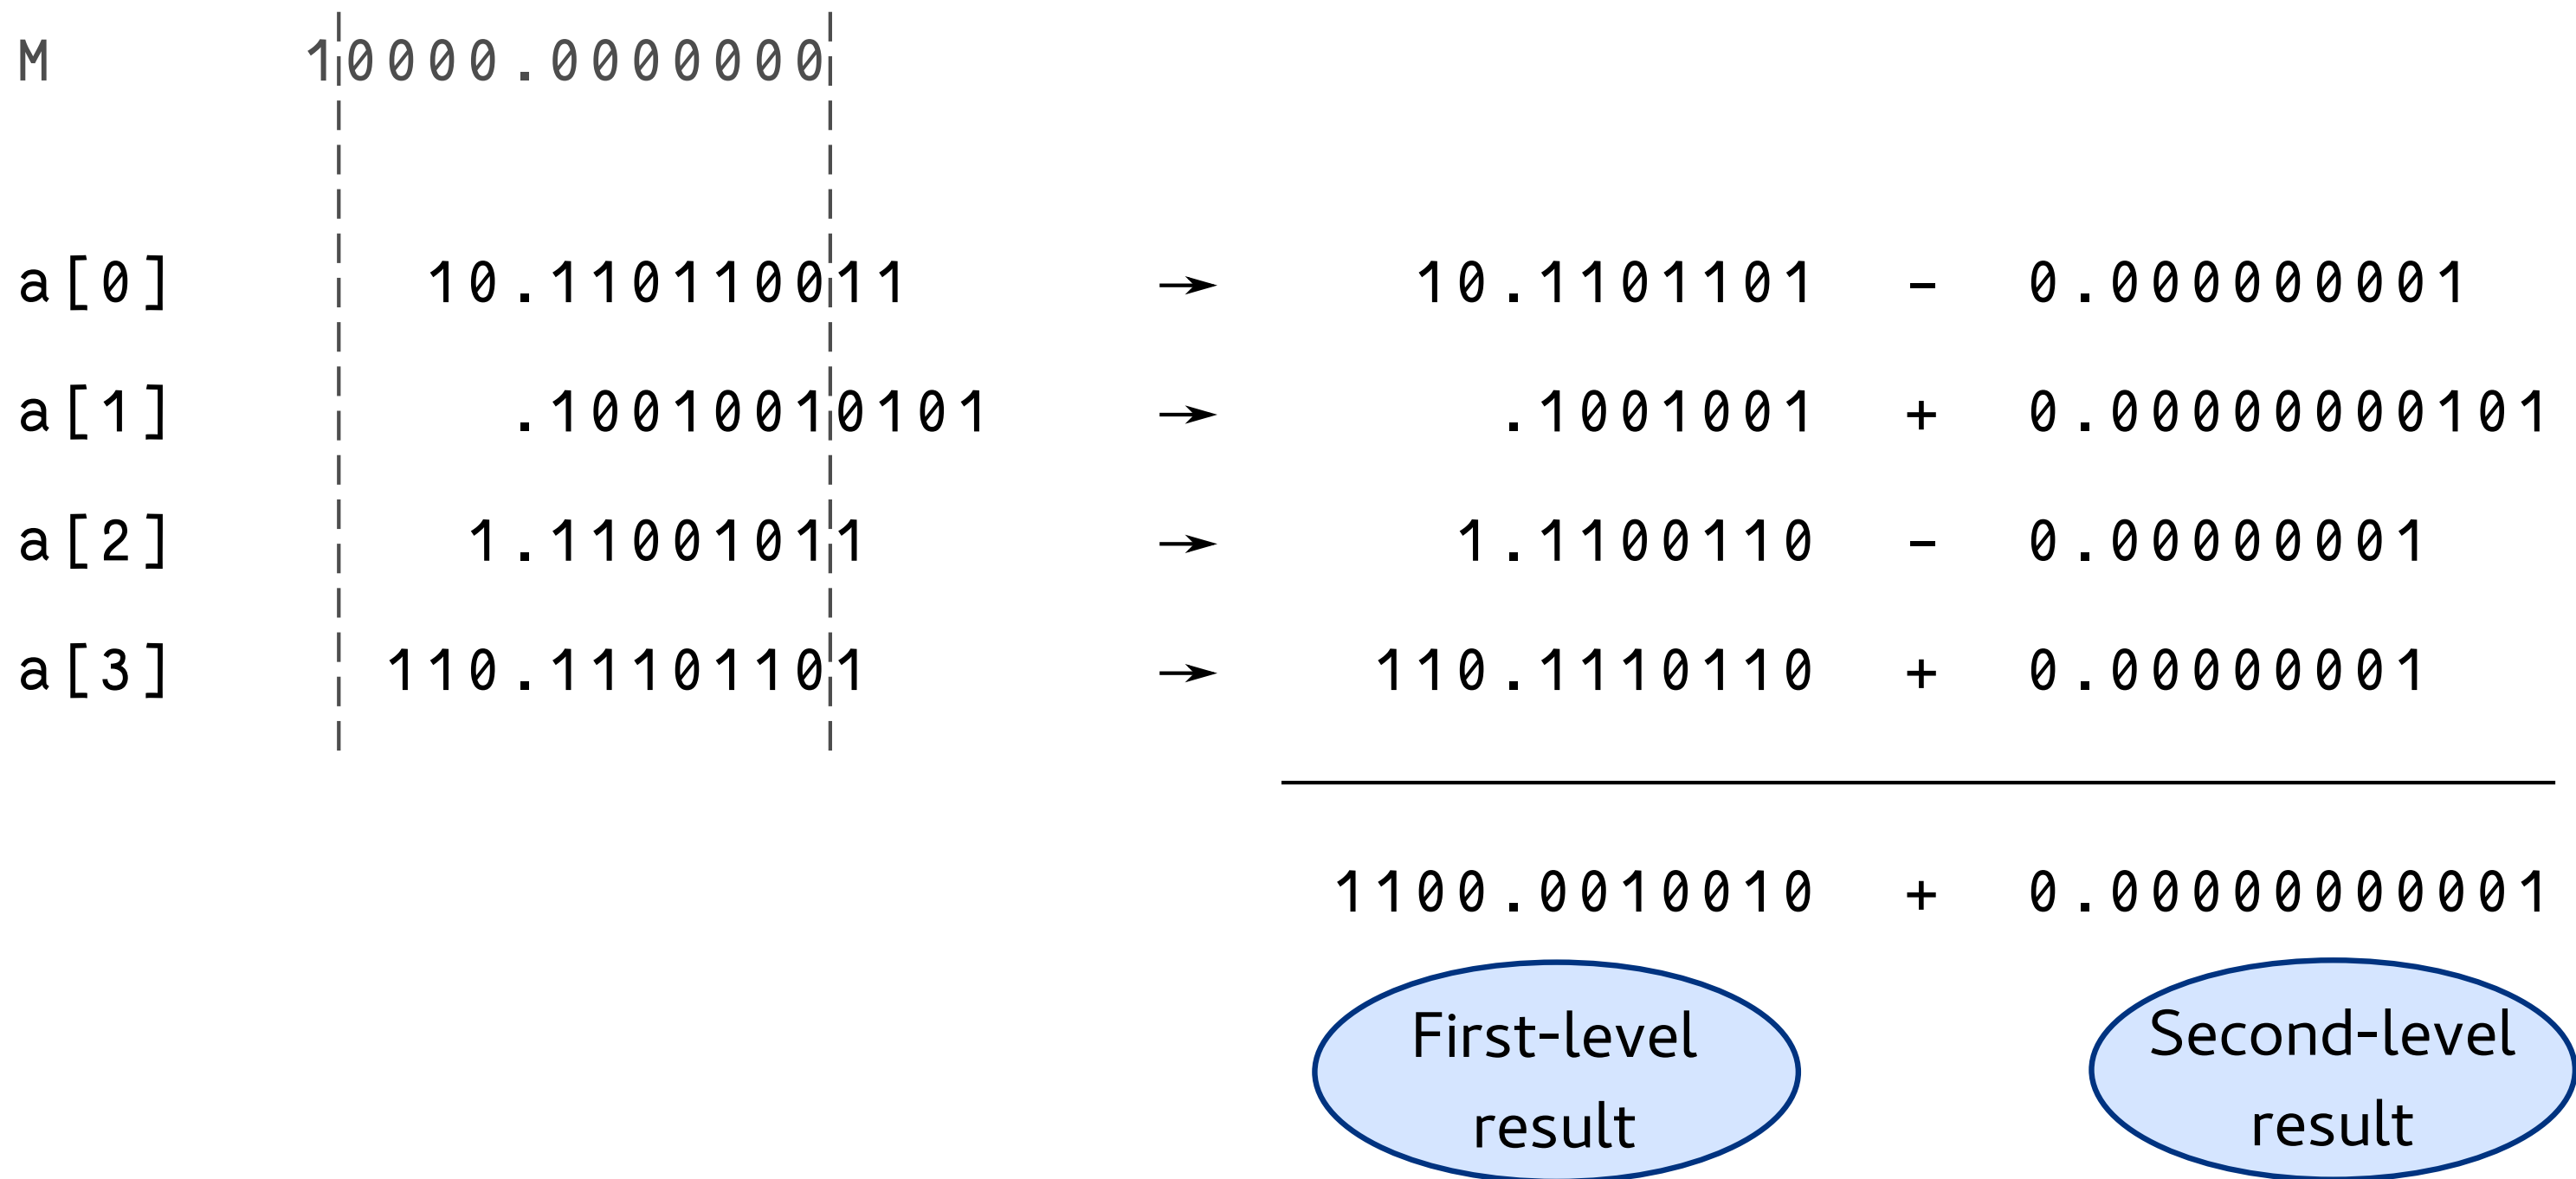
Reproducible Summation: Algorithm



Reproducible Summation: Algorithm



Reproducible Summation: Algorithm



Reproducible Summation: Performance

	Memory	Flops	Communication	Values per communication
Traditional	$N \cdot \text{sizeof}(\text{Type})$	N	$\log(p)$	1
Reproducible				

Reproducible Summation: Performance

	Memory	Flops	Communication	Values per communication
Traditional	$N \cdot \text{sizeof}(\text{Type})$	N	$\log(p)$	1
Reproducible	$N \cdot \text{sizeof}(\text{Type})$			

Reproducible Summation: Performance

	Memory	Flops	Communication	Values per communication
Traditional	$N \cdot \text{sizeof}(\text{Type})$	N	$\log(p)$	1
Reproducible	$N \cdot \text{sizeof}(\text{Type})$	$16 \cdot N$		

Reproducible Summation: Performance

	Memory	Flops	Communication	Values per communication
Traditional	$N \cdot \text{sizeof}(\text{Type})$	N	$\log(p)$	1
Reproducible	$N \cdot \text{sizeof}(\text{Type})$	$16 \cdot N$	$\log(p)$	

Reproducible Summation: Performance

	Memory	Flops	Communication	Values per communication
Traditional	$N \cdot \text{sizeof}(\text{Type})$	N	$\log(p)$	1
Reproducible	$N \cdot \text{sizeof}(\text{Type})$	$16 \cdot N$	$\log(p)$	3 - 4

Reproducible Summation: Performance

	Memory	Flops	Communication	Values per communication
Traditional	$N \cdot \text{sizeof}(\text{Type})$	N	$\log(p)$	1
Reproducible	$N \cdot \text{sizeof}(\text{Type})$	$16 \cdot N$	$\log(p)$	3 - 4

Memory-bound case: Good performance

Computation-bound case: Bad performance

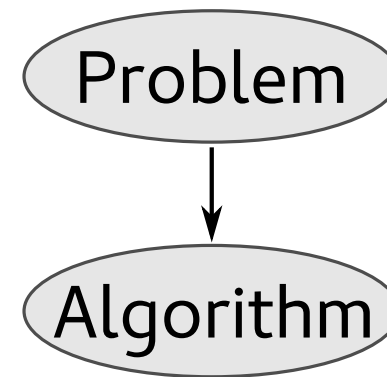
Communication-latency-bound case: Good performance

Thesis Project

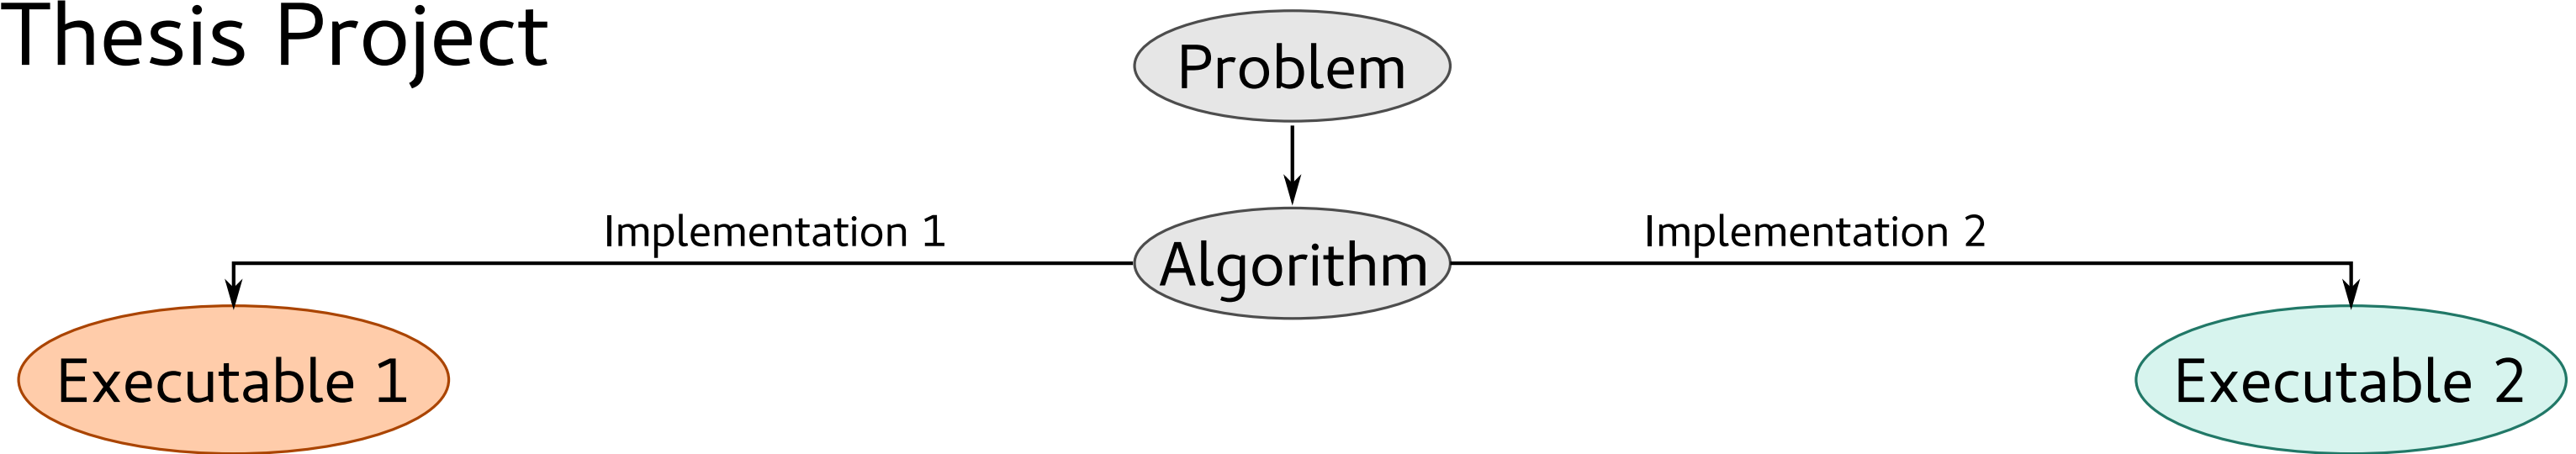
Thesis Project

Problem

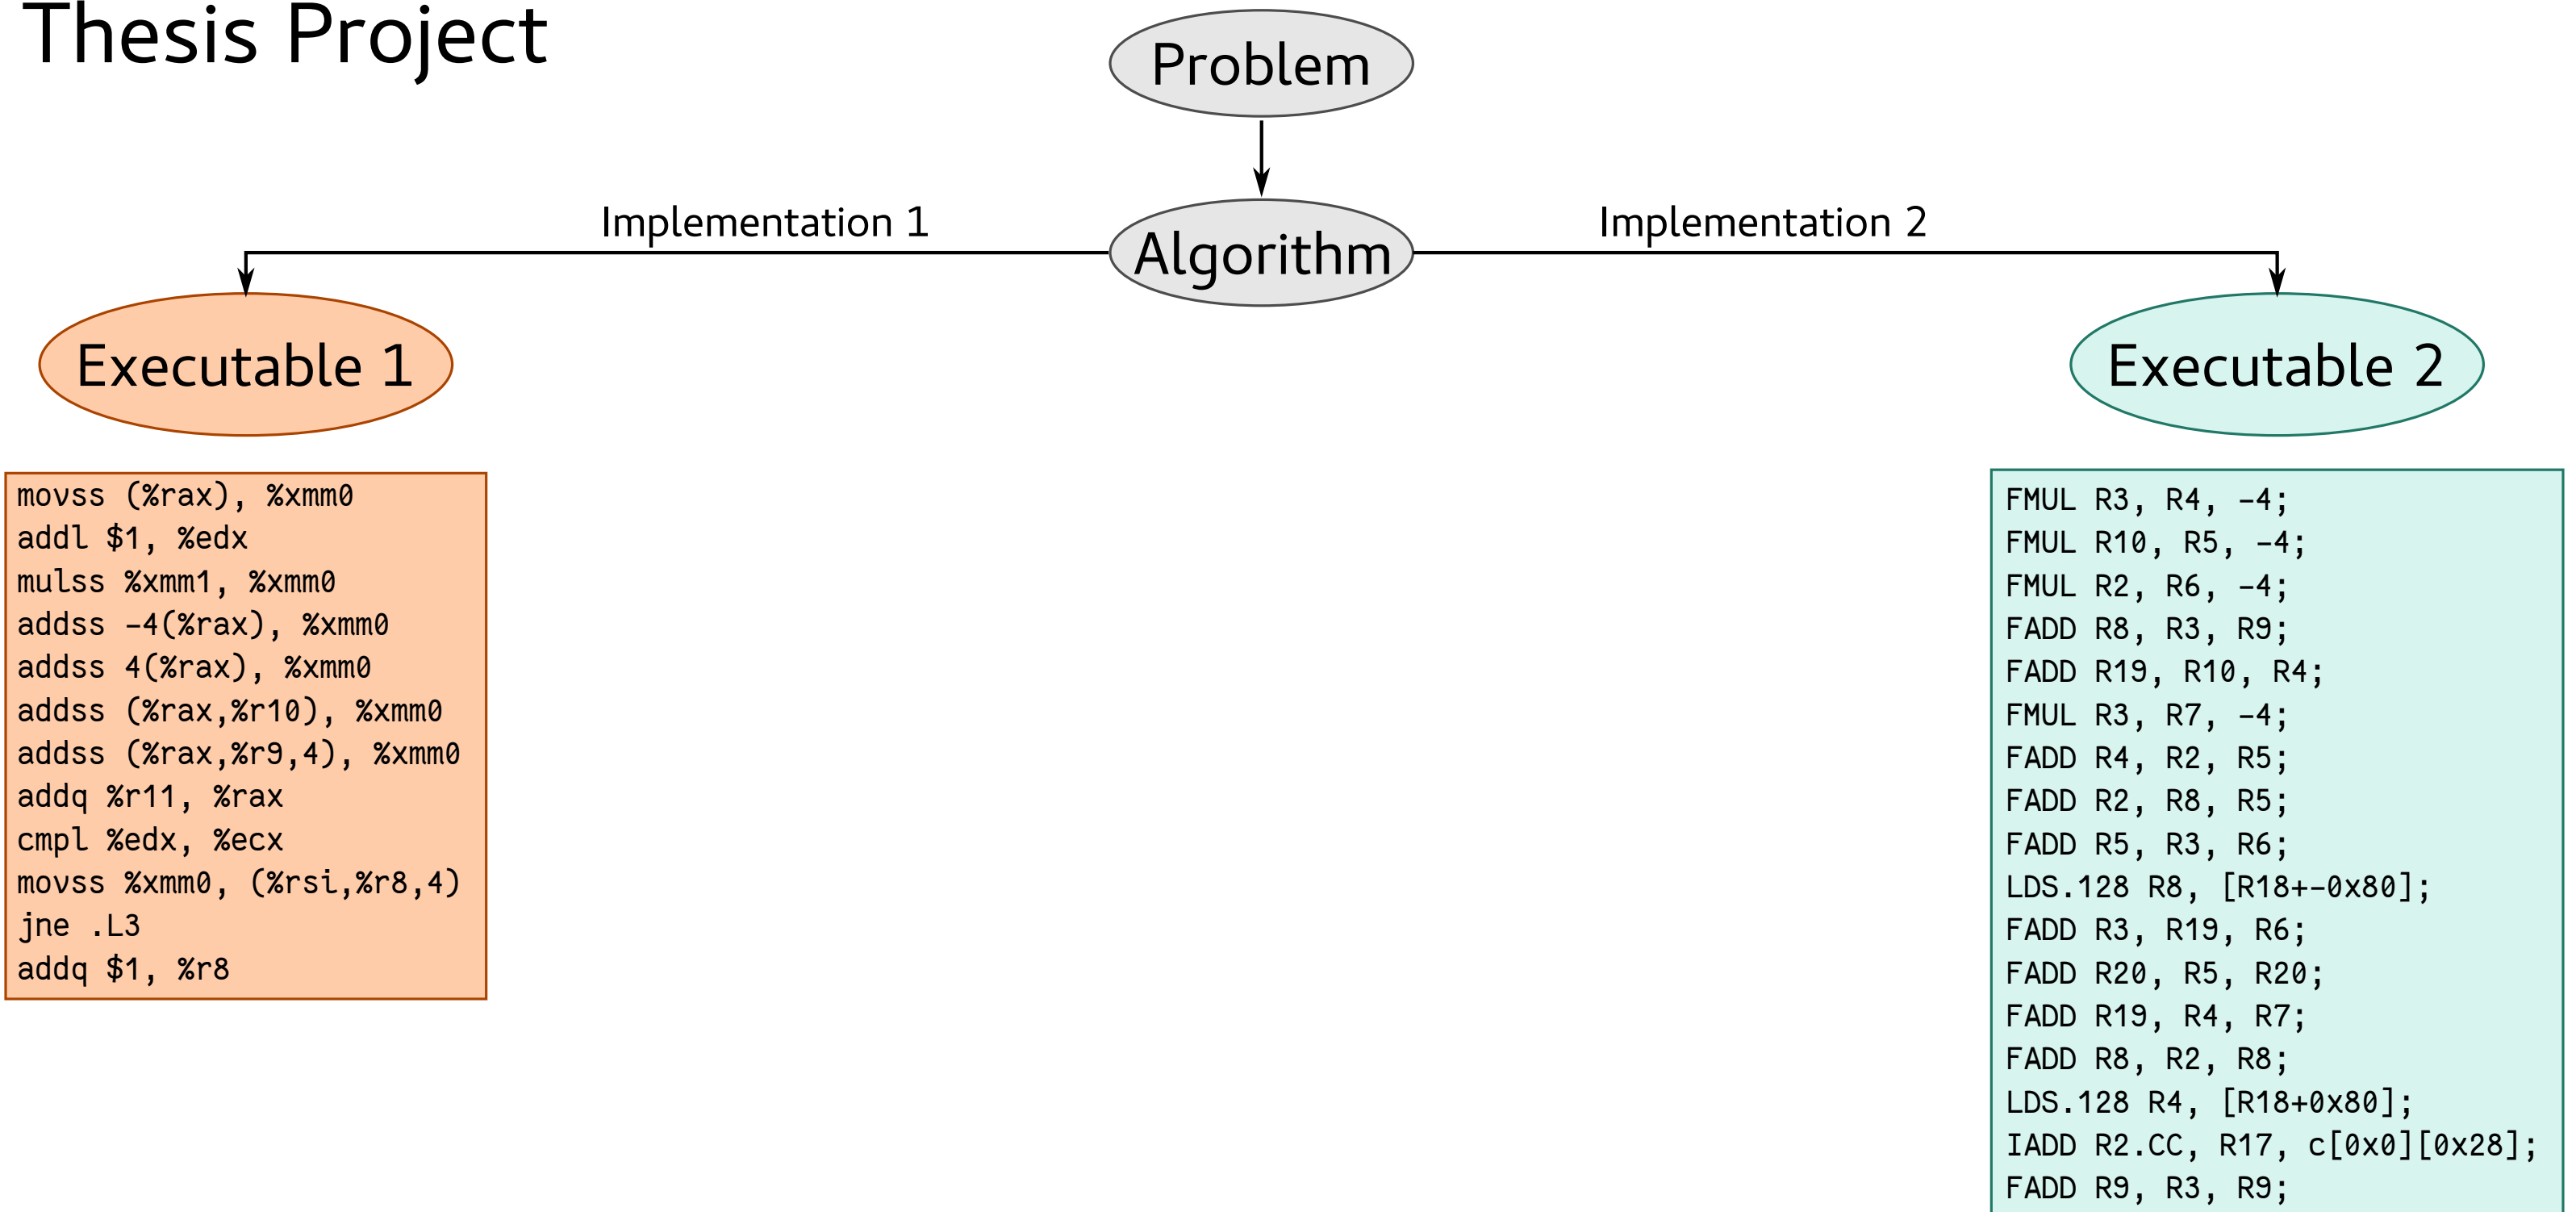
Thesis Project



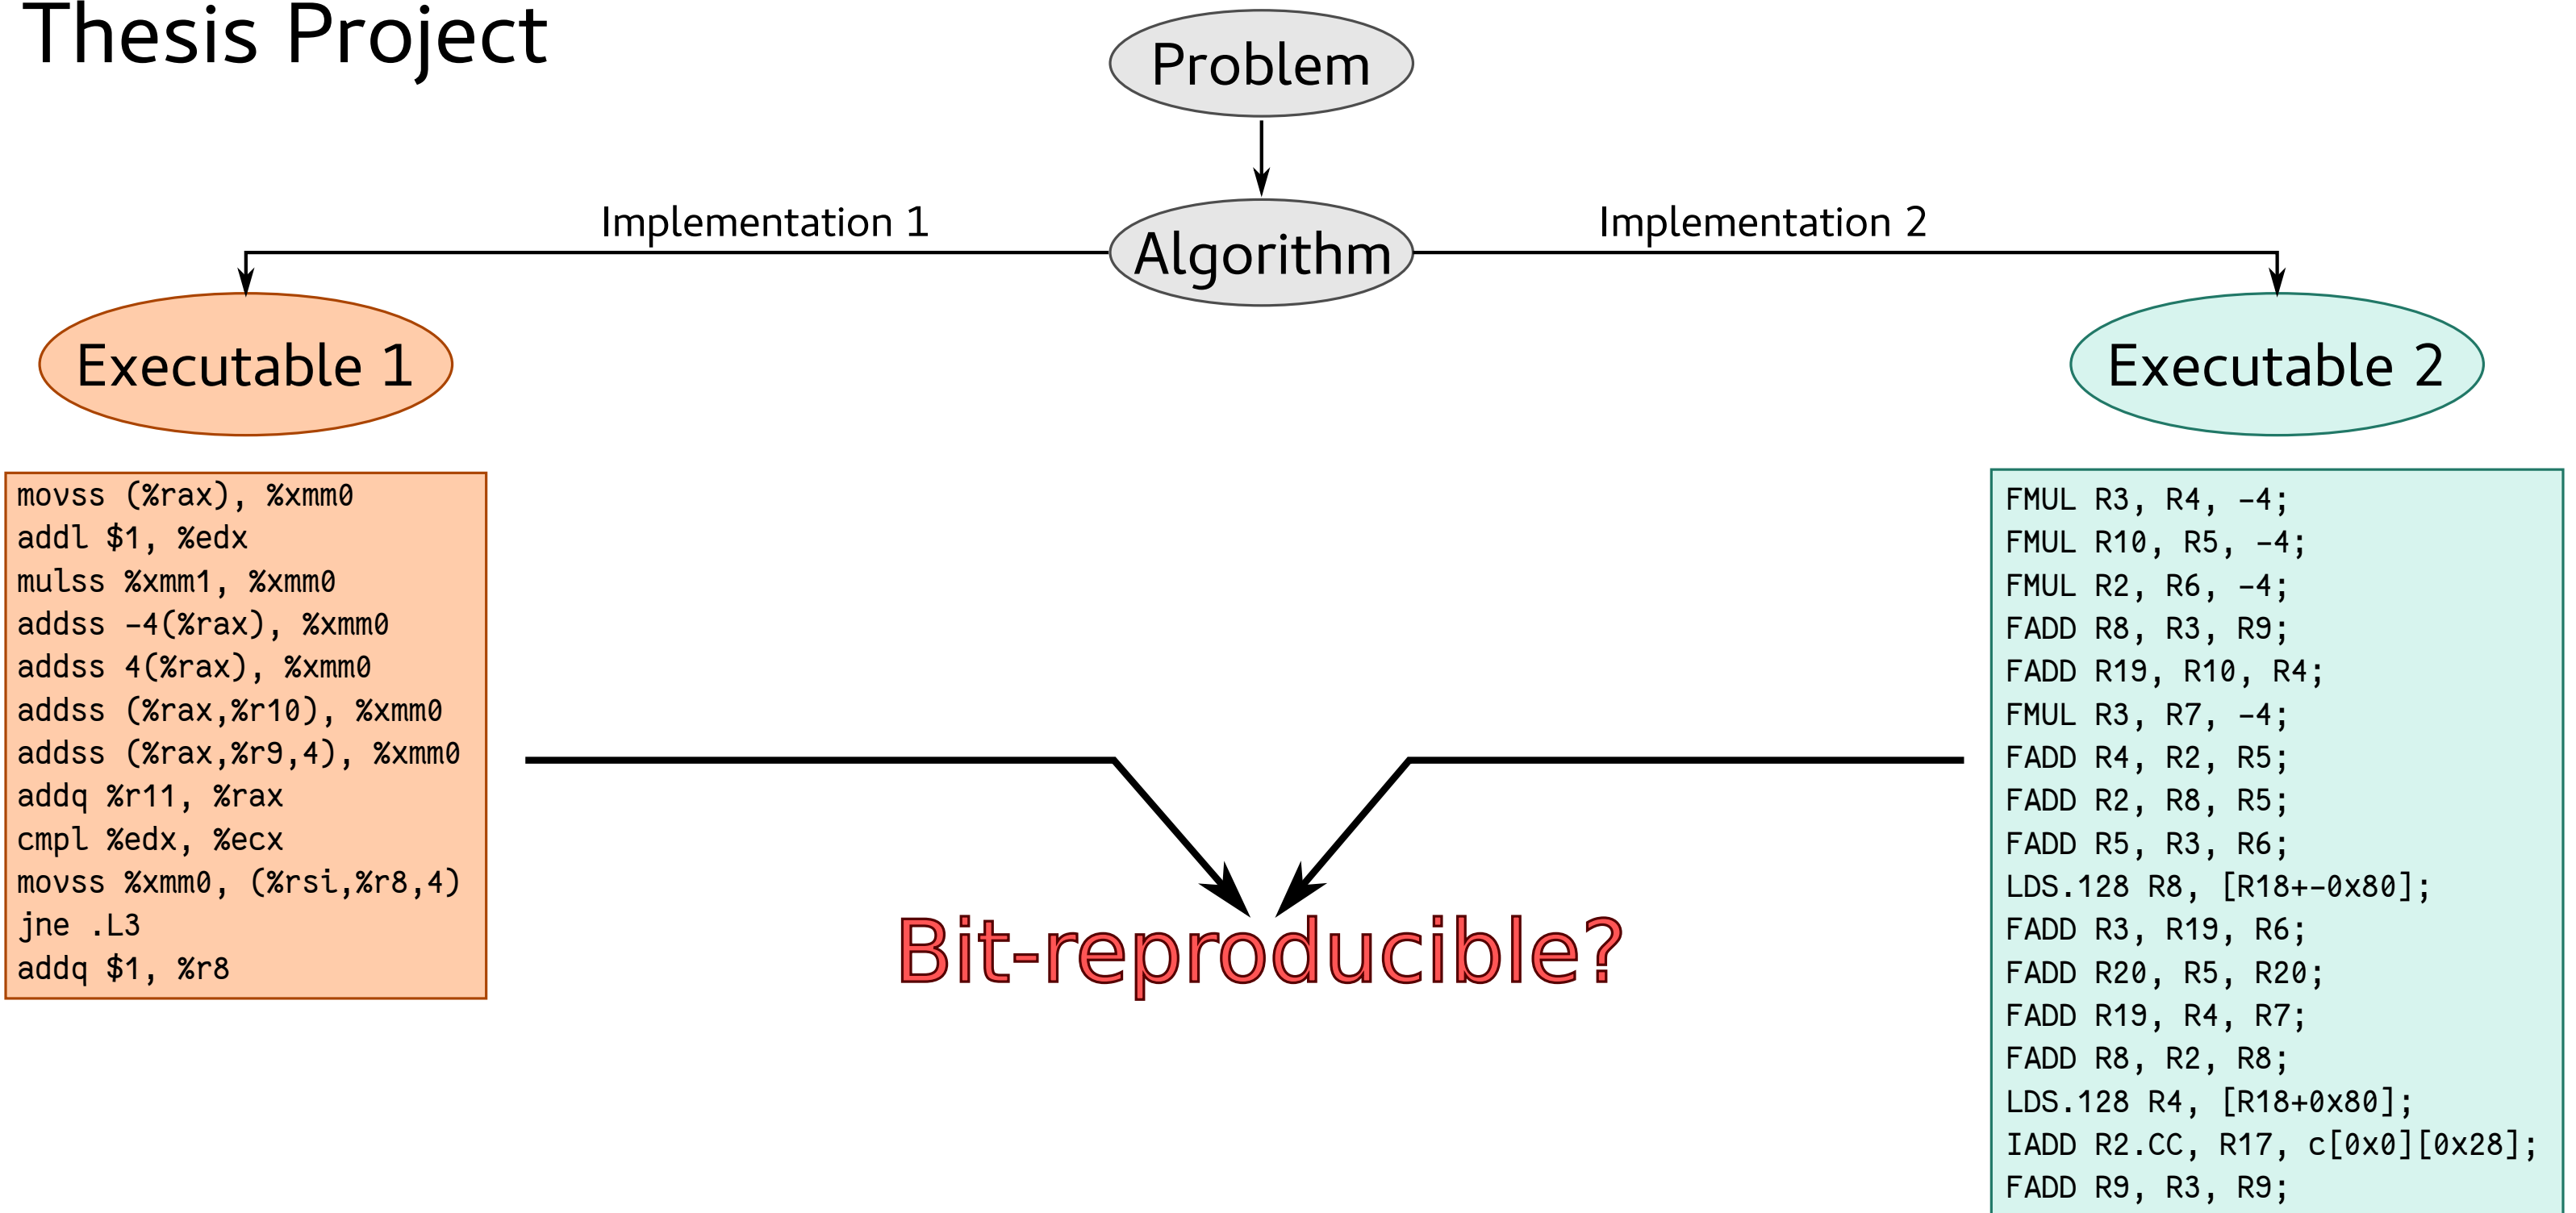
Thesis Project



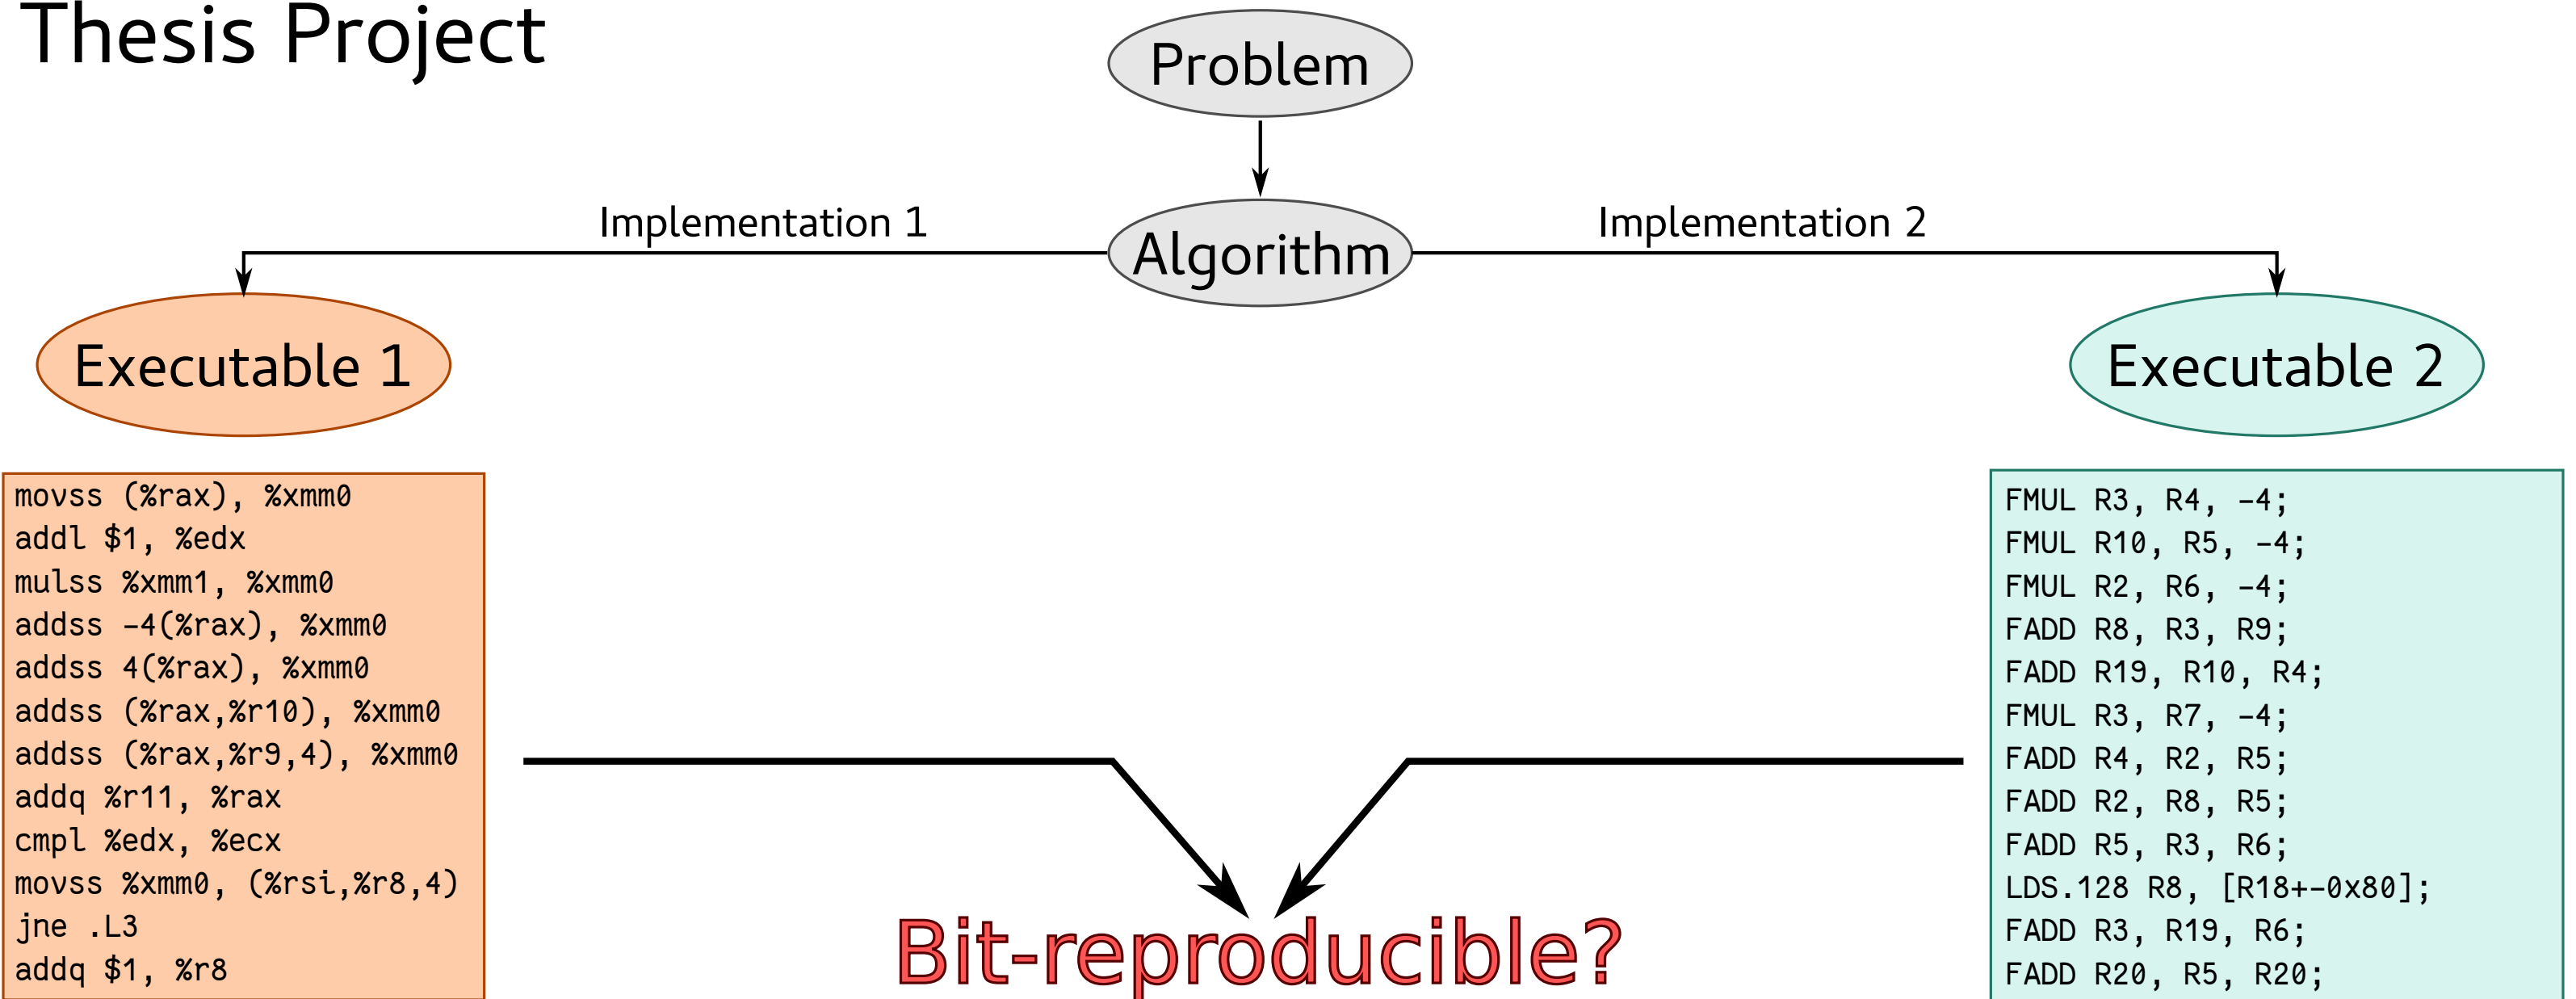
Thesis Project



Thesis Project



Thesis Project



Contact: andrea.arteaga@env.ethz.ch