

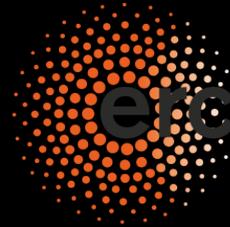
The JIL library (Justinline) for computations with SLPs

Joris van der Hoeven

Joint work with Albin AHLBÄCK, Ricardo BURING, Grégoire LECERF
CNRS, École polytechnique, France



Funded by
the European Union



European Research Council
Established by the European Commission

ODELIX thematic program, Palaiseau

December 8, 2025



Part I

Introduction and motivation

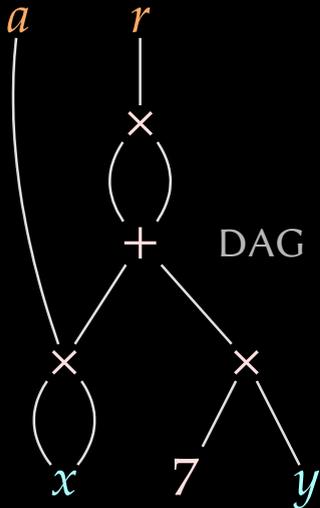
Example of an SLP (Straight Line Program)

3/27

```
in( $x, y$ )  
 $a := x \cdot x$   
 $b := 7 \cdot y$   
 $r := a + b$   
 $r := r \cdot r$   
out( $a, r$ )
```

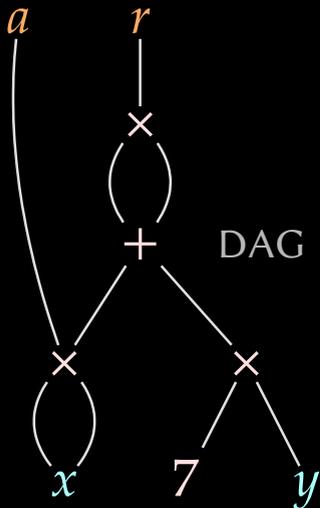
Example of an SLP (Straight Line Program)

$\text{in}(x, y)$
 $a := x \cdot x$
 $b := 7 \cdot y$
 $r := a + b$
 $r := r \cdot r$
 $\text{out}(a, r)$



Example of an SLP (Straight Line Program)

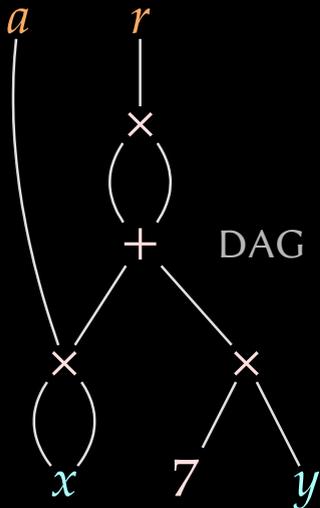
in(x, y)
 $a := x \cdot x$
 $b := 7 \cdot y$
 $r := a + b$
 $r := r \cdot r$
out(a, r)



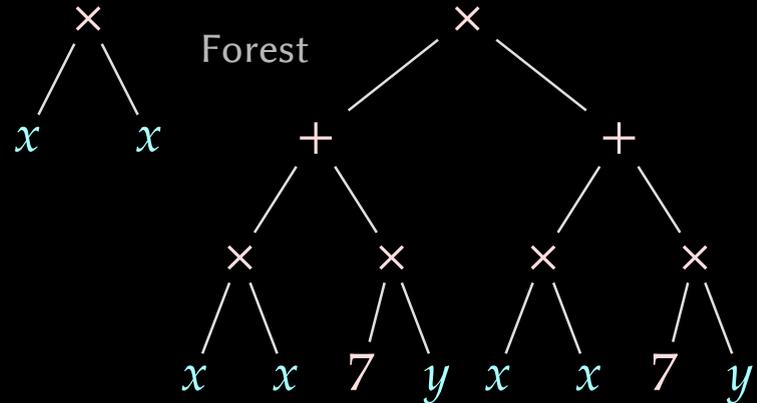
$(x^2, (x^2 + 7y)^2)$
Expression

Example of an SLP (Straight Line Program)

in(x, y)
 $a := x \cdot x$
 $b := 7 \cdot y$
 $r := a + b$
 $r := r \cdot r$
out(a, r)



$(x^2, (x^2 + 7y)^2)$
Expression



Theory

- Algebraic complexity analysis

Theory

- Algebraic complexity analysis
- More general than it seems → recording

Theory

- Algebraic complexity analysis
- More general than it seems → recording

Practice

- Evaluate same SLP many times → allows for massive parallelism

Theory

- Algebraic complexity analysis
- More general than it seems → recording

Practice

- Evaluate same SLP many times → allows for massive parallelism
- Trivial control → easier to compile and optimize

Theory

- Algebraic complexity analysis
- More general than it seems → recording

Practice

- Evaluate same SLP many times → allows for massive parallelism
- Trivial control → easier to compile and optimize
- Optimization techniques from both compilers and symbolic computation

Theory

- Algebraic complexity analysis
- More general than it seems → recording

Practice

- Evaluate same SLP many times → allows for massive parallelism
- Trivial control → easier to compile and optimize
- Optimization techniques from both compilers and symbolic computation

Applications

- Homotopy continuation

Theory

- Algebraic complexity analysis
- More general than it seems → recording

Practice

- Evaluate same SLP many times → allows for massive parallelism
- Trivial control → easier to compile and optimize
- Optimization techniques from both compilers and symbolic computation

Applications

- Homotopy continuation, geometric resolution

Theory

- Algebraic complexity analysis
- More general than it seems → recording

Practice

- Evaluate same SLP many times → allows for massive parallelism
- Trivial control → easier to compile and optimize
- Optimization techniques from both compilers and symbolic computation

Applications

- Homotopy continuation, geometric resolution, gradient descent

Theory

- Algebraic complexity analysis
- More general than it seems → recording

Practice

- Evaluate same SLP many times → allows for massive parallelism
- Trivial control → easier to compile and optimize
- Optimization techniques from both compilers and symbolic computation

Applications

- Homotopy continuation, geometric resolution, gradient descent
- Discrete Fourier Transforms

Theory

- Algebraic complexity analysis
- More general than it seems → recording

Practice

- Evaluate same SLP many times → allows for massive parallelism
- Trivial control → easier to compile and optimize
- Optimization techniques from both compilers and symbolic computation

Applications

- Homotopy continuation, geometric resolution, gradient descent
- Discrete Fourier Transforms, multi-precision arithmetic

Theory

- Algebraic complexity analysis
- More general than it seems → recording

Practice

- Evaluate same SLP many times → allows for massive parallelism
- Trivial control → easier to compile and optimize
- Optimization techniques from both compilers and symbolic computation

Applications

- Homotopy continuation, geometric resolution, gradient descent
- Discrete Fourier Transforms, multi-precision arithmetic, other “codelets”

Theory

- Algebraic complexity analysis
- More general than it seems → recording

Practice

- Evaluate same SLP many times → allows for massive parallelism
- Trivial control → easier to compile and optimize
- Optimization techniques from both compilers and symbolic computation

Applications

- Homotopy continuation, geometric resolution, gradient descent
- Discrete Fourier Transforms, multi-precision arithmetic, other “codelets”
- Solving ODEs via relaxed power series computations

Theory

- Bürgisser, Clausen, Shokrollahi: *Algebraic complexity theory*, 1997

Theory

- Bürgisser, Clausen, Shokrollahi: *Algebraic complexity theory*, 1997

Software

- Díaz, Kaltofen: *FoxBox*, 1998

Theory

- Bürgisser, Clausen, Shokrollahi: *Algebraic complexity theory*, 1997

Software

- Díaz, Kaltofen: *FoxBox*, 1998

Many *ad hoc* software for specific applications

- Geometric resolution → Aldaz, Castaño, Llovet, Martínez, Hägele, Bruno, Heintz, Matera, Giusti, Lecerf, Salvy, Durvy, ..., 2000–2008
- Codelets, DFTs → Frigo-Johnson, *FFTW3*, 1999, Püschel et al., *Spiral*, 2005
- Automatic differentiation → *Autodiff*, *Torch.autograd*, *Juliadiff*, ...

Theory

- Bürgisser, Clausen, Shokrollahi: *Algebraic complexity theory*, 1997

Software

- Díaz, Kaltofen: *FoxBox*, 1998

Many *ad hoc* software for specific applications

- Geometric resolution → Aldaz, Castaño, Llovet, Martínez, Hägele, Bruno, Heintz, Matera, Giusti, Lecerf, Salvy, Durvy, ..., 2000–2008
- Codelets, DFTs → Frigo-Johnson, *FFTW3*, 1999, Püschel et al., *Spiral*, 2005
- Automatic differentiation → *Autodiff*, *Torch.autograd*, *Juliadiff*, ...

More references

- vdH, Lecerf: *Towards a library for straight-line programs*, AAEECC, 2025

Towards a dedicated library

6/27

JUSTINLINE (in MATHEMAGIX)

(2015–2024)

JUSTINLINE (in MATHEMAGIX)

(2015–2024)

- High level source code and interface.

JUSTINLINE (in MATHEMAGIX)

(2015–2024)

- High level source code and interface.
- Reasonably fast on-the-fly compilation of SLPs.

JUSTINLINE (in MATHEMAGIX)

(2015–2024)

- High level source code and interface.
- Reasonably fast on-the-fly compilation of SLPs.
- Resulting machine code is fast.

JUSTINLINE (in MATHEMAGIX)

(2015–2024)

- High level source code and interface.
- Reasonably fast on-the-fly compilation of SLPs.
- Resulting machine code is fast.
- Suite of high level transformations and routines on SLPs.

JUSTINLINE (in MATHEMAGIX)

(2015–2024)

- High level source code and interface.
- Reasonably fast on-the-fly compilation of SLPs.
- Resulting machine code is fast.
- Suite of high level transformations and routines on SLPs.

JIL (in C++)

(2024–*)

JUSTINLINE (in MATHEMAGIX)

(2015–2024)

- High level source code and interface.
- Reasonably fast on-the-fly compilation of SLPs.
- Resulting machine code is fast.
- Suite of high level transformations and routines on SLPs.

JIL (in C++)

(2024–*)

- Stand-alone C++ library → lower threshold for developers and users.

JUSTINLINE (in MATHEMAGIX)

(2015–2024)

- High level source code and interface.
- Reasonably fast on-the-fly compilation of SLPs.
- Resulting machine code is fast.
- Suite of high level transformations and routines on SLPs.

JIL (in C++)

(2024–*)

- Stand-alone C++ library → lower threshold for developers and users.
- Low level source code, planned bindings for MATHEMAGIX, JULIA, FLINT, ...

JUSTINLINE (in MATHEMAGIX)

(2015–2024)

- High level source code and interface.
- Reasonably fast on-the-fly compilation of SLPs.
- Resulting machine code is fast.
- Suite of high level transformations and routines on SLPs.

JIL (in C++)

(2024–*)

- Stand-alone C++ library → lower threshold for developers and users.
- Low level source code, planned bindings for MATHEMAGIX, JULIA, FLINT, ...
- Very fast on-the-fly compilation of SLPs.

JUSTINLINE (in MATHEMAGIX)

(2015–2024)

- High level source code and interface.
- Reasonably fast on-the-fly compilation of SLPs.
- Resulting machine code is fast.
- Suite of high level transformations and routines on SLPs.

JIL (in C++)

(2024–*)

- Stand-alone C++ library → lower threshold for developers and users.
- Low level source code, planned bindings for MATHEMAGIX, JULIA, FLINT, ...
- Very fast on-the-fly compilation of SLPs.
- More architectures: X86, ARM, OPENCL, CUDA, SASS.

JUSTINLINE (in MATHEMAGIX)

(2015–2024)

- High level source code and interface.
- Reasonably fast on-the-fly compilation of SLPs.
- Resulting machine code is fast.
- Suite of high level transformations and routines on SLPs.

JIL (in C++)

(2024–*)

- Stand-alone C++ library → lower threshold for developers and users.
- Low level source code, planned bindings for MATHEMAGIX, JULIA, FLINT, ...
- Very fast on-the-fly compilation of SLPs.
- More architectures: X86, ARM, OPENCL, CUDA, SASS.
- Beyond SLPs...?

Motivation and spoiler

7/27

	sols	JUSTINLINE		JIL		GPU	
		jit	exe ₈	jit	exe ₈	jit	exe ₄₀₉₆
Katsura ₆	64	1.11 s	3 ms	7.13 ms	39.1 ms	1.00 s	68.9 ms
Katsura ₈	256	2.17 s	10 ms	13.6 ms	45.5 ms	3.21 s	0.42 s
Katsura ₁₀	1024	3.74 s	37 ms	23.8 ms	0.20 s	7.58 s	0.72 s
Katsura ₁₂	4096	6.16 s	0.23 s	42.4 ms	0.46 s	15.3 s	1.22 s
Katsura ₁₄	16384	9.34 s	1.59 s	90.1 ms	1.94 s	29.4 s	5.53 s
Katsura ₁₆	65536	13.4 s	11.3 s	0.27 s	11.9 s	49.6 s	29.1 s
Katsura ₁₈	262144	21.3 s	145 s	0.95 s	83.2 s	82.3 s	157 s
Katsura ₂₀	1048576	45.1 s	824 s	3.95 s	512 s	124 s	911 s
Posso _{3,3}	27	0.35 s	<1 ms	3.05 ms	0.26 ms	0.49 s	5.98 ms
Posso _{4,4}	256	1.50 s	3 ms	11.9 ms	1.5 ms	1.76 s	0.12 s
Posso _{5,5}	3125	8.22 s	78 ms	52.4 ms	0.31 s	19.1 s	0.96 s

INTEL XEON 3,2 GHz, AVX2 (JUSTINLINE), AVX512 (JIL), 8 threads
NVIDIA GEFORCE RTX 4070 SUPER GPU, OPENCL, 4096 threads
2nd order stepper (JUSTINLINE), 1st order stepper (JIL)

CPU

- Out of order execution
- Cache hierarchies
- Up to 8 or 16 wide SIMD parallelism
- Up to 100 cores

CPU

- Out of order execution
- Cache hierarchies
- Up to 8 or 16 wide SIMD parallelism
- Up to 100 cores

GPU

- Up to 10000 cores
- Often limited to 32 bits
- Large transfer times with CPU
- Hard to program and processors mostly undocumented

CPU

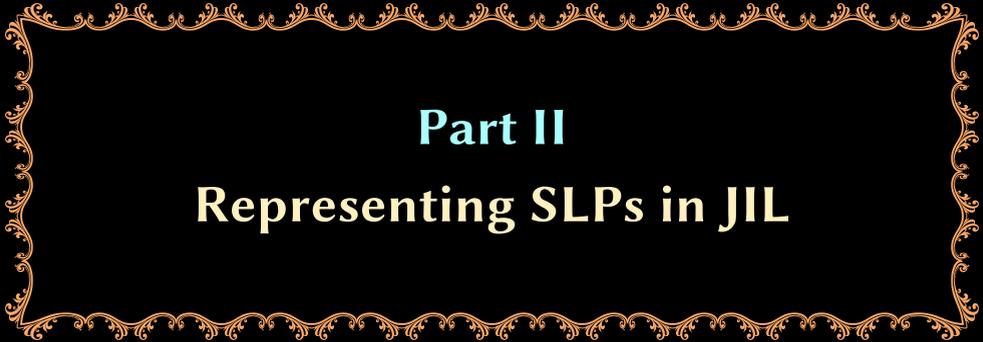
- Out of order execution
- Cache hierarchies
- Up to 8 or 16 wide SIMD parallelism
- Up to 100 cores

GPU

- Up to 10000 cores
- Often limited to 32 bits
- Large transfer times with CPU
- Hard to program and processors mostly undocumented

TPU

- Only matrix multiplication, but about 8 times faster than SIMD in theory



Part II

Representing SLPs in JIL

in(x, y)

$a := x \cdot x$

$b := 7 \cdot y$

$r := a + b$

$r := r \cdot r$

out(a, r)

$\text{in}(x, y)$

$a := x \cdot x$

$b := 7 \cdot y$

$r := a + b$

$r := r \cdot r$

$\text{out}(a, r)$

$\text{in}(x_0, x_1)$

$x_3 := x_0 \cdot x_0$

$x_4 := x_2 \cdot x_1$

$x_5 := x_3 + x_4$

$x_5 := x_5 \cdot x_5$

$\text{out}(x_3, x_5)$

Low level representation of SLPs

in(x, y)
 $a := x \cdot x$
 $b := 7 \cdot y$
 $r := a + b$
 $r := r \cdot r$
out(a, r)

in(x_0, x_1)
 $x_3 := x_0 \cdot x_0$
 $x_4 := x_2 \cdot x_1$
 $x_5 := x_3 + x_4$
 $x_5 := x_5 \cdot x_5$
out(x_3, x_5)

in

0	1
---	---

out

3	5
---	---

prg

×	3	0	0	×	4	2	1
+	5	3	4	+	5	5	5

data

0	0	7	0	0	0
---	---	---	---	---	---

Low level representation of SLPs

$\text{in}(x, y)$
 $a := x \cdot x$
 $b := 7 \cdot y$
 $r := a + b$
 $r := r \cdot r$
 $\text{out}(a, r)$

$\text{in}(x_0, x_1)$
 $x_3 := x_0 \cdot x_0$
 $x_4 := x_2 \cdot x_1$
 $x_5 := x_3 + x_4$
 $x_5 := x_5 \cdot x_5$
 $\text{out}(x_3, x_5)$

in

0	1
---	---

 out

3	5
---	---

 prg

\times	3	0	0	\times	4	2	1
$+$	5	3	4	$+$	5	5	5

 data

0	0	7	0	0	0
---	---	---	---	---	---

in , out , prg : arrays of 32 bit integers

Low level representation of SLPs

$\text{in}(x, y)$
 $a := x \cdot x$
 $b := 7 \cdot y$
 $r := a + b$
 $r := r \cdot r$
 $\text{out}(a, r)$

$\text{in}(x_0, x_1)$
 $x_3 := x_0 \cdot x_0$
 $x_4 := x_2 \cdot x_1$
 $x_5 := x_3 + x_4$
 $x_5 := x_5 \cdot x_5$
 $\text{out}(x_3, x_5)$

in

0	1
---	---

out

3	5
---	---

prg

×	3	0	0	×	4	2	1
+	5	3	4	+	5	5	5

data

0	0	7	0	0	0
---	---	---	---	---	---

in , out , prg : arrays of 32 bit integers

\Rightarrow operations $+$, \times , ... encoded as 32 bit integers

Low level representation of SLPs

10/27

$\text{in}(x, y)$
 $a := x \cdot x$
 $b := 7 \cdot y$
 $r := a + b$
 $r := r \cdot r$
 $\text{out}(a, r)$

$\text{in}(x_0, x_1)$
 $x_3 := x_0 \cdot x_0$
 $x_4 := x_2 \cdot x_1$
 $x_5 := x_3 + x_4$
 $x_5 := x_5 \cdot x_5$
 $\text{out}(x_3, x_5)$

in

0	1
---	---

 out

3	5
---	---

 prg

\times	3	0	0	\times	4	2	1
$+$	5	3	4	$+$	5	5	5

 data

0	0	7	0	0	0
---	---	---	---	---	---

in , out , prg : arrays of 32 bit integers

\Rightarrow operations $+$, \times , ... encoded as 32 bit integers

data : array of elements in some “scalar domain” \mathbb{K}

Hardware domains

- $\mathbb{Z}_8, \mathbb{Z}_{16}, \mathbb{Z}_{32}, \mathbb{Z}_{64}, \mathbb{N}_8, \mathbb{N}_{16}, \mathbb{N}_{32}, \mathbb{N}_{64}, \mathbb{R}_{32}, \mathbb{R}_{64}$

Hardware domains

- $\mathbb{Z}_8, \mathbb{Z}_{16}, \mathbb{Z}_{32}, \mathbb{Z}_{64}, \mathbb{N}_8, \mathbb{N}_{16}, \mathbb{N}_{32}, \mathbb{N}_{64}, \mathbb{R}_{32}, \mathbb{R}_{64}$
- SIMD variants, e.g. $\mathbb{Z}_8^{64}, \mathbb{Z}_{16}^{32}, \mathbb{Z}_{32}^{16}, \mathbb{Z}_{64}^8, \mathbb{N}_8^{64}, \mathbb{N}_{16}^{32}, \mathbb{N}_{32}^{16}, \mathbb{N}_{64}^8, \mathbb{R}_{32}^{16}, \mathbb{R}_{64}^8$

Hardware domains

- $\mathbb{Z}_8, \mathbb{Z}_{16}, \mathbb{Z}_{32}, \mathbb{Z}_{64}, \mathbb{N}_8, \mathbb{N}_{16}, \mathbb{N}_{32}, \mathbb{N}_{64}, \mathbb{R}_{32}, \mathbb{R}_{64}$
- SIMD variants, e.g. $\mathbb{Z}_8^{64}, \mathbb{Z}_{16}^{32}, \mathbb{Z}_{32}^{16}, \mathbb{Z}_{64}^8, \mathbb{N}_8^{64}, \mathbb{N}_{16}^{32}, \mathbb{N}_{32}^{16}, \mathbb{N}_{64}^8, \mathbb{R}_{32}^{16}, \mathbb{R}_{64}^8$

Software domains

- $\mathbb{K}[i], \text{Ball}(\mathbb{K}_{\text{cen}}, \mathbb{K}_{\text{rad}}), \dots$

Hardware domains

- $\mathbb{Z}_8, \mathbb{Z}_{16}, \mathbb{Z}_{32}, \mathbb{Z}_{64}, \mathbb{N}_8, \mathbb{N}_{16}, \mathbb{N}_{32}, \mathbb{N}_{64}, \mathbb{R}_{32}, \mathbb{R}_{64}$
- SIMD variants, e.g. $\mathbb{Z}_8^{64}, \mathbb{Z}_{16}^{32}, \mathbb{Z}_{32}^{16}, \mathbb{Z}_{64}^8, \mathbb{N}_8^{64}, \mathbb{N}_{16}^{32}, \mathbb{N}_{32}^{16}, \mathbb{N}_{64}^8, \mathbb{R}_{32}^{16}, \mathbb{R}_{64}^8$

Software domains

- $\mathbb{K}[i], \text{Ball}(\mathbb{K}_{\text{cen}}, \mathbb{K}_{\text{rad}}), \dots$
- $\mathbb{K}^n, \mathbb{K}^{r \times c}, \mathbb{K}[x] / (x^n), \dots$

Hardware domains

- $\mathbb{Z}_8, \mathbb{Z}_{16}, \mathbb{Z}_{32}, \mathbb{Z}_{64}, \mathbb{N}_8, \mathbb{N}_{16}, \mathbb{N}_{32}, \mathbb{N}_{64}, \mathbb{R}_{32}, \mathbb{R}_{64}$
- SIMD variants, e.g. $\mathbb{Z}_8^{64}, \mathbb{Z}_{16}^{32}, \mathbb{Z}_{32}^{16}, \mathbb{Z}_{64}^8, \mathbb{N}_8^{64}, \mathbb{N}_{16}^{32}, \mathbb{N}_{32}^{16}, \mathbb{N}_{64}^8, \mathbb{R}_{32}^{16}, \mathbb{R}_{64}^8$

Software domains

- $\mathbb{K}[i], \text{Ball}(\mathbb{K}_{\text{cen}}, \mathbb{K}_{\text{rad}}), \dots$
- $\mathbb{K}^n, \mathbb{K}^{r \times c}, \mathbb{K}[x] / (x^n), \dots$
- $\mathbb{Z}_{64} / (m \mathbb{Z}_{64}), \dots$

Hardware domains

- $\mathbb{Z}_8, \mathbb{Z}_{16}, \mathbb{Z}_{32}, \mathbb{Z}_{64}, \mathbb{N}_8, \mathbb{N}_{16}, \mathbb{N}_{32}, \mathbb{N}_{64}, \mathbb{R}_{32}, \mathbb{R}_{64}$
- SIMD variants, e.g. $\mathbb{Z}_8^{64}, \mathbb{Z}_{16}^{32}, \mathbb{Z}_{32}^{16}, \mathbb{Z}_{64}^8, \mathbb{N}_8^{64}, \mathbb{N}_{16}^{32}, \mathbb{N}_{32}^{16}, \mathbb{N}_{64}^8, \mathbb{R}_{32}^{16}, \mathbb{R}_{64}^8$

Software domains

- $\mathbb{K}[i], \text{Ball}(\mathbb{K}_{\text{cen}}, \mathbb{K}_{\text{rad}}), \dots$
- $\mathbb{K}^n, \mathbb{K}^{r \times c}, \mathbb{K}[x] / (x^n), \dots$
- $\mathbb{Z}_{64} / (m \mathbb{Z}_{64}), \dots$
- $\text{Recorder}(\mathbb{K})$

Signature Σ : the supported operations $\sigma: \mathbb{K}^{|\sigma|} \rightarrow \mathbb{K}$

$$\Sigma := \Sigma_{\text{basic}} \cup \Sigma_{\text{ext}}$$

Signature Σ : the supported operations $\sigma: \mathbb{K}^{|\sigma|} \rightarrow \mathbb{K}$

$$\Sigma := \Sigma_{\text{basic}} \cup \Sigma_{\text{ext}}$$

Σ_{basic} : operations that most domains \mathbb{K} supports

$$\Sigma_{\text{basic}} \supseteq \{\text{move, neg, add, sub, sqr, mul, nmul, fma, fnma, \dots}\}$$

Signature Σ : the supported operations $\sigma: \mathbb{K}^{|\sigma|} \rightarrow \mathbb{K}$

$$\Sigma := \Sigma_{\text{basic}} \cup \Sigma_{\text{ext}}$$

Σ_{basic} : operations that most domains \mathbb{K} supports

$$\Sigma_{\text{basic}} \supseteq \{\text{move, neg, add, sub, sqr, mul, nmul, fma, fnma, \dots}\} \\ \cup \{\text{inv, div, \dots}\}$$

Signature Σ : the supported operations $\sigma: \mathbb{K}^{|\sigma|} \rightarrow \mathbb{K}$

$$\Sigma := \Sigma_{\text{basic}} \cup \Sigma_{\text{ext}}$$

Σ_{basic} : operations that most domains \mathbb{K} supports

$$\begin{aligned} \Sigma_{\text{basic}} \supseteq & \{ \text{move, neg, add, sub, sqr, mul, nmul, fma, fnma, \dots} \} \\ & \cup \{ \text{inv, div, \dots} \} \\ & \cup \{ \text{min, max, abs, \dots} \} \cup \{ \text{floor, ceil, round, trunc} \} \end{aligned}$$

Signature Σ : the supported operations $\sigma: \mathbb{K}^{|\sigma|} \rightarrow \mathbb{K}$

$$\Sigma := \Sigma_{\text{basic}} \cup \Sigma_{\text{ext}}$$

Σ_{basic} : operations that most domains \mathbb{K} supports

$$\begin{aligned} \Sigma_{\text{basic}} \supseteq & \{ \text{move, neg, add, sub, sqr, mul, nmul, fma, fnma, \dots} \} \\ & \cup \{ \text{inv, div, \dots} \} \\ & \cup \{ \text{min, max, abs, \dots} \} \cup \{ \text{floor, ceil, round, trunc} \} \\ & \cup \{ \text{not, or, xor, and} \} \cup \{ \text{eq, neq, lt, le, gt, ge} \} \end{aligned}$$

Signature Σ : the supported operations $\sigma: \mathbb{K}^{|\sigma|} \rightarrow \mathbb{K}$

$$\Sigma := \Sigma_{\text{basic}} \cup \Sigma_{\text{ext}}$$

Σ_{basic} : operations that most domains \mathbb{K} supports

$$\begin{aligned} \Sigma_{\text{basic}} \supseteq & \{ \text{move, neg, add, sub, sqr, mul, nmul, fma, fnma, \dots} \} \\ & \cup \{ \text{inv, div, \dots} \} \\ & \cup \{ \text{min, max, abs, \dots} \} \cup \{ \text{floor, ceil, round, trunc} \} \\ & \cup \{ \text{not, or, xor, and} \} \cup \{ \text{eq, neq, lt, le, gt, ge} \} \end{aligned}$$

Σ_{ext} : further operations for special domains \mathbb{K} and user extensions

$$\begin{aligned} \Sigma_{\text{ext}} \supseteq & \{ \text{duplicate, permute} \} \\ & \cup \{ \text{shl, shr, \dots} \} \\ & \cup \{ \text{addc, subc, fmac, \dots} \} \end{aligned}$$

Embedding Boolean $\hookrightarrow \mathbb{K}$

- Needed for eq, neq, ..., not, or, ...

Embedding Boolean $\hookrightarrow \mathbb{K}$

- Needed for eq, neq, ..., not, or, ...
- Usually, natural implementation: for $\mathbb{K} = \mathbb{Z}_{64}$, take false $\mapsto 0$, true $\mapsto -1$.

Embedding Boolean $\hookrightarrow \mathbb{K}$

- Needed for eq, neq, ..., not, or, ...
- Usually, natural implementation: for $\mathbb{K} = \mathbb{Z}_{64}$, take false $\mapsto 0$, true $\mapsto -1$.

Embedding $\mathbb{Z} \hookrightarrow \mathbb{K}$ and $\mathbb{Z}^w \hookrightarrow \mathbb{K}^w$

- For shl, shr, ..., where $\text{shl}(a, n) := a 2^n$

Embedding Boolean $\hookrightarrow \mathbb{K}$

- Needed for eq, neq, ..., not, or, ...
- Usually, natural implementation: for $\mathbb{K} = \mathbb{Z}_{64}$, take false $\mapsto 0$, true $\mapsto -1$.

Embedding $\mathbb{Z} \hookrightarrow \mathbb{K}$ and $\mathbb{Z}^w \hookrightarrow \mathbb{K}^w$

- For shl, shr, ..., where $\text{shl}(a, n) := a 2^n$
- For permute on SIMD types, with $\text{permute}(a, \pi) := (a_{\pi(0)}, \dots, a_{\pi(w-1)})$

Embedding Boolean $\hookrightarrow \mathbb{K}$

- Needed for eq, neq, ..., not, or, ...
- Usually, natural implementation: for $\mathbb{K} = \mathbb{Z}_{64}$, take false $\mapsto 0$, true $\mapsto -1$.

Embedding $\mathbb{Z} \hookrightarrow \mathbb{K}$ and $\mathbb{Z}^w \hookrightarrow \mathbb{K}^w$

- For shl, shr, ..., where $\text{shl}(a, n) := a 2^n$
- For permute on SIMD types, with $\text{permute}(a, \pi) := (a_{\pi(0)}, \dots, a_{\pi(w-1)})$

Multi-sorted signatures

- We can also create union domains like $\mathbb{K} \cup \mathbb{L}$

Embedding Boolean $\hookrightarrow \mathbb{K}$

- Needed for eq, neq, ..., not, or, ...
- Usually, natural implementation: for $\mathbb{K} = \mathbb{Z}_{64}$, take false $\mapsto 0$, true $\mapsto -1$.

Embedding $\mathbb{Z} \hookrightarrow \mathbb{K}$ and $\mathbb{Z}^w \hookrightarrow \mathbb{K}^w$

- For shl, shr, ..., where $\text{shl}(a, n) := a 2^n$
- For permute on SIMD types, with $\text{permute}(a, \pi) := (a_{\pi(0)}, \dots, a_{\pi(w-1)})$

Multi-sorted signatures

- We can also create union domains like $\mathbb{K} \cup \mathbb{L}$
- And introduce variants $\sigma_{\mathbb{K}}, \sigma_{\mathbb{L}}, \dots$ of $\sigma \in \Sigma$ depending on the sort \mathbb{K}, \mathbb{L}

Embedding Boolean $\hookrightarrow \mathbb{K}$

- Needed for eq, neq, ..., not, or, ...
- Usually, natural implementation: for $\mathbb{K} = \mathbb{Z}_{64}$, take false $\mapsto 0$, true $\mapsto -1$.

Embedding $\mathbb{Z} \hookrightarrow \mathbb{K}$ and $\mathbb{Z}^w \hookrightarrow \mathbb{K}^w$

- For shl, shr, ..., where $\text{shl}(a, n) := a 2^n$
- For permute on SIMD types, with $\text{permute}(a, \pi) := (a_{\pi(0)}, \dots, a_{\pi(w-1)})$

Multi-sorted signatures

- We can also create union domains like $\mathbb{K} \cup \mathbb{L}$
- And introduce variants $\sigma_{\mathbb{K}}, \sigma_{\mathbb{L}}, \dots$ of $\sigma \in \Sigma$ depending on the sort \mathbb{K}, \mathbb{L}
- And versions $\sigma_{\text{conditional}}$ with an extra Boolean argument (if $\mathbb{L} = \text{Boolean}$)

Why include an operation like fma in Σ_{basic} ?

14/27

Cons

- Need to implement fma for many software domains

Why include an operation like `fma` in Σ_{basic} ?

Cons

- Need to implement `fma` for many software domains
- We could write a “clever” routine to simplify $ab + c \rightarrow \text{fma}(a, b, c)$

Why include an operation like `fma` in Σ_{basic} ?

Cons

- Need to implement `fma` for many software domains
- We could write a “clever” routine to simplify $ab + c \rightarrow \text{fma}(a, b, c)$

Pros

- `fma` corresponds to an important instruction in hardware

Why include an operation like fma in Σ_{basic} ?

Cons

- Need to implement fma for many software domains
- We could write a “clever” routine to simplify $ab + c \rightarrow \text{fma}(a, b, c)$

Pros

- fma corresponds to an important instruction in hardware
- $\text{fma}(a, b, c)$ does more than $ab + c$ for \mathbb{R}_{32} and \mathbb{R}_{64}

Why include an operation like fma in Σ_{basic} ?

Cons

- Need to implement fma for many software domains
- We could write a “clever” routine to simplify $ab + c \rightarrow \text{fma}(a, b, c)$

Pros

- fma corresponds to an important instruction in hardware
- $\text{fma}(a, b, c)$ does more than $ab + c$ for \mathbb{R}_{32} and \mathbb{R}_{64}
- Systematic support of $\pm x$, $\pm xy$, $\pm xy \pm z$ tends to yield better simplifications

Why include an operation like fma in Σ_{basic} ?

Cons

- Need to implement fma for many software domains
- We could write a “clever” routine to simplify $ab + c \rightarrow \text{fma}(a, b, c)$

Pros

- fma corresponds to an important instruction in hardware
- $\text{fma}(a, b, c)$ does more than $ab + c$ for \mathbb{R}_{32} and \mathbb{R}_{64}
- Systematic support of $\pm x$, $\pm xy$, $\pm xy \pm z$ tends to yield better simplifications
- On $\mathbb{Z}/p\mathbb{Z}$ via \mathbb{R}_{64} , better implementation of $\text{fma}(a, b, c)$ than $ab + c$:

$\text{reduce}(a)$	$ab + c$	$\text{fma}(a, b, c)$
$q := a \cdot u$	$h := a \cdot b$	$h := \text{fma}(a, b, c)$
$q := \text{round}(q)$	$l := \text{fms}(a, b, h)$	$l := \text{fms}(a, b, h)$
$r := \text{fnma}(p, q, r)$	$r := \text{reduce}(h)$	$l := l + c$
	$r := l + r$	$r := \text{reduce}(h)$
	$r := r + c$	$r := l + r$

A decorative gold border with intricate scrollwork and floral patterns, framing the central text.

Part III
Using JIL



Example: fast code for the cofactor matrix of a 4×4 matrix

- **Record** a program to compute a generic 4×4 determinant \longrightarrow Initial SLP f_1
- Compute gradient of f_1 \longrightarrow **algebraically transformed** SLP f_2
- Simplify the result \longrightarrow **optimized** SLP f_3
- **Compile** f_3 for a specific architecture \longrightarrow binary code inside memory



Example: fast code for the cofactor matrix of a 4×4 matrix

- **Record** a program to compute a generic 4×4 determinant → Initial SLP f_1
- Compute gradient of f_1 → **algebraically transformed** SLP f_2
- Simplify the result → **optimized** SLP f_3
- **Compile** f_3 for a specific architecture → binary code inside memory



Example: fast code for the cofactor matrix of a 4×4 matrix

- **Record** a program to compute a generic 4×4 determinant \longrightarrow Initial SLP f_1
- **Compute gradient of f_1** \longrightarrow **algebraically transformed SLP f_2**
- Simplify the result \longrightarrow **optimized SLP f_3**
- **Compile f_3** for a specific architecture \longrightarrow binary code inside memory



Example: fast code for the cofactor matrix of a 4×4 matrix

- **Record** a program to compute a generic 4×4 determinant \longrightarrow Initial SLP f_1
- Compute gradient of f_1 \longrightarrow **algebraically transformed** SLP f_2
- **Simplify the result** \longrightarrow **optimized SLP f_3**
- **Compile** f_3 for a specific architecture \longrightarrow binary code inside memory



Example: fast code for the cofactor matrix of a 4×4 matrix

- **Record** a program to compute a generic 4×4 determinant \longrightarrow Initial SLP f_1
- Compute gradient of f_1 \longrightarrow **algebraically transformed** SLP f_2
- Simplify the result \longrightarrow **optimized** SLP f_3
- **Compile** f_3 for a specific architecture \longrightarrow binary code inside memory

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

We may record an SLP for `f` as follows:

```
slp  
record_f (const domain& tp) {  
    recorder_start (tp);  
    slp_variable in = input_variable ();  
    slp_variable out= f (in);  
    specify_output (out);  
    return recorder_end ();  
}
```

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

We may record an SLP for `f` as follows:

```
slp  
record_f (const domain& tp) {  
    recorder_start (tp);  
    slp_variable in = input_variable ();  
    slp_variable out= f (in);  
    specify_output (out);  
    return recorder_end ();  
}
```

Idea: an instance of `slp_variable` specifies a data field on which to operate operations on `slp_variable` are recorded instead of being executed

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

Step by step execution

```
recorder_start (tp)
```

in
out
prg
data

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

Step by step execution

```
recorder_start (tp)  
in = input_variable ()
```

in

out

prg

data

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

Step by step execution

```
recorder_start (tp)  
in = input_variable ()  
out= foo (in)  
    x * x
```

in

0

out

prg

x	1	0	0
---	---	---	---

data

0	0
---	---

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

Step by step execution

```
recorder_start (tp)  
in = input_variable ()  
out= foo (in)  
    x * x  
    x * x + x
```

in

0

out

prg

x	1	0	0	+	2	1	0
---	---	---	---	---	---	---	---

data

0	0	0
---	---	---

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

Step by step execution

```
recorder_start (tp)  
in = input_variable ()  
out= foo (in)  
  x * x  
  x * x + x  
  3
```

in

0

out

prg

x	1	0	0	+	2	1	0
---	---	---	---	---	---	---	---

data

0	0	0	3
---	---	---	---

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

Step by step execution

```
recorder_start (tp)  
in = input_variable ()  
out= foo (in)  
  x * x  
  x * x + x  
  3  
  x * x + x - 3
```

in

0

out

prg

x	1	0	0	+	2	1	0	-	4	2	3
---	---	---	---	---	---	---	---	---	---	---	---

data

0	0	0	3	0
---	---	---	---	---

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

Step by step execution

```
recorder_start (tp)  
in = input_variable ()  
out= foo (in)  
  x * x  
  x * x + x  
  3  
  x * x + x - 3  
output_variable (out)
```

in

0

out

4

prg

×	1	0	0	+	2	1	0	-	4	2	3
---	---	---	---	---	---	---	---	---	---	---	---

data

0	0	0	3	0
---	---	---	---	---

Assume that we have a generic C++ function

```
template<typename C> C  
f (const C& x) { return x * x + x - 3; }
```

Step by step execution

```
recorder_start (tp)  
in = input_variable ()  
out= foo (in)  
  x * x  
  x * x + x  
  3  
  x * x + x - 3  
output_variable (out)
```

in

0

out

4

prg

×	1	0	0	+	2	1	0	-	4	2	3
---	---	---	---	---	---	---	---	---	---	---	---

data

0	0	0	3	0
---	---	---	---	---

Note: clean execution trace more important than efficient execution of **f**

Important: by design, all transformations of SLP take linear time in JIL

Important: by design, all transformations of SLP take linear time in JIL

Simplification

- Common subexpression elimination
- CSE + algebraic simplifications ($0 + x \rightarrow x$, etc.)
- Dead code elimination
- Rewrite $ab + c \rightarrow \text{fma}(a, b, c)$

Important: by design, all transformations of SLP take linear time in JIL

Simplification

- Common subexpression elimination
- CSE + algebraic simplifications ($0 + x \rightarrow x$, etc.)
- Dead code elimination
- Rewrite $ab + c \rightarrow \text{fma}(a, b, c)$

Backend

- Emulate missing instructions
- Optimizations for immediate arguments
- Rescheduling
- Register allocation

General transformations

- Forward and backward differentiation: $f \mapsto \text{Jac}(f)$
- If f is linear (w.r.t. some of its inputs), then compute the transposed map
- $P \in \mathbb{K}[x_1, \dots, x_n] \mapsto$ homogeneous $\tilde{P} \in \mathbb{K}[x_1, \dots, x_n, t]$ with $P(\mathbf{x}) = \tilde{P}(\mathbf{x}, 1)$
- Add just enough reductions to avoid overflows for redundant representations

General transformations

- Forward and backward differentiation: $f \mapsto \text{Jac}(f)$
- If f is linear (w.r.t. some of its inputs), then compute the transposed map
- $P \in \mathbb{K}[x_1, \dots, x_n] \mapsto$ homogeneous $\tilde{P} \in \mathbb{K}[x_1, \dots, x_n, t]$ with $P(\mathbf{x}) = \tilde{P}(\mathbf{x}, 1)$
- Add just enough reductions to avoid overflows for redundant representations

Lifting and related transformations

- Lift SLP over \mathbb{K} to SLP over \mathbb{A} for \mathbb{K} -algebra \mathbb{A}
- Reinterpret SLP over \mathbb{A} as SLP over \mathbb{K}
- Specific vector and ball lifts
- Reduce the number of divisions in SLPs

Example of complexification

$\mathbb{K}[i]$ is an **SLP algebra**: operations in Σ_{basic} can be implemented using SLPs:

add

in(x_1, y_1, x_2, y_2)

$x_3 := x_1 + x_2$

$y_3 := y_1 + y_2$

out(x_3, y_3)

subtract

in(x_1, y_1, x_2, y_2)

$x_3 := x_1 + x_2$

$y_3 := y_1 + y_2$

out(x_3, y_3)

multiply

in(x_1, y_1, x_2, y_2)

$x_3 := x_1 \cdot x_2$

$y_3 := x_1 \cdot y_2$

$x_3 := \text{fms}(y_1, y_2, x_3)$

$y_3 := \text{fma}(x_2, y_1, y_3)$

out(x_3, y_3)

Be careful with *aliasing* like in $z := u \cdot z$

Example of complexification

$\mathbb{K}[i]$ is an **SLP algebra**: operations in Σ_{basic} can be implemented using SLPs:

add	subtract	multiply
$x_3 := x_1 + x_2$	$x_3 := x_1 + x_2$	$x_3 := x_1 \cdot x_2$
$y_3 := y_1 + y_2$	$y_3 := y_1 + y_2$	$y_3 := x_1 \cdot y_2$
		$x_3 := \text{fms}(y_1, y_2, x_3)$
		$y_3 := \text{fma}(x_2, y_1, y_3)$

$\text{in}(z_0)$

$z_1 := z_0 \cdot z_0$

$z_2 := z_1 + z_0$

$z_4 := z_2 - 7$

$\text{out}(z_4)$

lift
→

$\text{in}(x_0, y_0)$

$x_1 := x_0 \cdot x_0$

$y_1 := x_0 \cdot y_0$

$x_1 := \text{fms}(y_0, y_0, x_1)$

$y_1 := \text{fma}(x_0, y_0, y_1)$

$x_2 := x_1 + x_0$

$y_2 := y_1 + y_0$

$x_4 := x_2 - 7$

$y_4 := y_2 - 0$

$\text{out}(x_4, y_4)$

simplify
→

$\text{in}(x_0, y_0)$

$x_1 := x_0 \cdot x_0$

$y_1 := x_0 \cdot y_0$

$x_1 := \text{fms}(y_0, y_0, x_1)$

$y_1 := \text{fma}(x_0, y_0, y_1)$

$x_2 := x_1 + x_0$

$y_2 := y_1 + y_0$

$x_4 := x_2 - 7$

$\text{out}(x_4, y_2)$

n	len	cse	sim	∇	lift	reg	jit	exe
2	2	1460	2687	1221	1796	1436	12784	10.00
3	9	543	825	407	614	562	2963	2.222
4	40	229	341	184	262	291	924	0.550
5	205	126	196	125	163	237	452	0.424
6	1236	97	149	112	138	221	455	0.422
7	8659	89	144	127	134	229	447	0.419
8	69280	94	171	125	160	261	470	0.424
9	623529	133	207	188	178	347	472	0.865
10	6235300	158	242	296	245	391	522	3.308

Table. Timings in cycles per instruction on INTEL XEON for (very) naive $n \times n$ determinants.

A decorative gold border with intricate scrollwork and floral patterns, framing the central text.

Part IV
Upcoming features

We can efficiently generate multiple implementations of a function and determine which one is best by benching each implementation

We can efficiently generate multiple implementations of a function and determine which one is best by benching each implementation

This requires

- Good random sample generators
- Mechanism to list alternative implementations and bench them
- Mechanism to cache results on disk
- Mechanism to predict timings without running any code

Typical use cases:

- Run SLP $f: \mathbb{K}^m \rightarrow \mathbb{K}^n$ on vectors: $\tilde{f}: (\mathbb{K}^m)^N \rightarrow (\mathbb{K}^n)^N$ with $N \gg 1$

Typical use cases:

- Run SLP $f: \mathbb{K}^m \rightarrow \mathbb{K}^n$ on vectors: $\tilde{f}: (\mathbb{K}^m)^N \rightarrow (\mathbb{K}^n)^N$ with $N \gg 1$
- Iterate SLP $f: \mathbb{K}^n \rightarrow \mathbb{K}^n$ until condition is met (e.g. homotopy continuation)

Typical use cases:

- Run SLP $f: \mathbb{K}^m \rightarrow \mathbb{K}^n$ on vectors: $\tilde{f}: (\mathbb{K}^m)^N \rightarrow (\mathbb{K}^n)^N$ with $N \gg 1$
- Iterate SLP $f: \mathbb{K}^n \rightarrow \mathbb{K}^n$ until condition is met (e.g. homotopy continuation)
- SLP over “big” SLP algebra \mathbb{A} over $\mathbb{K} \rightarrow$ subroutines for operations in \mathbb{A}

Typical use cases:

- Run SLP $f: \mathbb{K}^m \rightarrow \mathbb{K}^n$ on vectors: $\tilde{f}: (\mathbb{K}^m)^N \rightarrow (\mathbb{K}^n)^N$ with $N \gg 1$
- Iterate SLP $f: \mathbb{K}^n \rightarrow \mathbb{K}^n$ until condition is met (e.g. homotopy continuation)
- SLP over “big” SLP algebra \mathbb{A} over $\mathbb{K} \rightarrow$ subroutines for operations in \mathbb{A}

Need for *shallow* control structures

Typical use cases:

- Run SLP $f: \mathbb{K}^m \rightarrow \mathbb{K}^n$ on vectors: $\tilde{f}: (\mathbb{K}^m)^N \rightarrow (\mathbb{K}^n)^N$ with $N \gg 1$
- Iterate SLP $f: \mathbb{K}^n \rightarrow \mathbb{K}^n$ until condition is met (e.g. homotopy continuation)
- SLP over “big” SLP algebra \mathbb{A} over $\mathbb{K} \rightarrow$ subroutines for operations in \mathbb{A}

Need for *shallow* control structures

Dedicated support for various frequent patterns...

... or better to use general purpose techniques?

Input: $N \in \mathbb{N}$ and vectors $u, v \in \mathbb{R}_{64}[\mathbf{i}]^N$

Output: $w \in \mathbb{R}_{64}[\mathbf{i}]^N$ with $w_i = u_i v_i$ for all $0 \leq i < N$

Input: $N \in \mathbb{N}$ and vectors $u, v \in \mathbb{R}_{64}[\mathbf{i}]^N$

Output: $w \in \mathbb{R}_{64}[\mathbf{i}]^N$ with $w_i = u_i v_i$ for all $0 \leq i < N$

First assume that $N = 32$.

Input: $N \in \mathbb{N}$ and vectors $u, v \in \mathbb{R}_{64}[\mathbf{i}]^N$

Output: $w \in \mathbb{R}_{64}[\mathbf{i}]^N$ with $w_i = u_i v_i$ for all $0 \leq i < N$

First assume that $N = 32$.

- Reinterpret u, v as $\tilde{u}, \tilde{v} \in (\mathbb{R}_{64}^8)^4[\mathbf{i}]$ (this requires SIMD matrix transposition)
- Lift complex multiplication over \mathbb{R}_{64} to multiplication in $(\mathbb{R}_{64}^8)^4[\mathbf{i}]$ over \mathbb{R}_{64}^8
- Reinterpret $\tilde{w} := \tilde{u} \tilde{v}$ as an element of $\mathbb{R}_{64}[\mathbf{i}]$ (another SIMD transposition)

Input: $N \in \mathbb{N}$ and vectors $u, v \in \mathbb{R}_{64}[\mathbf{i}]^N$

Output: $w \in \mathbb{R}_{64}[\mathbf{i}]^N$ with $w_i = u_i v_i$ for all $0 \leq i < N$

First assume that $N = 32$.

- Reinterpret u, v as $\tilde{u}, \tilde{v} \in (\mathbb{R}_{64}^8)^4[\mathbf{i}]$ (this requires SIMD matrix transposition)
- Lift complex multiplication over \mathbb{R}_{64} to multiplication in $(\mathbb{R}_{64}^8)^4[\mathbf{i}]$ over \mathbb{R}_{64}^8
- Reinterpret $\tilde{w} := \tilde{u} \tilde{v}$ as an element of $\mathbb{R}_{64}[\mathbf{i}]$ (another SIMD transposition)

Create a loop which repeats this code until $N < 32$.

Input: $N \in \mathbb{N}$ and vectors $u, v \in \mathbb{R}_{64}[\mathbf{i}]^N$

Output: $w \in \mathbb{R}_{64}[\mathbf{i}]^N$ with $w_i = u_i v_i$ for all $0 \leq i < N$

First assume that $N = 32$.

- Reinterpret u, v as $\tilde{u}, \tilde{v} \in (\mathbb{R}_{64}^8)^4[\mathbf{i}]$ (this requires SIMD matrix transposition)
- Lift complex multiplication over \mathbb{R}_{64} to multiplication in $(\mathbb{R}_{64}^8)^4[\mathbf{i}]$ over \mathbb{R}_{64}^8
- Reinterpret $\tilde{w} := \tilde{u} \tilde{v}$ as an element of $\mathbb{R}_{64}[\mathbf{i}]$ (another SIMD transposition)

Create a loop which repeats this code until $N < 32$. Next

Input: $N \in \mathbb{N}$ and vectors $u, v \in \mathbb{R}_{64}[\mathbf{i}]^N$

Output: $w \in \mathbb{R}_{64}[\mathbf{i}]^N$ with $w_i = u_i v_i$ for all $0 \leq i < N$

First assume that $N = 32$.

- Reinterpret u, v as $\tilde{u}, \tilde{v} \in (\mathbb{R}_{64}^8)^4[\mathbf{i}]$ (this requires SIMD matrix transposition)
- Lift complex multiplication over \mathbb{R}_{64} to multiplication in $(\mathbb{R}_{64}^8)^4[\mathbf{i}]$ over \mathbb{R}_{64}^8
- Reinterpret $\tilde{w} := \tilde{u} \tilde{v}$ as an element of $\mathbb{R}_{64}[\mathbf{i}]$ (another SIMD transposition)

Create a loop which repeats this code until $N < 32$. Next

- Reduce further to the cases when $N < 16$ and then $N < 8$
- Reduce the case when $N < 8 \rightarrow$ case $N = 8$ with an appropriate mask

Input: $N \in \mathbb{N}$ and vectors $u, v \in \mathbb{R}_{64}[\mathbf{i}]^N$

Output: $w \in \mathbb{R}_{64}[\mathbf{i}]^N$ with $w_i = u_i v_i$ for all $0 \leq i < N$

First assume that $N = 32$.

- Reinterpret u, v as $\tilde{u}, \tilde{v} \in (\mathbb{R}_{64}^8)^4[\mathbf{i}]$ (this requires SIMD matrix transposition)
- Lift complex multiplication over \mathbb{R}_{64} to multiplication in $(\mathbb{R}_{64}^8)^4[\mathbf{i}]$ over \mathbb{R}_{64}^8
- Reinterpret $\tilde{w} := \tilde{u} \tilde{v}$ as an element of $\mathbb{R}_{64}[\mathbf{i}]$ (another SIMD transposition)

Create a loop which repeats this code until $N < 32$. Next

- Reduce further to the cases when $N < 16$ and then $N < 8$
- Reduce the case when $N < 8 \rightarrow$ case $N = 8$ with an appropriate mask

This whole implementation can be run using only AVX512 vector instructions.

Thank you !



<http://www.TEXMACS.org>