

✓ Swiss Verification Day 2025

13 February 2025, USI East Campus, Lugano

# Reasoning about out-of-order execution in dataflow circuits

---

Yann Herklotz, Ayatallah Elakhras, Martina Camaioni, Lana Josipović, Paolo lenne, Thomas Bourgeat

February 13, 2025

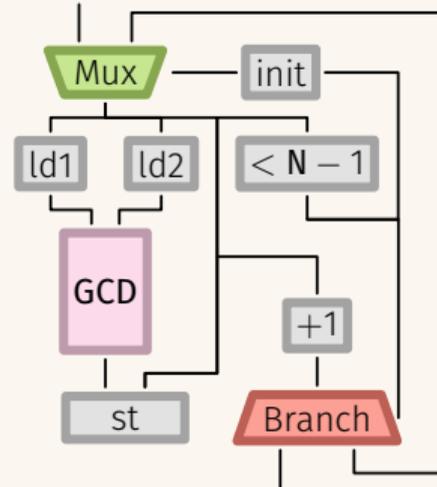
**VCA**

**EPFL**

# High-level synthesis with dataflow graphs

```
for (int i = 0; i < N; i++) {  
    result[i] = gcd(arr1[i], arr2[i]);  
}
```

Dynamic high-level synthesis tools translate C code into dataflow circuits.



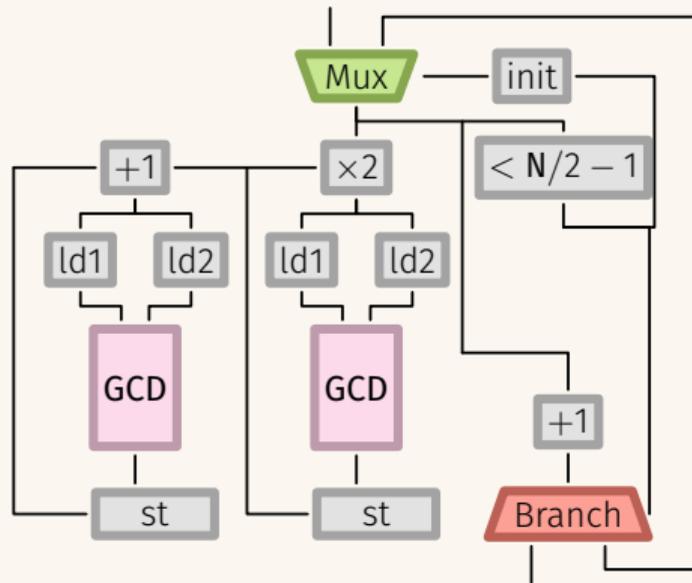
max 1 GCD operation at a time.

# High-level synthesis with dataflow graphs

```
for (int i = 0; i < N/2; i++) {  
    result[2*i] =  
        gcd(arr1[2*i], arr2[2*i]);  
    result[2*i+1] =  
        gcd(arr1[2*i+1], arr2[2*i+1]);  
}
```

Dynamic high-level synthesis tools translate C code into dataflow circuits.

Manual optimisations can help throughput.

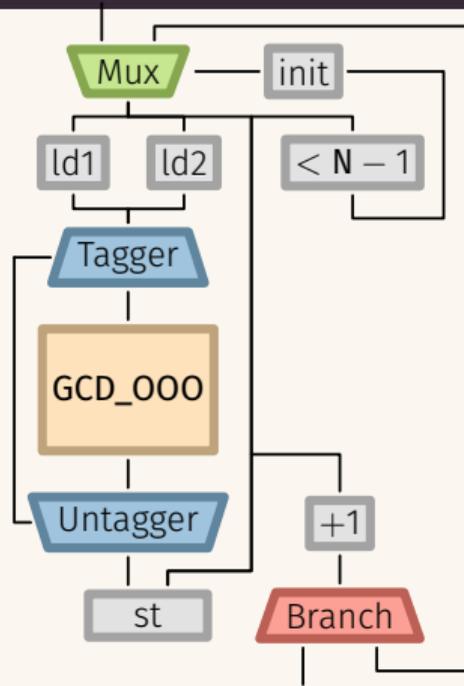


up to 2 GCD operation at a time.

## Out-of-order in dataflow circuits

We can generalise the arbiter idea by defining an out-of-order GCD component.

A more general tagger/untagger pair keeps track of the correct order of tokens.



$n$  GCD operation at a time.

# Out-of-order in dataflow circuits

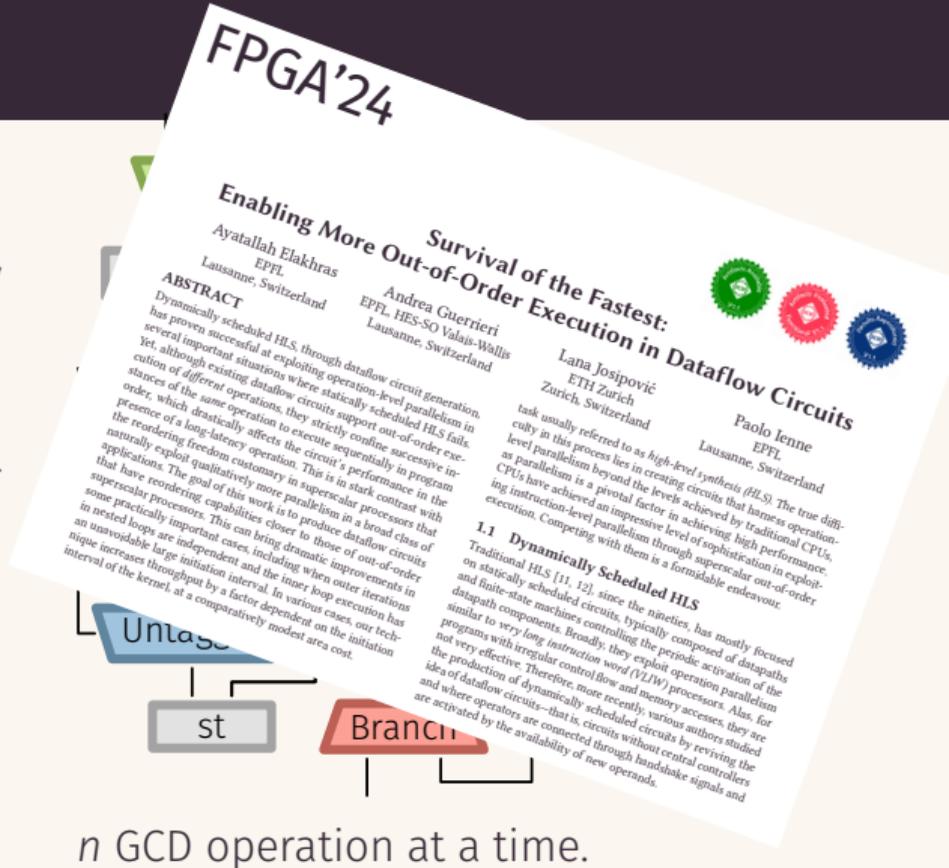
We can generalise the arbiter idea by defining and out-of-order GCD component.

A more general tagger/untagger pair keeps track of the correct order of tokens.

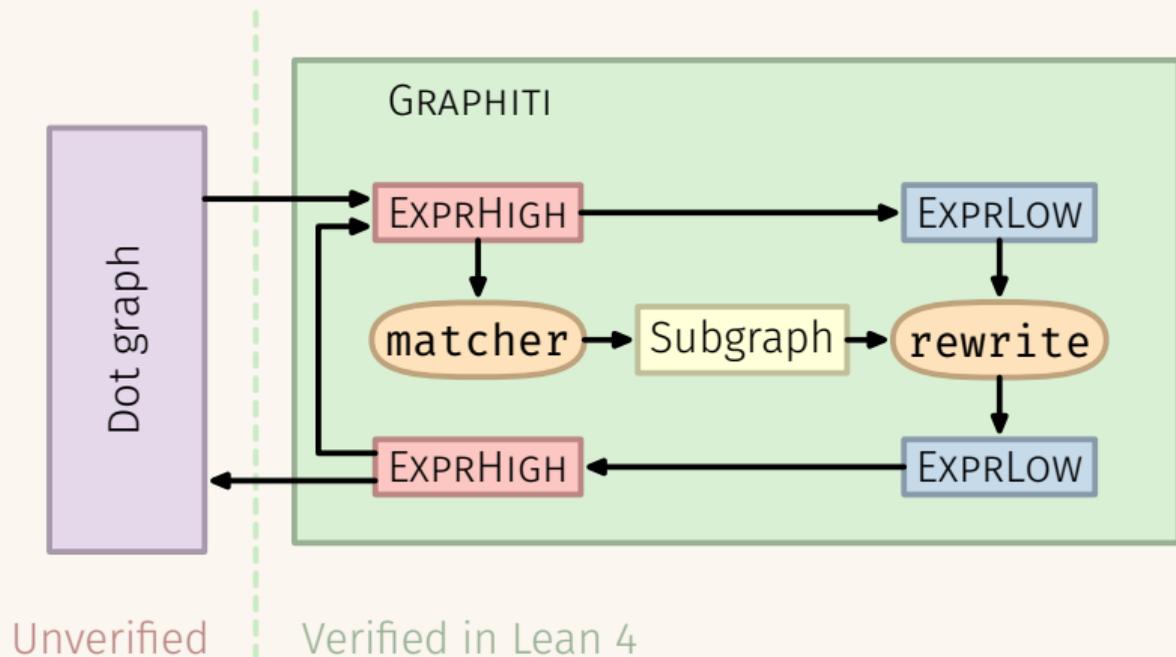
Systematic approach developed by the Dynamic lab.

$n$  GCD operation at a time.

Ayatallah Elakhras, Andrea Guerrieri, Lana Josipovic, and Paolo lenne. “Survival of the Fastest: Enabling More Out-of-Order Execution in Dataflow Circuits”. In: *Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays*. FPGA '24. Monterey, CA, USA: Association for Computing Machinery, 2024, pp. 44–54. ISBN: 9798400704185. doi: 10.1145/3626202.3637556

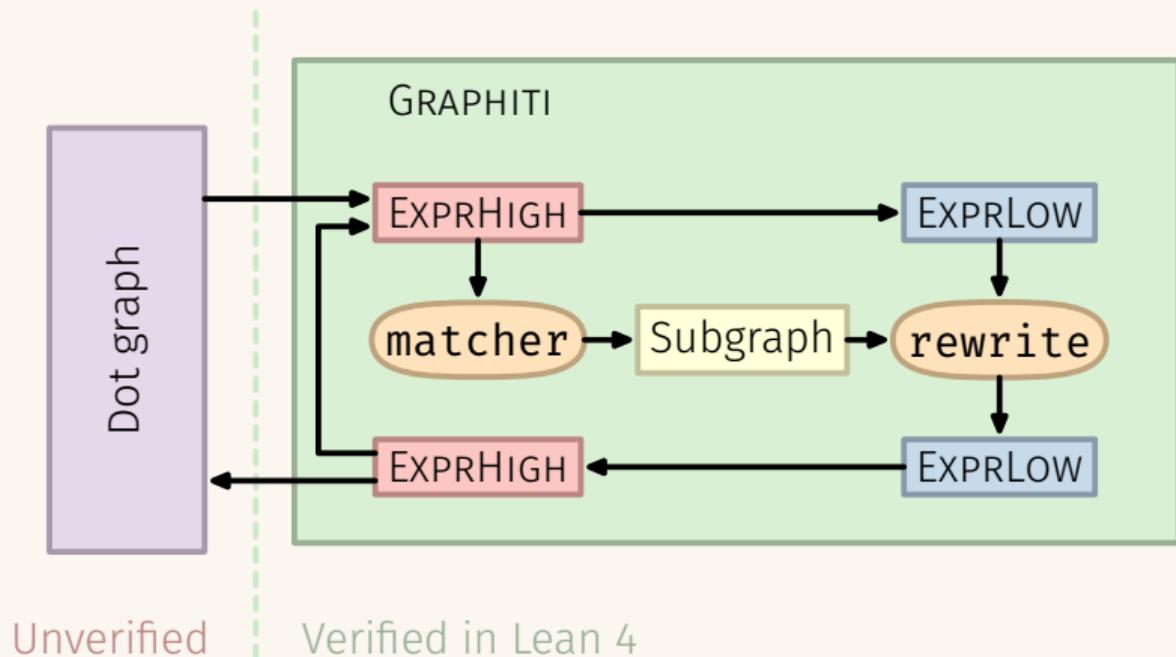


# Rewrites are a nice way to perform optimisations in a verified framework



Started collaborating on a verified dataflow rewrite framework.

# Rewrites are a nice way to perform optimisations in a verified framework



Started collaborating on a verified dataflow rewrite framework. Use the rewrite framework to implement the FPGA'24 algorithm.

# Quick overview of the syntax and semantics

## Language descriptions

**ExprHigh** Graph language represented by sets of nodes and edges.

**ExprLow** Inductive graph language with products and connections.

## Graph denotation and semantics

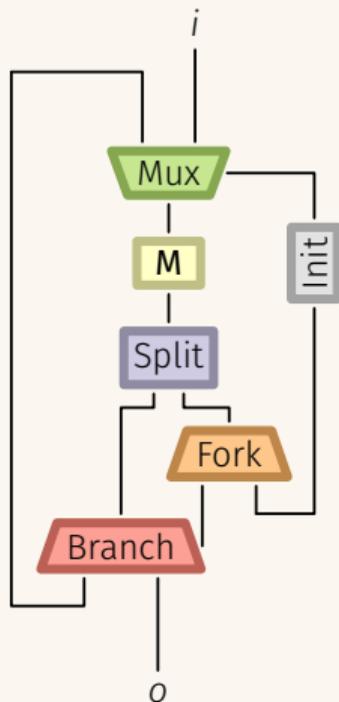
```
def join T T' : NatModule (List T × List T') :=
  { inputs := [ (0, ( T, fun ol el nl => nl.1 = ol.1.concat el /\ nl.2 = ol.2))
                , (1, ( T', fun ol el nl => nl.2 = ol.2.concat el /\ nl.1 = ol.1))]
    .toAssocList,
    outputs := [(0, ( T × T', fun ol el nl => ol.1 = el.1 :: nl.1 /\ ol.2 = el.2 :: nl.2 ))]
    .toAssocList }
```

## Top-level correctness theorem

$$(rw_{\text{rhs}} \Downarrow \varepsilon) \sqsubseteq (rw_{\text{lhs}} \Downarrow \varepsilon) \rightarrow$$
$$(e[rw_{\text{lhs}} := rw_{\text{lhs}}] \Downarrow \varepsilon) \sqsubseteq (e \Downarrow \varepsilon)$$

# In particular we only need one general rewrite replicate many of the FPGA'24 results

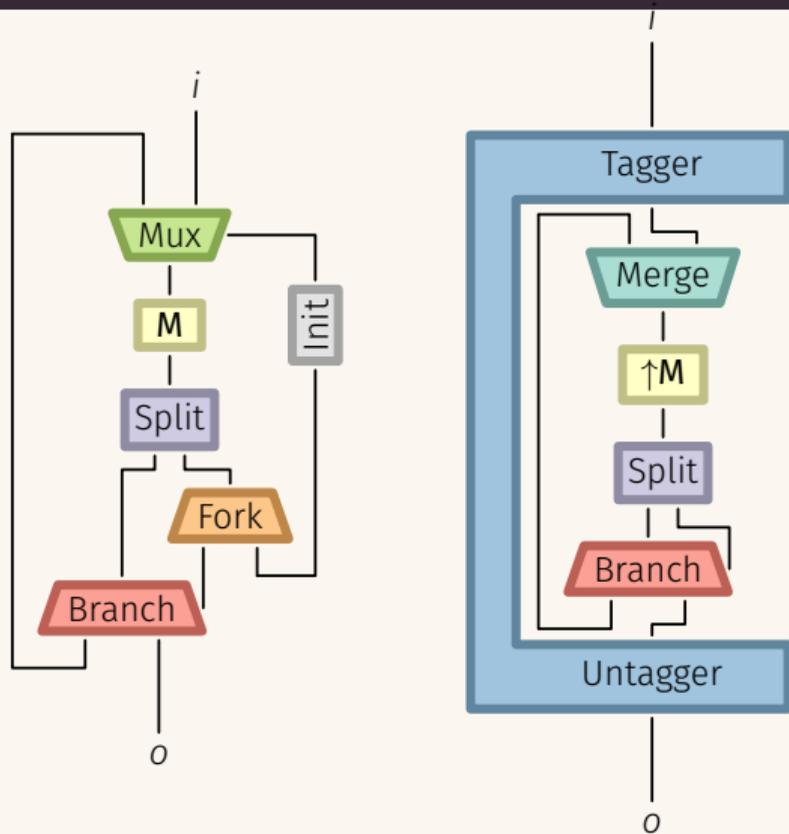
Given a **sequential** loop over an arbitrary loop body  $M$ .



# In particular we only need one general rewrite replicate many of the FPGA'24 results

Given a **sequential** loop over an arbitrary loop body  $M$ .

We can generate an **out-of-order** version of the loop by **overlapping** loop executions.

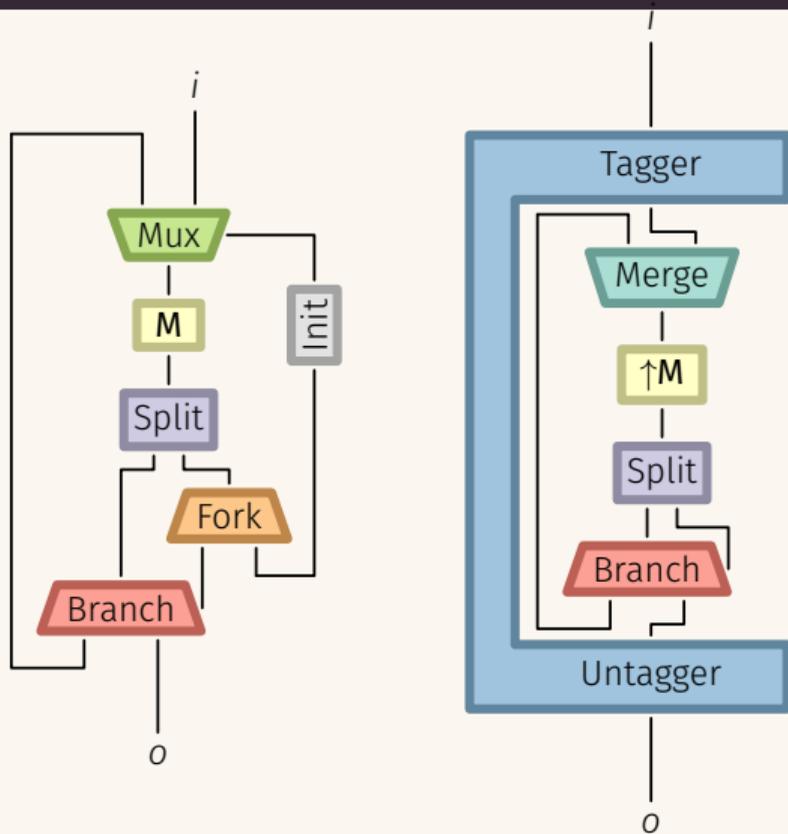


# In particular we only need one general rewrite replicate many of the FPGA'24 results

Given a **sequential** loop over an arbitrary loop body  $M$ .

We can generate an **out-of-order** version of the loop by **overlapping** loop executions.

Does this hold for an **arbitrary** module  $M$ ?



# In particular we only need one general rewrite replicate many of the FPGA'24 results

Given a **sequential** loop over an arbitrary loop body  $M$ .

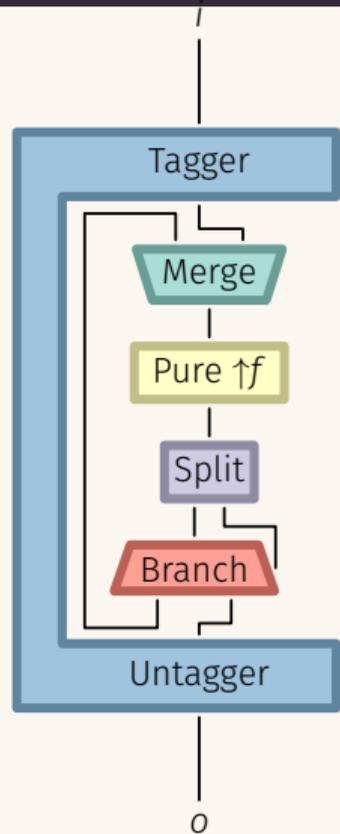
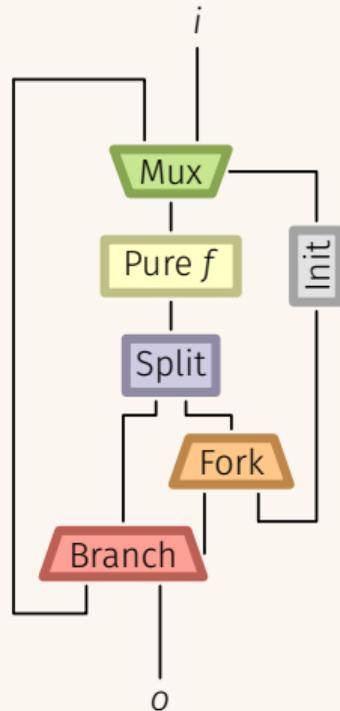
We can generate an **out-of-order** version of the loop by **overlapping** loop executions.

Does this hold for an **arbitrary** module  $M$ ?

**No**, we need:

**1-in-1-out** behaviour for the module.

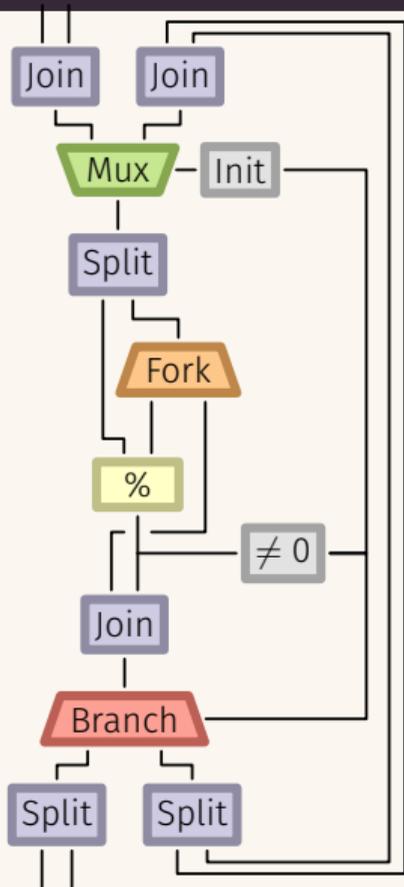
It needs to be **stateless**.







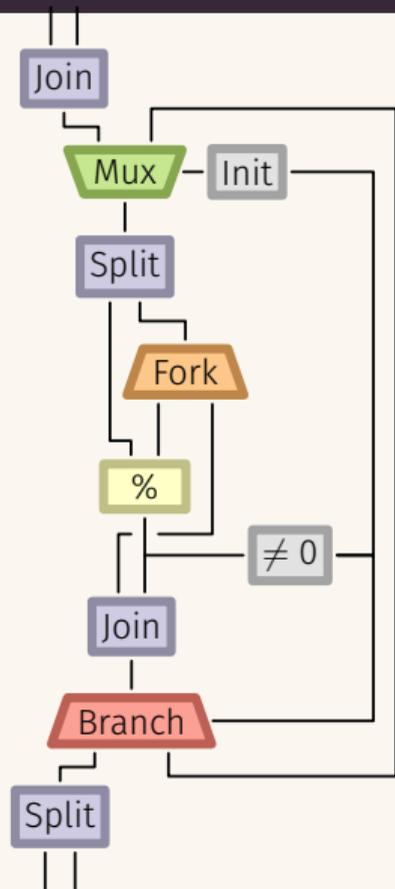
## Applying the rewrite to the GCD example



We **have a problem**: too many muxes and branches, and we don't know whether the body can be modelled as a pure function.

**Combine** muxes and branches.

## Applying the rewrite to the GCD example

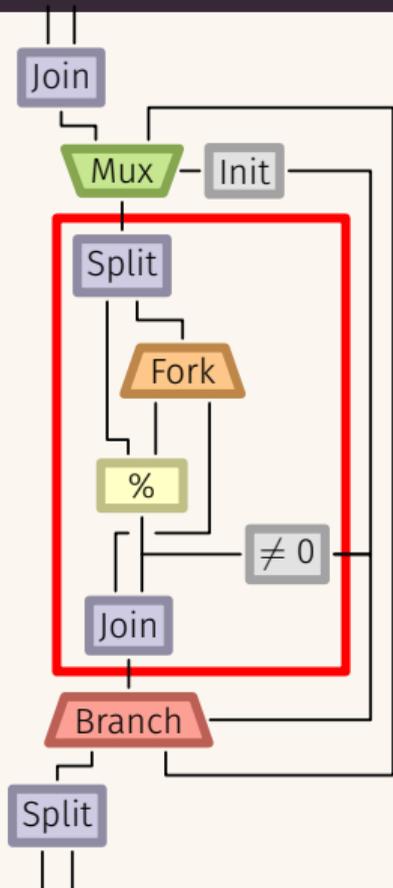


We **have a problem**: too many muxes and branches, and we don't know whether the body can be modelled as a pure function.

**Combine** muxes and branches.

**Remove** unnecessary splits and joins.

## Applying the rewrite to the GCD example



We **have a problem**: too many muxes and branches, and we don't know whether the body can be modelled as a pure function.

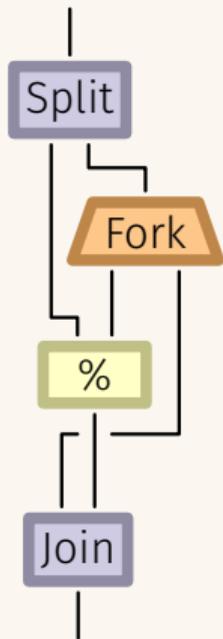
**Combine** muxes and branches.

**Remove** unnecessary splits and joins.

Prove that the body of the loop is

**Pure: More rewrites!**

## Applying the rewrite to the GCD example



We **have a problem**: too many muxes and branches, and we don't know whether the body can be modelled as a pure function.

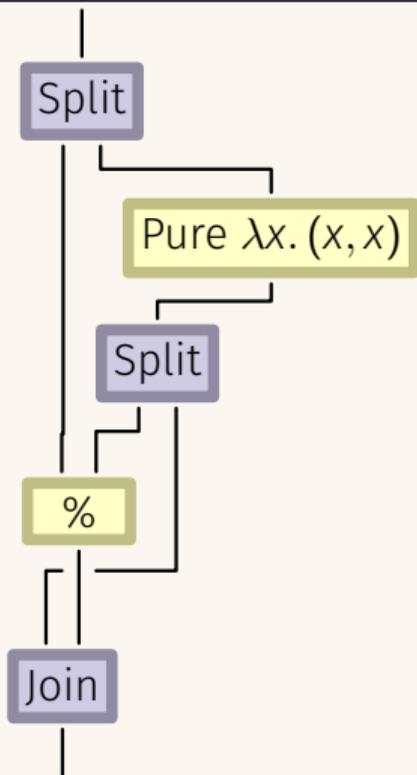
**Combine** muxes and branches.

**Remove** unnecessary splits and joins.

Prove that the body of the loop is

**Pure: More rewrites!**

## Applying the rewrite to the GCD example



We **have a problem**: too many muxes and branches, and we don't know whether the body can be modelled as a pure function.

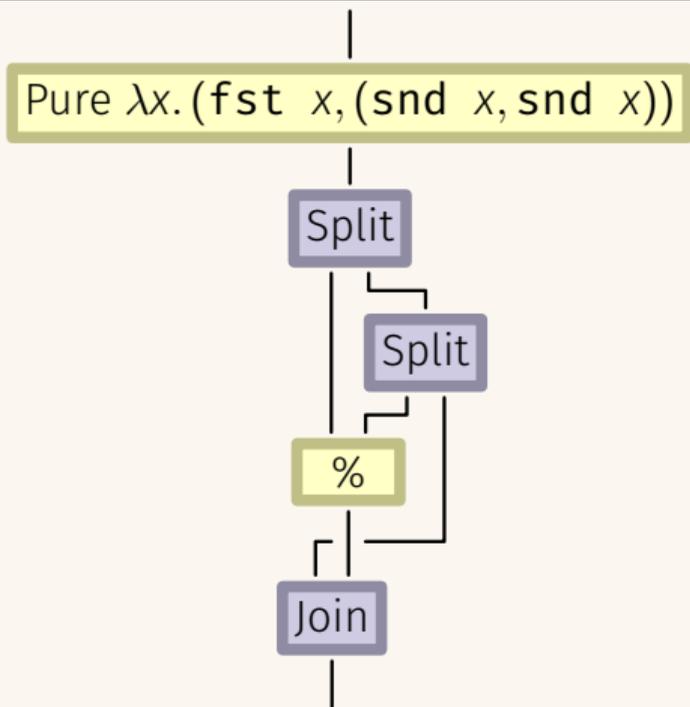
**Combine** muxes and branches.

**Remove** unnecessary splits and joins.

Prove that the body of the loop is

**Pure: More rewrites!**

## Applying the rewrite to the GCD example



We **have a problem**: too many muxes and branches, and we don't know whether the body can be modelled as a pure function.

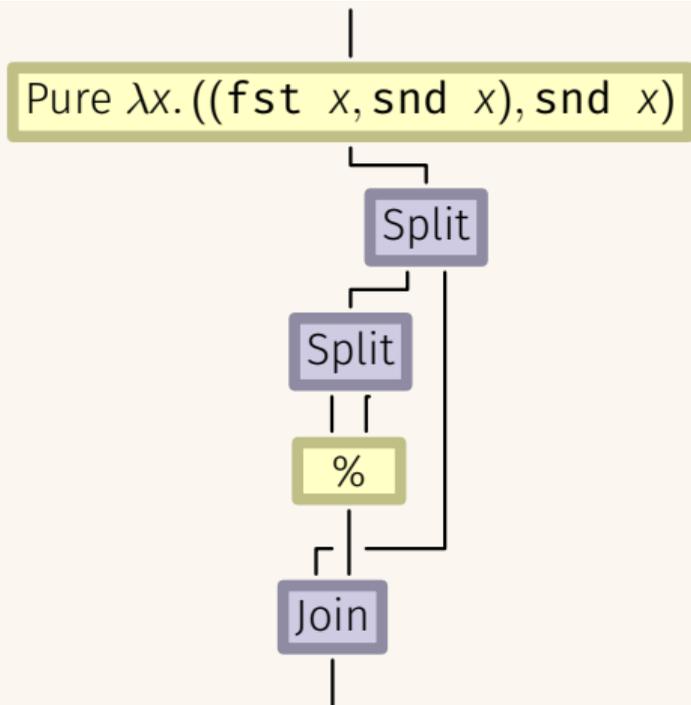
**Combine** muxes and branches.

**Remove** unnecessary splits and joins.

Prove that the body of the loop is

**Pure: More rewrites!**

## Applying the rewrite to the GCD example



We **have a problem**: too many muxes and branches, and we don't know whether the body can be modelled as a pure function.

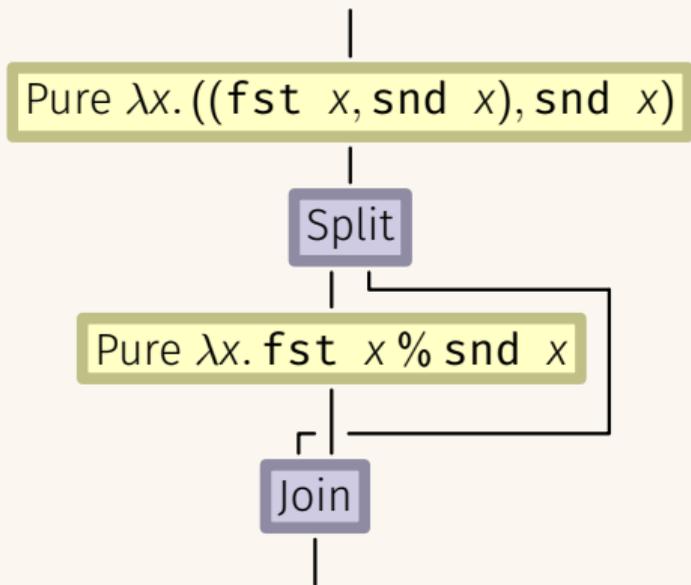
**Combine** muxes and branches.

**Remove** unnecessary splits and joins.

Prove that the body of the loop is

**Pure: More rewrites!**

## Applying the rewrite to the GCD example



We **have a problem**: too many muxes and branches, and we don't know whether the body can be modelled as a pure function.

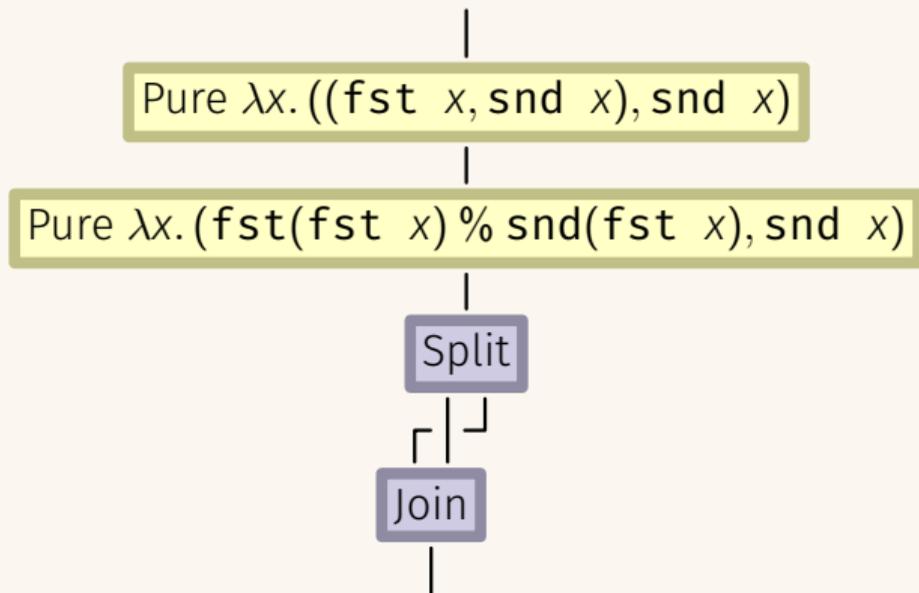
**Combine** muxes and branches.

**Remove** unnecessary splits and joins.

Prove that the body of the loop is

**Pure: More rewrites!**

## Applying the rewrite to the GCD example



We **have a problem**: too many muxes and branches, and we don't know whether the body can be modelled as a pure function.

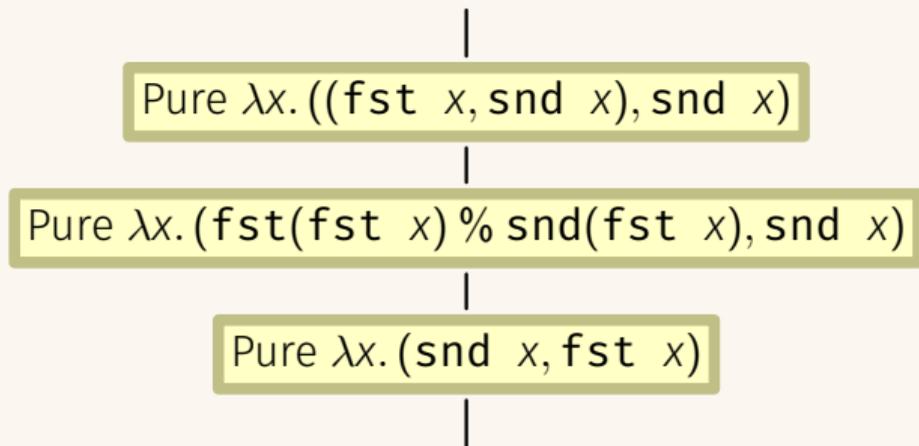
**Combine** muxes and branches.

**Remove** unnecessary splits and joins.

Prove that the body of the loop is

**Pure: More rewrites!**

## Applying the rewrite to the GCD example



We **have a problem**: too many muxes and branches, and we don't know whether the body can be modelled as a pure function.

**Combine** muxes and branches.

**Remove** unnecessary splits and joins.

Prove that the body of the loop is

**Pure: More rewrites!**

## Applying the rewrite to the GCD example

Pure  $\lambda x. (\text{snd } x, \text{fst } x \% \text{snd } x)$

We **have a problem**: too many muxes and branches, and we don't know whether the body can be modelled as a pure function.

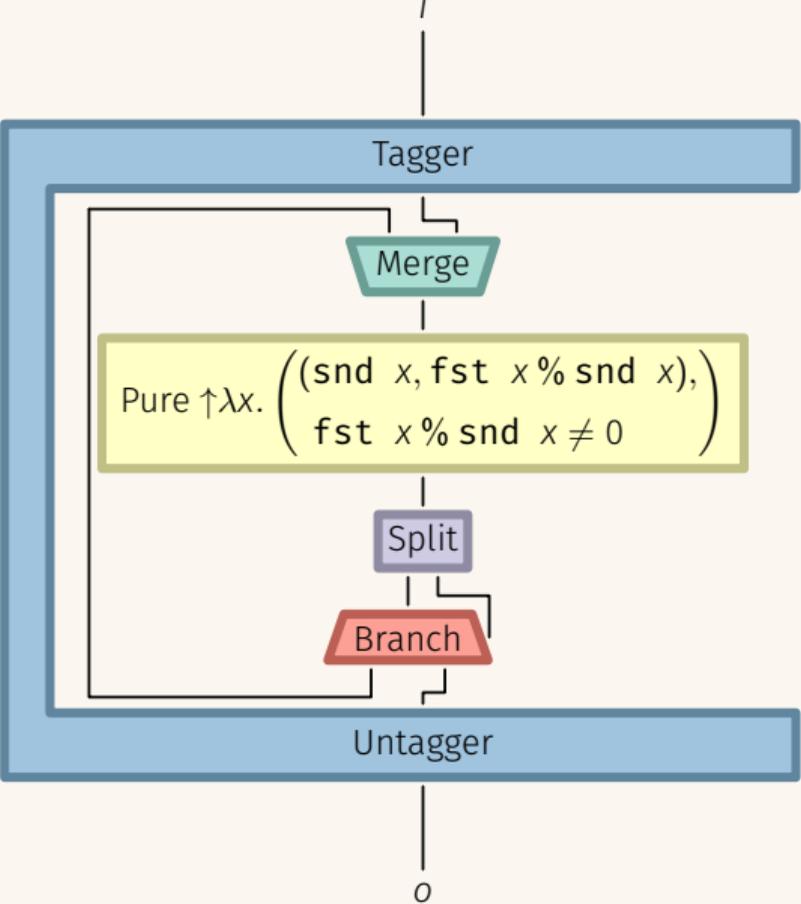
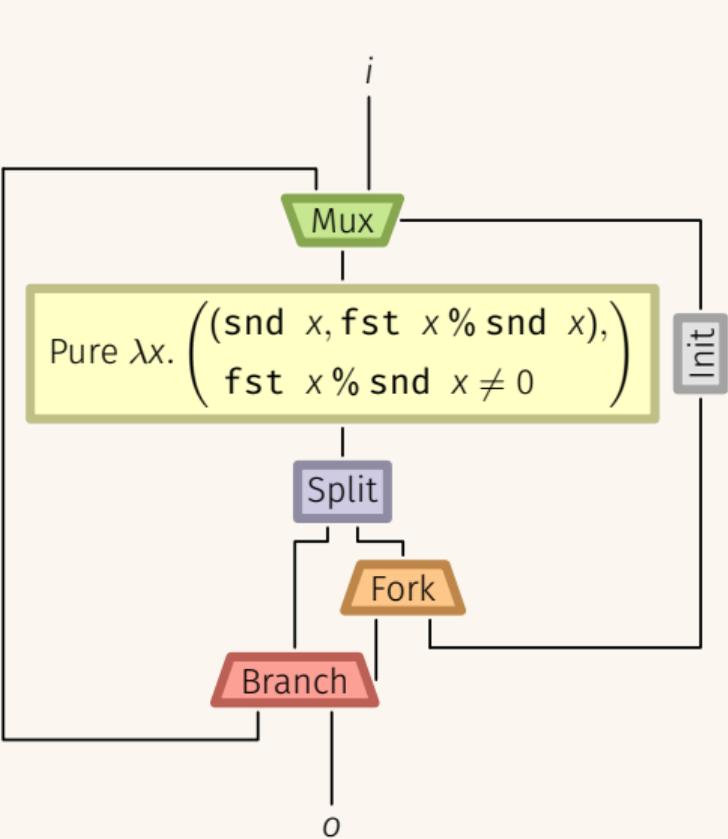
**Combine** muxes and branches.

**Remove** unnecessary splits and joins.

Prove that the body of the loop is

**Pure: More rewrites!**

# Finally we can parallelise our GCD!



# Hypothetical results

Benchmark	N	Cycle count		CP (ns)		Execution time ( $\mu$ s)			LUTs			FFs		DSPs	
		[18]	Ours	[18]	Ours	[18]	Ours		[18]	Ours		[18]	Ours		
img-avg	4	1,722	634	6.42	8.21	11.05	5.21	<b>2.1x</b>	1,415	1,593	<b>+13%</b>	1,320	1,206	<b>-9%</b>	5
gsum_many	10	68,523	32,874	7.91	9.95	541.81	327.0	<b>1.7x</b>	2,835	3,657	<b>+29%</b>	3,256	3,725	<b>+14%</b>	22
gsum_single	10	6,703	9,333	6.90	9.01	46.25	84.11	<b>0.55x</b>	2,736	2,677	<b>-2%</b>	3,142	3,114	<b>-1%</b>	22
gemm	20	68,825	8,144	6.75	12.81	464.78	104.30	<b>4.5x</b>	3,214	5,937	<b>+85%</b>	2,693	3,688	<b>+37%</b>	11
matvec	50	7,936	918	6.41	13.51	50.86	12.40	<b>4.1x</b>	1,272	4,396	<b>+246%</b>	1,373	3,423	<b>+149%</b>	5
mvt	10	7,940	2,044	6.27	11.70	49.79	23.92	<b>2.1x</b>	2,886	5,544	<b>+92%</b>	2,701	3,730	<b>+38%</b>	10
bicg	10	7,936	1,000	6.43	11.27	51.06	11.27	<b>4.5x</b>	2,051	3,229	<b>+57%</b>	2,182	2,737	<b>+25%</b>	10

Table 1. Ayatallah Elakhras, Andrea Guerrieri, Lana Josipovic, and Paolo Ienne. “Survival of the Fastest: Enabling More Out-of-Order Execution in Dataflow Circuits”. In: *Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays*. FPGA '24. Monterey, CA, USA: Association for Computing Machinery, 2024, pp. 44–54. ISBN: 9798400704185. doi: 10.1145/3626202.3637556

# Conclusion and future (current) work

## Conclusion

We verify a rewrite framework for dataflow circuits in Lean 4.

We implement one main rewrite to introduce out-of-order execution.

## Future work

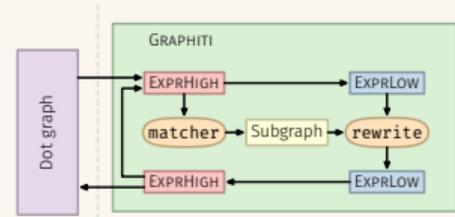
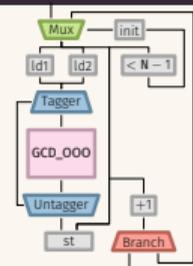
Still a work in progress, finish proofs and gather proper results.

Develop better proof methods for verifying rewrites.

Hopeful we can extend the framework to other dataflow graphs (circuits?).

We can generalise the arbiter idea by defining and out-of-order GCD component.

A more general tagger/untagger pair keeps track of the correct order of tokens.



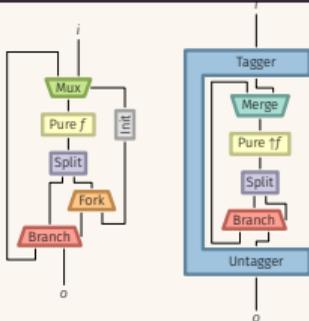
Started collaborating on a verified dataflow rewrite framework. Use the rewrite framework to implement the FPGA'24 algorithm.

Unverified | Verified in Lean 4

# Thanks! Any questions?

In particular we only need one general rewrite replicate many of the FPGA'24 results

Given a sequential loop over an arbitrary loop body  $M$ .  
 We can generate an out-of-order version of the loop by overlapping loop executions.  
 Does this hold for an arbitrary module  $M$ ?  
**No**, we need:  
 1-in-1-out behaviour for the module.  
 It needs to be stateless.



Hypothetical results

Benchmark	N	Cycle count		CF (ns)		Execution time ( $\mu$ s)		LUTs		FFs		DSPs			
		[18]	Ours	[18]	Ours	[18]	Ours	[18]	Ours	[18]	Ours				
img-avg	4	1,722	634	6.42	8.21	11.85	5.21	2.1x	1,415	1,591	+13%	1,320	1,200	-9%	5
gemm_many	10	68,523	32,874	7.91	9.95	541.81	327.0	1.7x	2,835	3,657	+29%	3,256	3,725	+14%	22
gemm_single	10	6,703	9,333	6.90	9.01	46.25	84.11	0.55x	2,736	2,677	-2%	3,142	3,114	-1%	22
germm	20	68,825	8,144	6.75	12.81	464.78	104.30	4.5x	3,214	5,937	+83%	2,693	3,688	+37%	11
matvec	50	7,936	918	6.41	13.51	58.86	12.40	4.1x	3,272	4,396	+246%	1,373	3,423	+149%	5
mv1	10	7,940	2,044	6.27	11.70	49.79	23.92	2.1x	2,866	5,544	+92%	2,701	3,730	+38%	10
hbcg	10	7,936	1,000	6.43	11.27	51.66	11.27	4.5x	2,051	3,229	+57%	2,182	2,737	+25%	10

Table 1. Ayatallah Elabbas, Andrea Guerrieri, Lana Jospovic, and Paolo Ienne. "Survival of the Fastest: Enabling More Out-of-Order Execution in Dataflow Circuits". In: Proceedings of The 2024 ACM/SGDA International Symposium on Field Programmable Gate Arrays, FPGA '24, Monterey, CA, USA: Association for Computing Machinery, 2024, pp. 44-54. isbn: 9789400704185. doi: 10.1145/3626262.3637556

-  Elakhras, Ayatallah, Andrea Guerrieri, Lana Josipovic, and Paolo Ienne. **“Survival of the Fastest: Enabling More Out-of-Order Execution in Dataflow Circuits”**. In: *Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays*. FPGA '24. Monterey, CA, USA: Association for Computing Machinery, 2024, pp. 44–54. ISBN: 9798400704185. DOI: [10.1145/3626202.3637556](https://doi.org/10.1145/3626202.3637556).