



μ CFI: Formal Verification of Microarchitectural Control-flow Integrity

Katharina Ceesay-Seitz, Flavien Solt,
Kaveh Razavi

COMSEC, Computer Security Group,
ETH Zurich

Verification Futures UK, 01.07.2025


COMSEC

ETH zürich



<https://www.linkedin.com/in/katharina-ceesay-seitz-ba521087/>

CPU Verification



Testing, e.g., fuzzing

CPU Verification




Testing, e.g., fuzzing,
is incomplete



Security: need guarantee of
absence of bugs


CPU Verification



Formal verification:


- Provides formal guarantees for all inputs

CPU Verification




Formal verification:

- Provides formal guarantees for all inputs




- Often a CPU-specific, manual effort

Formal Property Verification

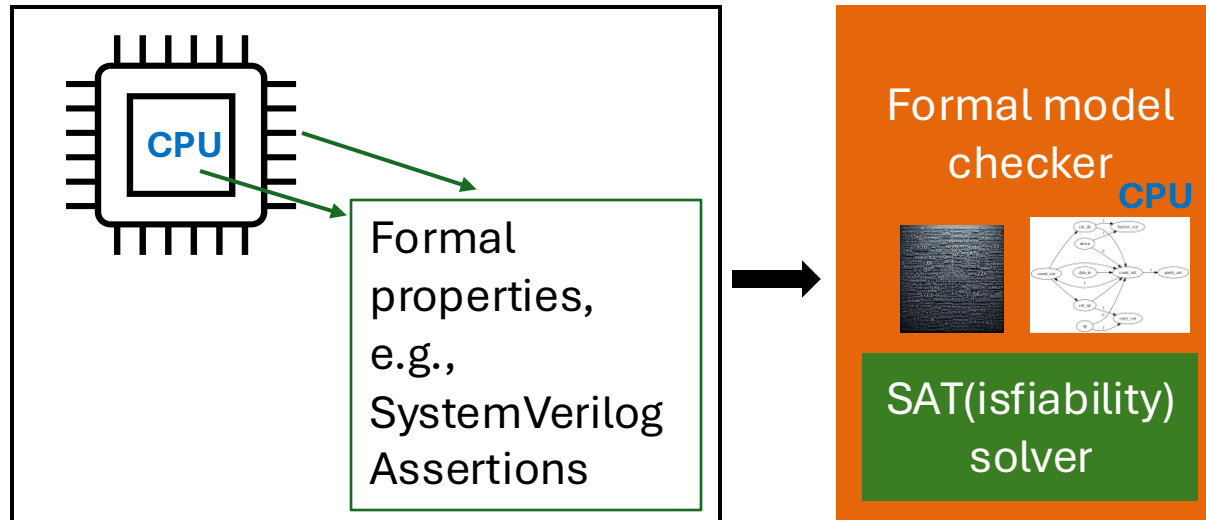


HDL
(Hardware
Description
Language)
Design

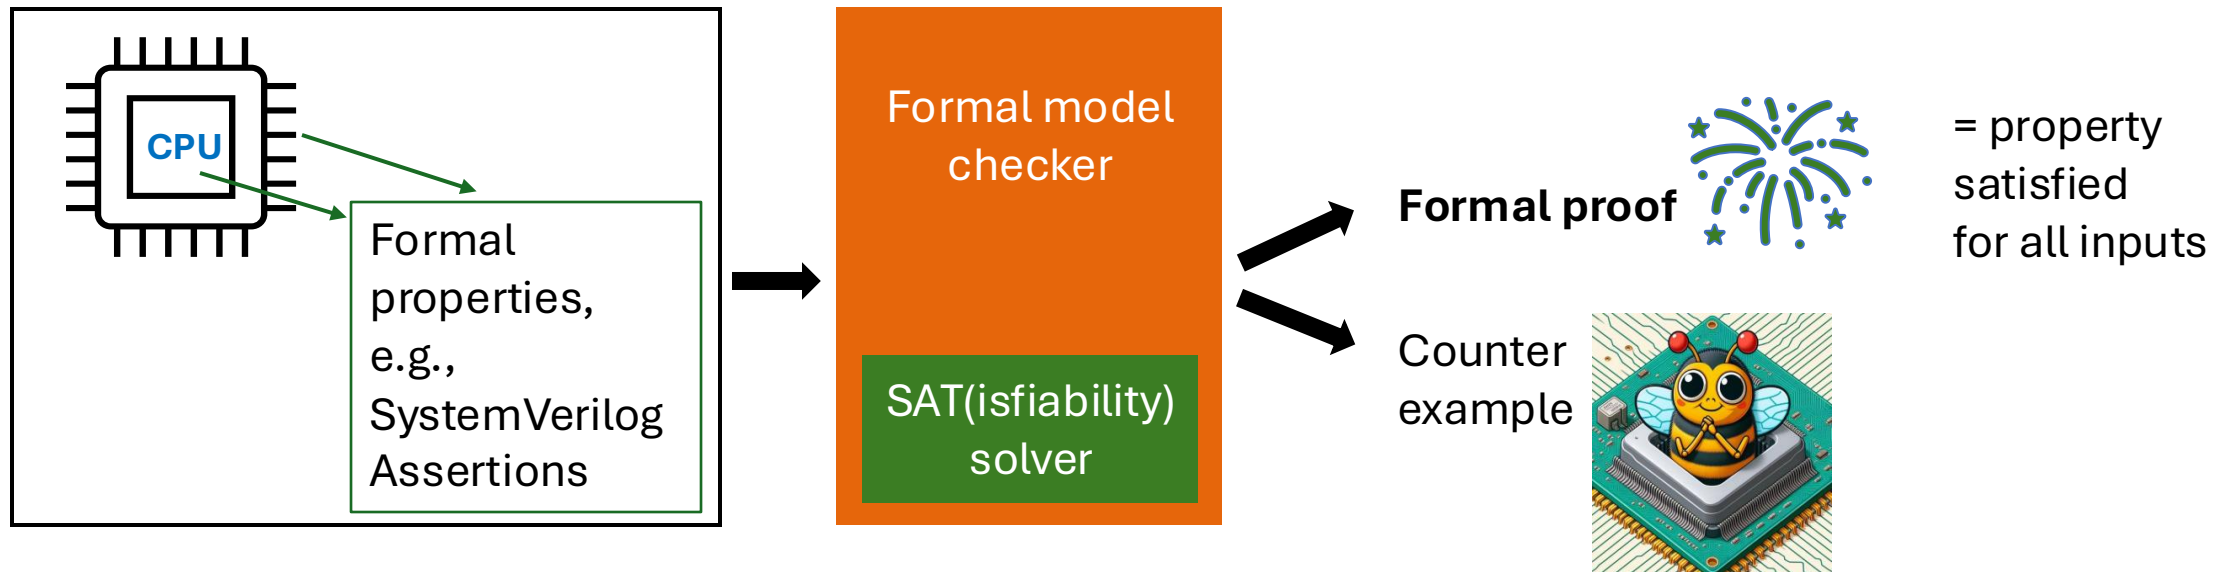
Formal properties,
e.g.,
SystemVerilog
Assertions
describe desired
behavior




Formal Property Verification



Formal Property Verification






SIMPLER SOLUTION?



CPU Verification



Formal verification:

- Provides formal guarantees for all inputs

μCFI - Generalized security property

- Easy application and reuse
- Independent of CPU's verification state
=> apply it early in the design cycle
- Captures multiple threat models


Definition: Architectural Control Flow

Architectural
(software)
Program
Counter
(PC)



Software program (assembly instructions)

```
80000000 <_start>:  
80000000: 00010337      lui t1,0x10  
80000004: 010eaf83      lw  t6,16(t4)  
80000008: 01f32823      sw  t6,16(t1)  
8000000c: 400b0b13      addi s6,s6,1024  
80000010: 34319073      csrw mtval,gp  
80000014: 341020f3      csrr ra,mepc  
80000018: 0030c133      xor sp,ra,gp
```



CPU

Architectural PC decides the order of instructions

Software 'if'
=
Branch instruction

```
If condition  
    PC = Branch target = A  
Else  
    PC = Branch target = PC + 4
```


Definition: Microarchitectural Control Flow (μ CF)

Architectural
(software)
Program
Counter
(PC)




Software program (assembly instructions)

```
80000000 <_start>:  
80000000: 00010337      lui t1,0x10  
80000004: 010eaf83      lw  t6,16(t4)  
80000008: 01f32823      sw  t6,16(t1)  
8000000c: 400b0b13      addi s6,s6,1024  
80000010: 34319073      csrw mtval,gp  
80000014: 341020f3      csrr ra,mepc  
80000018: 0030c133      xor sp,ra,gp
```



CPU

Microarchitectural control flow (μ CF)



update time

Microarchitectural PC
= a register inside
the CPU

0x80000004

0x80000008


0x80001000

value

Microarchitectural Control Flow Violations

Constant Time (CT) RISC-V program

Architectural control flow




reads
secret
data


Microarchitectural Control Flow Violations

Constant Time (CT) RISC-V program

Architectural control flow





reads
secret
data



Microarchitectural control flow



operand: 0



Microarchitectural Control Flow Violations


Constant Time (CT) program

Architectural control flow



Microarchitectural control flow

operand: 0



operand: '1





PROBLEM?



Microarchitectural Control Flow Violations



Constant Time (CT) program

Architectural control flow



Microarchitectural control flow

operand: 0



operand: '1



Secret influences μ CF

Execution takes longer = timing side channel


Delayed PC update



Microarchitectural Control Flow Violations


Constant Time (CT) program

Architectural control flow



Microarchitectural control flow

operand: :0




operand: '1



Control-flow integrity secure program

Architectural control flow





reads
attacker
controlled
input
data

Exception

Microarchitectural Control Flow Violations


Constant Time (CT) program

Architectural control flow



Microarchitectural control flow

operand: :0



clock

operand: '1





clock




Control-flow integrity secure program

Architectural control flow




Microarchitectural control flow

operand: 0x80



clock




Input influences μ CF
by changing PC value



μCFI - Microarchitectural Control-flow Integrity

- Proof that the μCF has only ISA-specified data dependencies
- Detect non-ISA specified flows




μCFI - Microarchitectural Control-flow Integrity

- Proof that the μCF has only ISA-specified data dependencies
- Detect non-ISA specified flows



One property

Operand



Two threat
models
captured



μCFI - Microarchitectural Control-flow Integrity

- Proof that the μCF has only ISA-specified data dependencies
- Detect non-ISA specified flows




One property

Two threat
models
captured

Operand

secret



PC



Information leakage




μCFI - Microarchitectural Control-flow Integrity

- Proof that the μCF has only ISA-specified data dependencies
- Detect non-ISA specified flows



One property

Operand



≠ PC

**Two threat
models
captured**

secret

attacker-
controlled

Information leakage

Control-flow hijack




μCFI - Microarchitectural Control-flow Integrity

- Proof that the μCF has only ISA-specified data dependencies
- Detect non-ISA specified flows



Information flow property




μCFI - Microarchitectural Control-flow Integrity

- Proof that the μCF has only ISA-specified data dependencies
- Detect non-ISA specified flows



Information flow property



Information =

≠ data flows

🕒 time & control flows




μCFI - Microarchitectural Control-flow Integrity

- Proof that the μCF has only ISA-specified data dependencies
- Detect non-ISA specified flows



Information flow property



Information =



data flows




time & control flows



Formal Verification of Information Flow

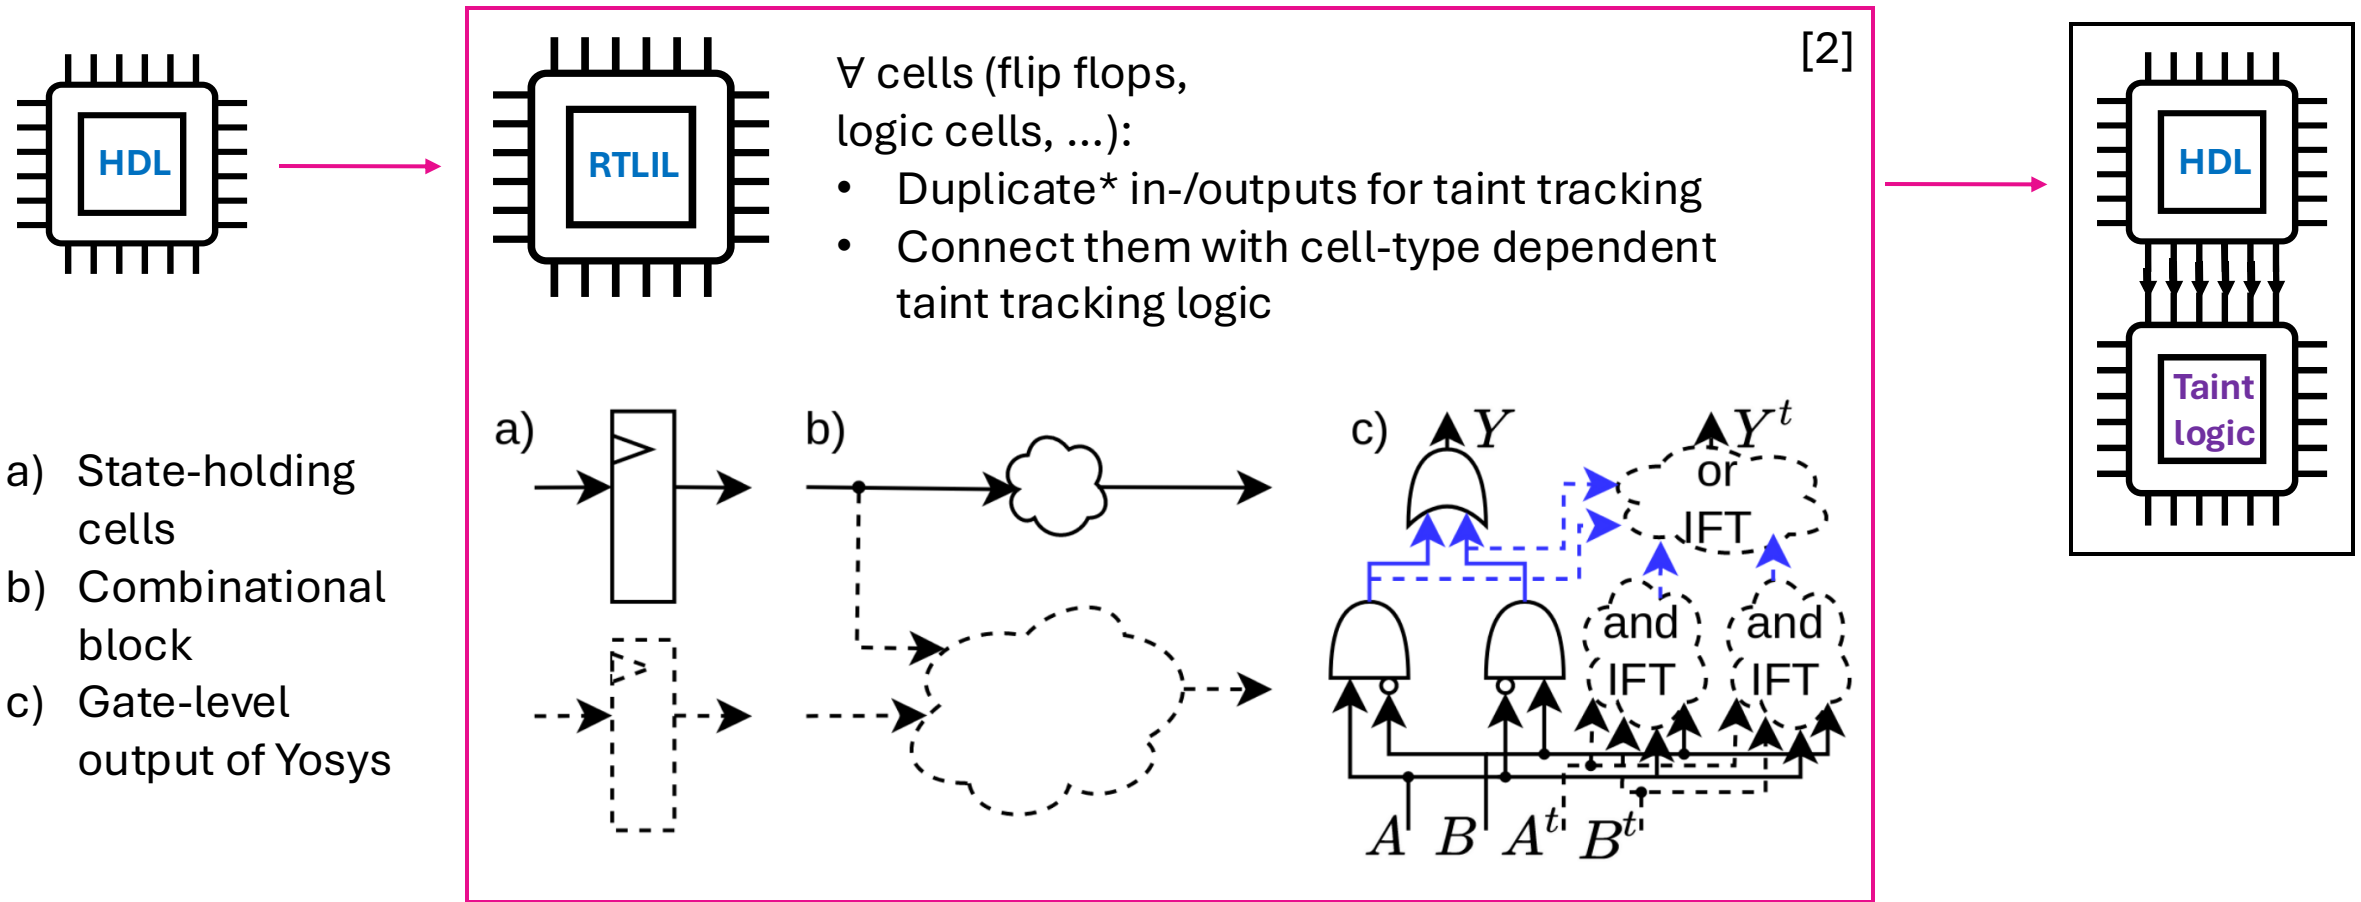
Information flow tracking with
taint logic – CellIFT [1]



taint = secret or attacker-controlled information

[1] F. Solt, B. Gras, K. Razavi, "CELLIFT: Leveraging Cells for Scalable and Precise Dynamic Information Flow Tracking in RTL", USENIX Security 2022

CellIFT Yosys [1] pass




*it is possible to add multiple independent taint instrumentations. Each in-/output gets a taint representation per instrumentation.

[1] Yosys Open SYnthesis Suite - <https://github.com/YosysHQ/yosys>


[2] F. Solt, B. Gras, K. Razavi, "CELLIFT: Leveraging Cells for Scalable and Precise Dynamic Information Flow Tracking in RTL", USENIX Security 2022

CellIFT




Taint logic (CellIFT [1]) tracks information flows

Information flow tracking with **taint logic** – CellIFT [1]




taint = **secret** or **attacker-controlled**

CellIFT




Taint logic (CellIFT [1]) tracks information flows

Information flow tracking with **taint logic** – CellIFT [1]




taint = **secret** or **attacker-controlled**

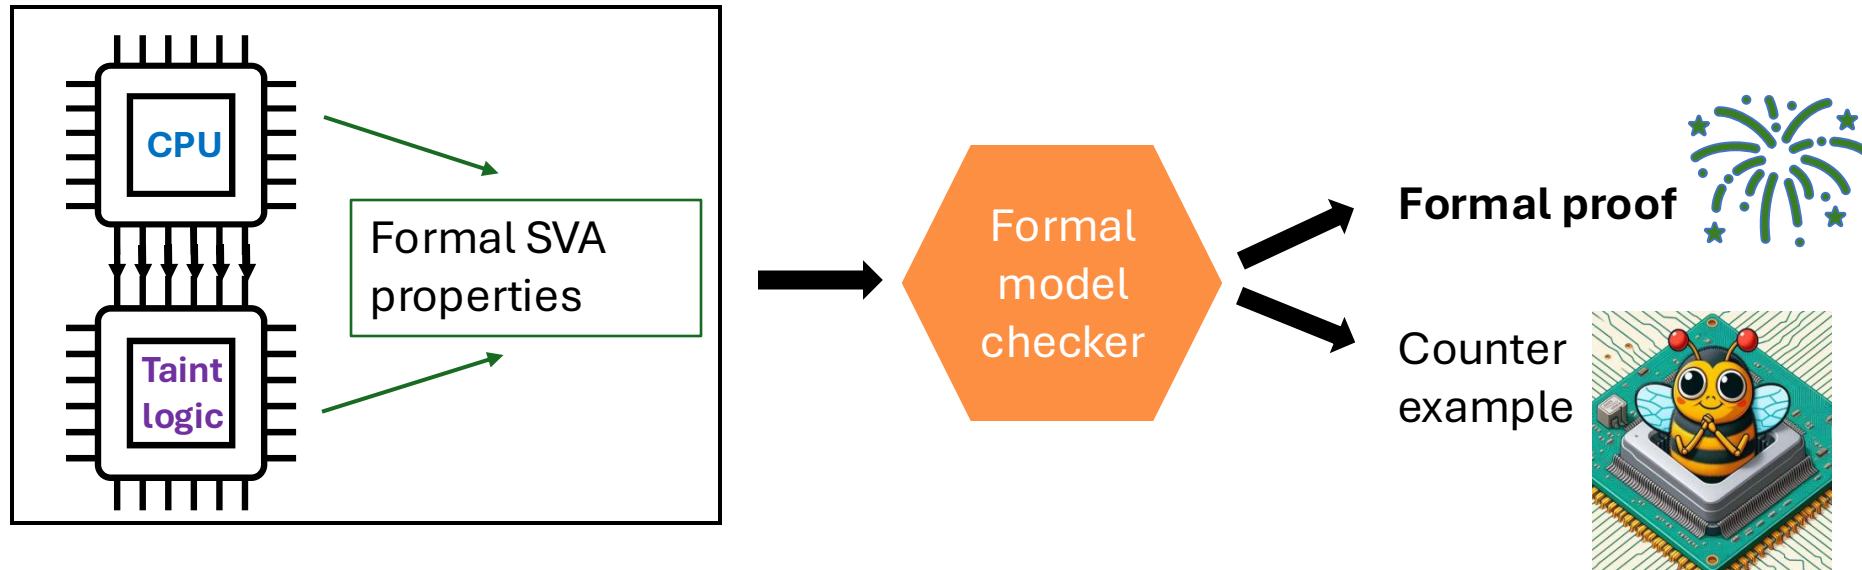
Formal Verification of Information Flow



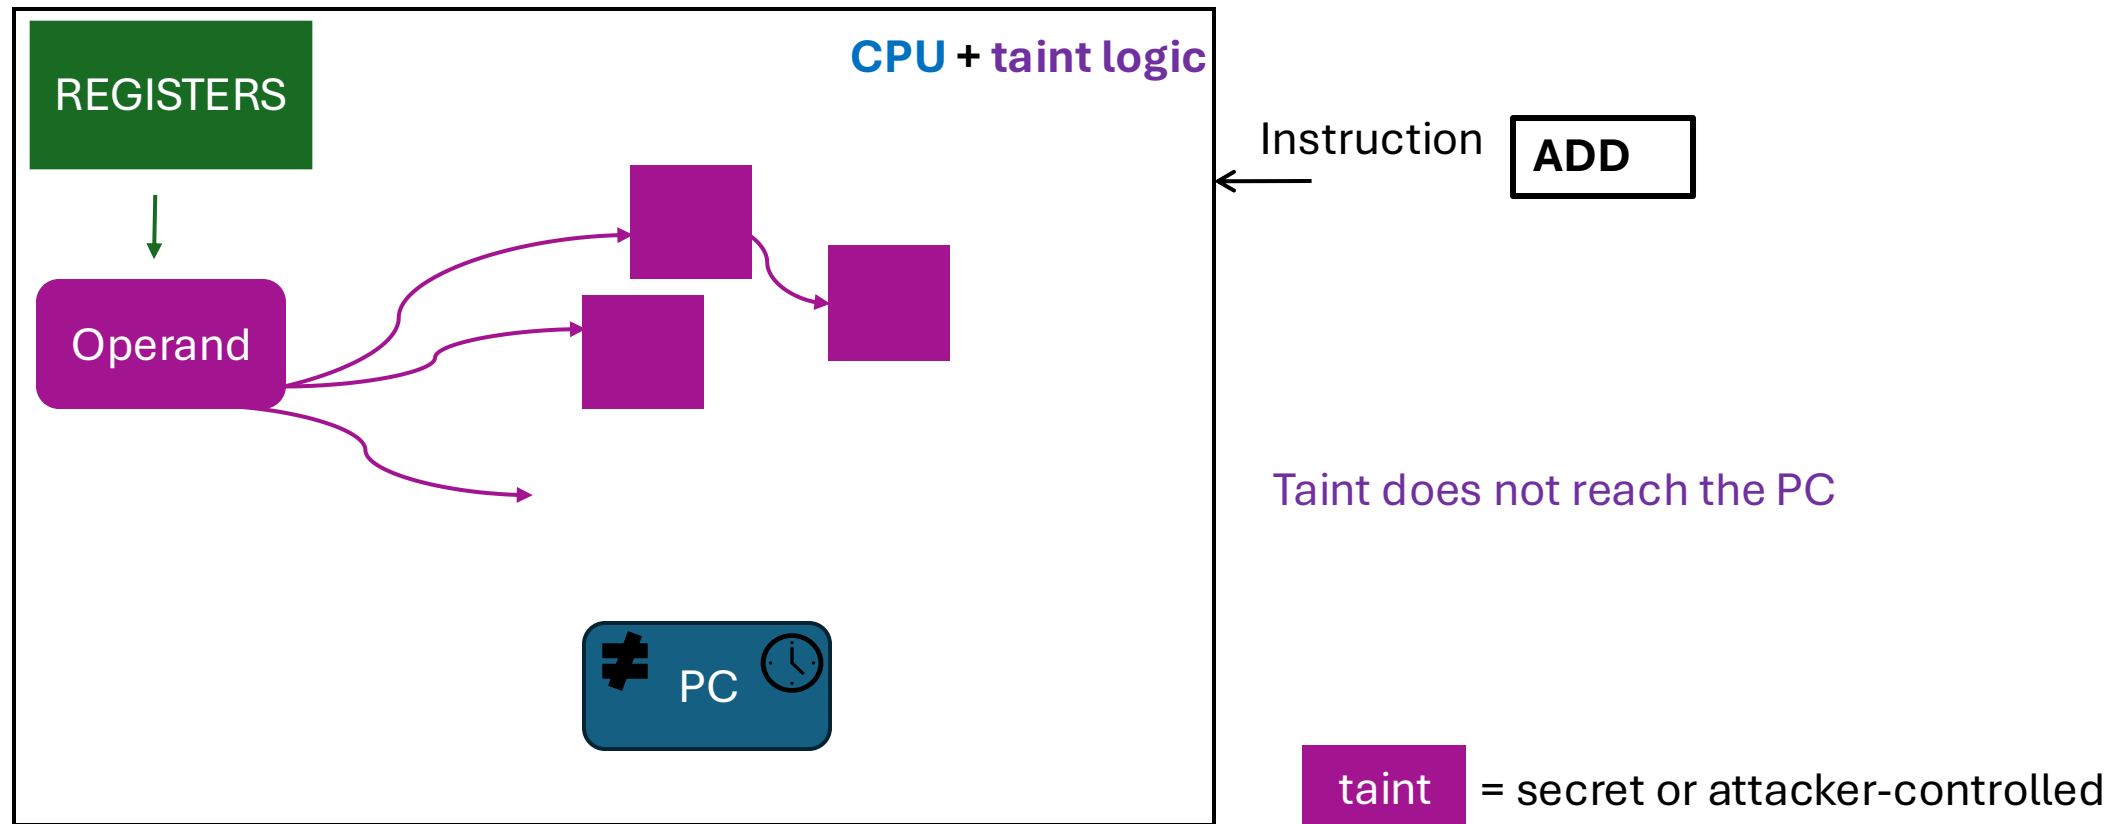
Formal Verification of Information Flow



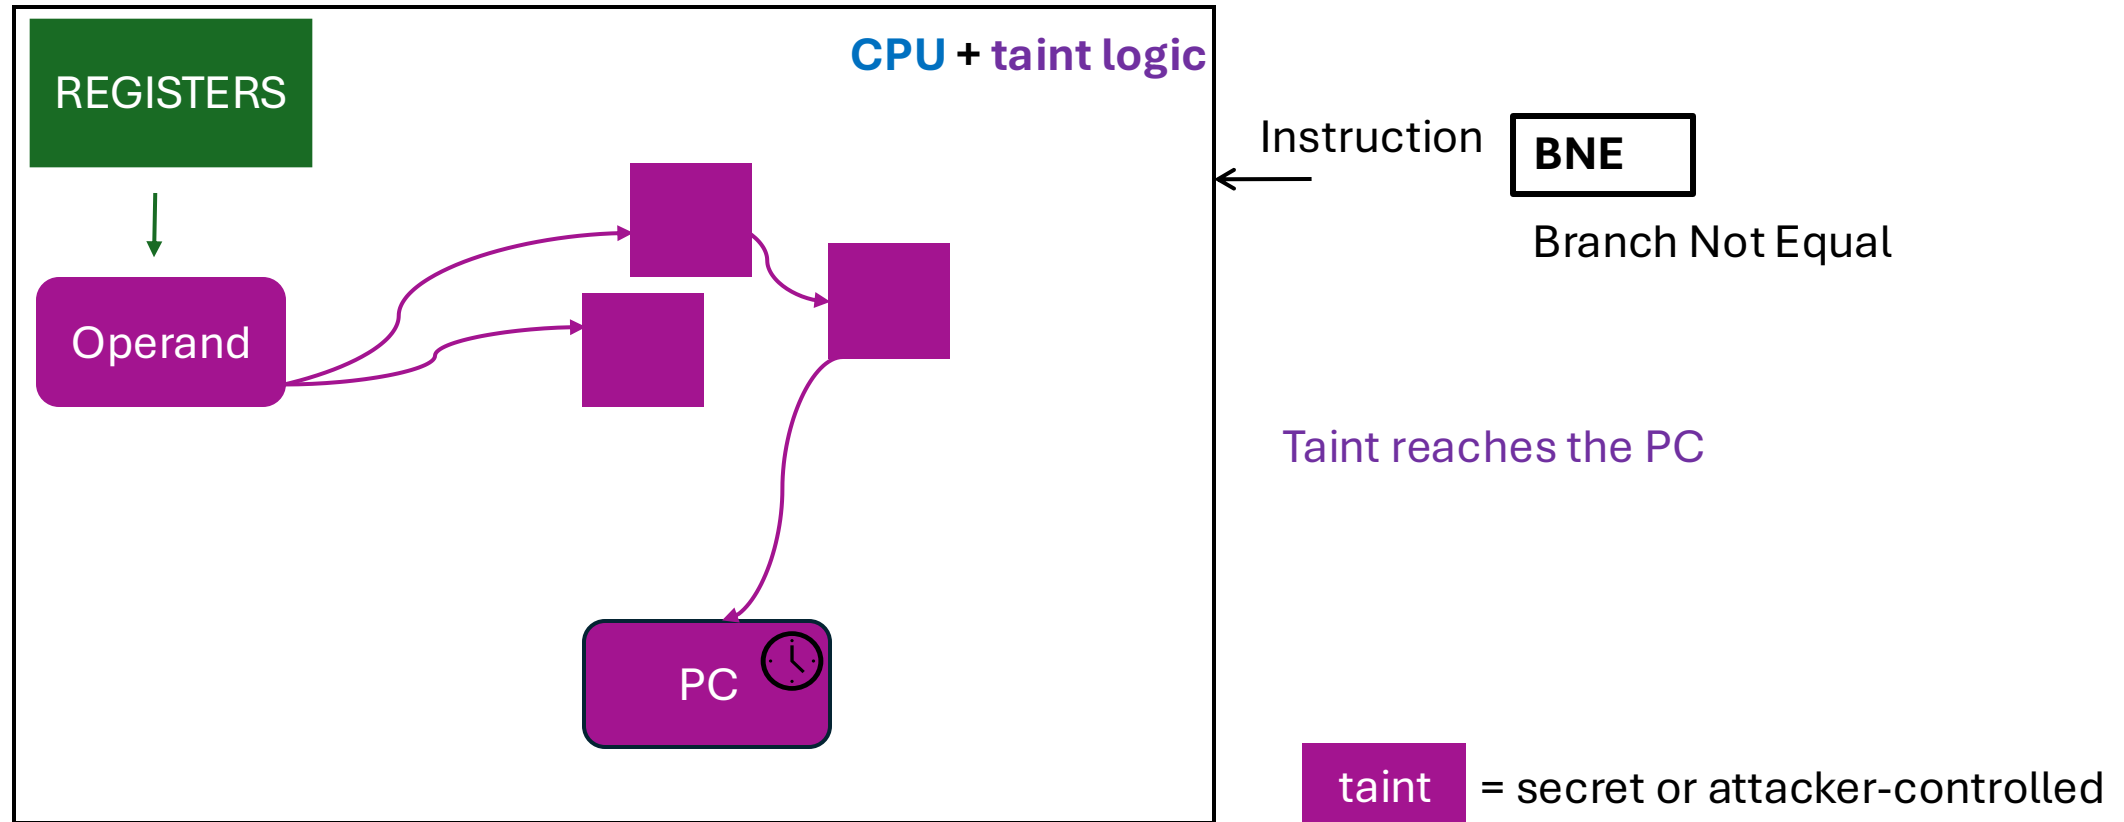
Formal Verification of Information Flow



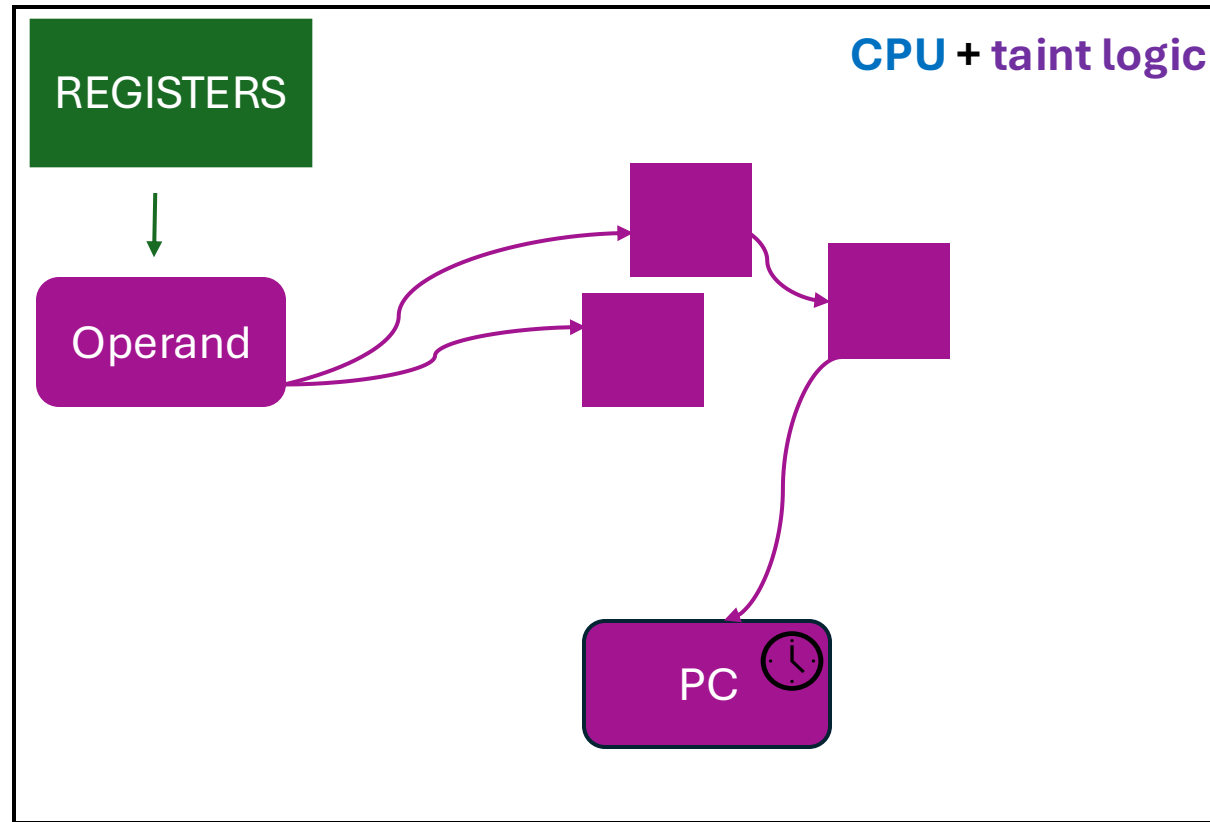
Formally Verifying μ CFI



Formally Verifying μ CFI



Formally Verifying μ CFI



Instruction **BNE**
Branch Not Equal

Taint reaches the PC


μ CFI violated??

Instruction Classification

Control-influencing:

direct branches,
instructions with
exceptions, ...

are expected
to influence
the program counter



```
If reg[t1] == reg[t2]
    Branch target = A
Else
    Branch target = PC + 4
```

Instruction Classification

Control-influencing:

branches,
instructions with
exceptions, ...

are expected
to influence
the program counter



via implicit/control paths only

```
If reg[t1] == reg[t2]
    Branch target = A
Else
    Branch target = PC + 4
```

not via data paths

~~Program Counter = reg[t1]~~
~~Program Counter = reg[t2]~~


Instruction Classification

Control-influencing:


branches,
instructions with
exceptions, ...

are expected
to influence
the program counter

via implicit/control
paths only




```
If reg[t1] == reg[t2]
    Branch target = A
Else
    Branch target = PC + 4
```



~~Program Counter = reg[t1]
Program Counter = reg[t2]~~

CellIFT



CellIFT [1]


tracks information =




data,




control & timing flows



Non-influencing:
arithmetic, logic, ...






CellDFT – Data Flow Tracking



CellIFT [1]
tracks information =
 data,
 control & timing flows

New:




CellDFT – Data Flow Tracking

CellDFT

- Only tracks **data flows** \neq
- Blocks control flows

CellDFT rules implemented in Yosys pass

Cell Name	Definition
State elems. with enable (EN)	$Q_n^t = (EN \wedge D^t) \vee (\neg EN \wedge Q_{n-1}^t)$
2-input mux, aldff cells [106]	$Y^t = (\neg S \wedge A^t) \vee (S \wedge B^t)$
pmux cells [106]	$Y^t = A[S]^t$
Comparison/reduction cells	$Y^t = 0$
Shift cells	$Y^t = A^t \circ B$



Data taint propagates

No taint propagates

Data taint is shifted

Instruction classification

Non-influencing:
arithmetic, logic, ...

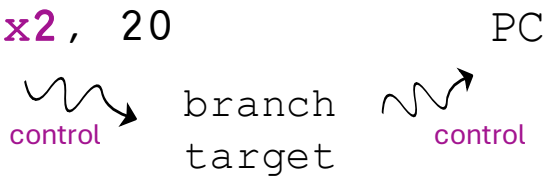
add x4, x5, x6



PC

Control-influencing:
branches, exceptions

bne x1, x2, 20




branch target

PC

Value-influencing:
jumps



jalr ra, x1, 80




jump target

PC

μCFI



μCFI - Verification Goals

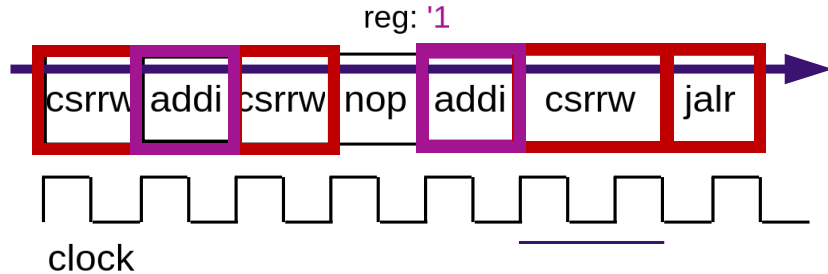


For communication with software engineers/tools:

- Security classification per instruction

Is a constant time (CT) software really CT on an actual CPU implementation?


μCFI - Verification Goals



For communication with software engineers/tools:

- Security classification per instruction,
- surrounded by arbitrary, **potentially insecure**, instructions

μCFI - Verification Goals




For communication with software:

- Security classification per instruction

To ease debugging:

- Identify the specific instruction that leaks

μCFI - Verification Goals



..... ∞

For communication with software:

- Security classification per instruction

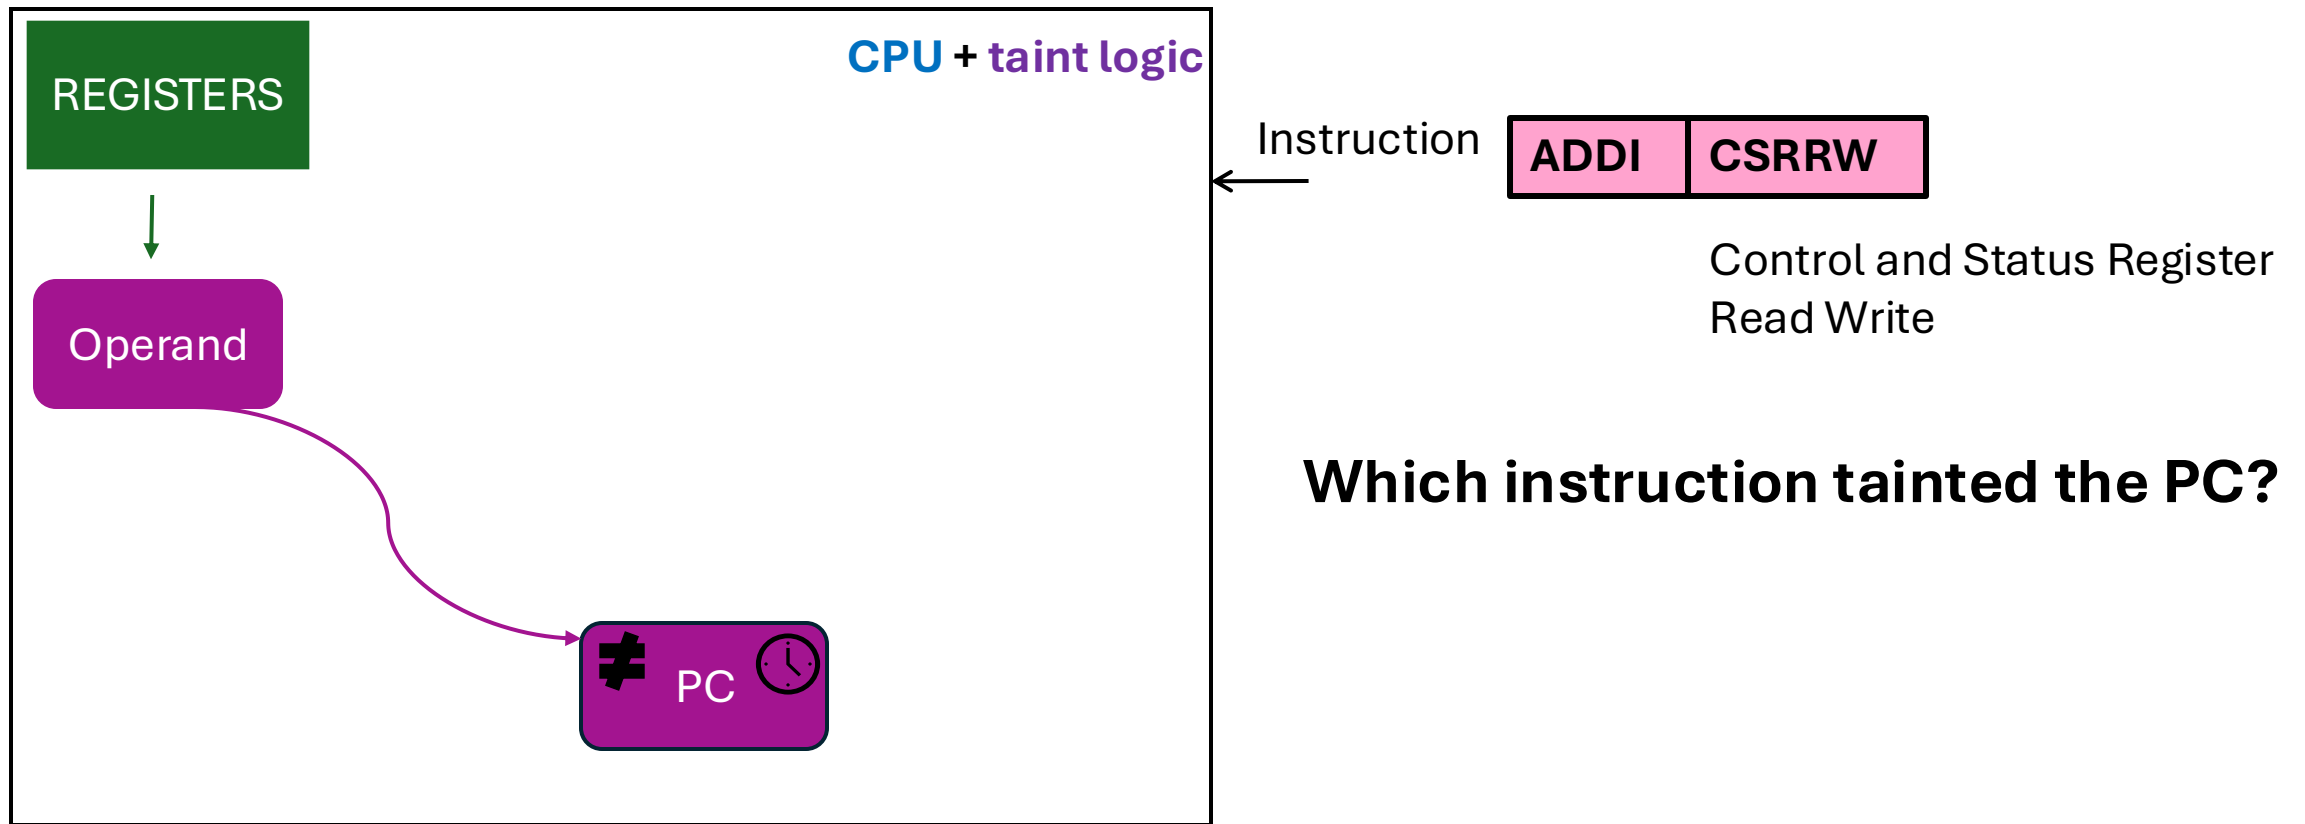
To ease debugging:

- Identify the specific instruction that leaks

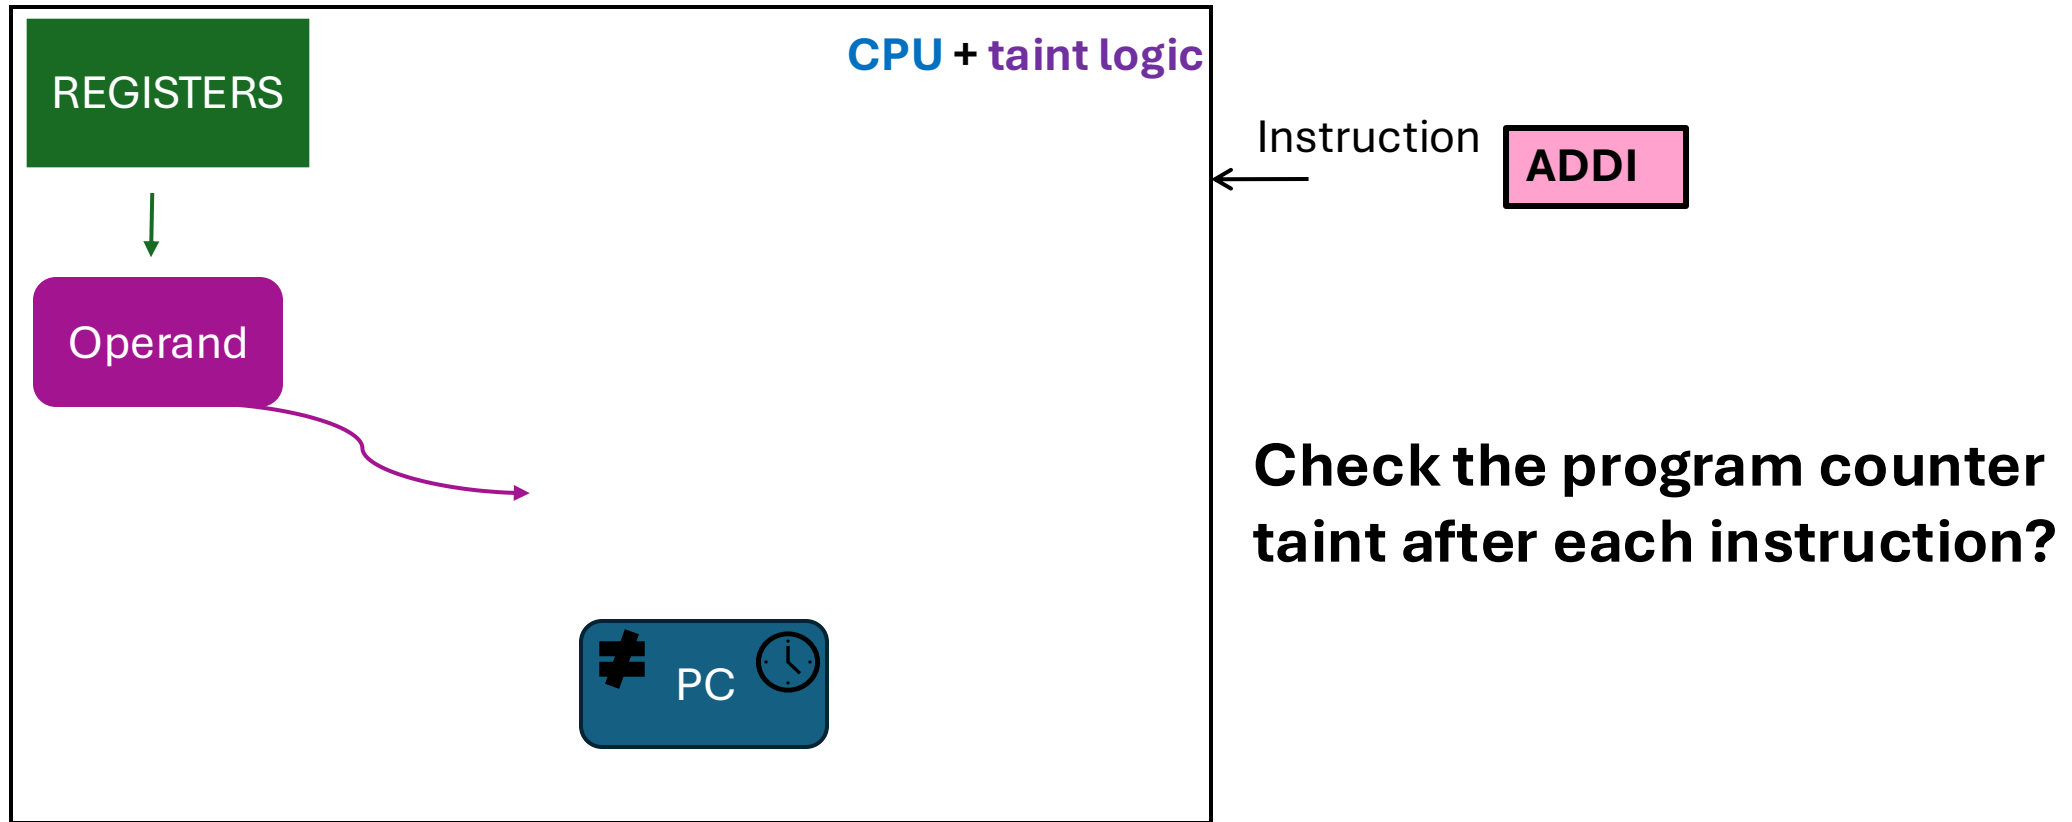
For strong security guarantees:

- consider influences on younger instructions
- over arbitrary, infinitely long programs

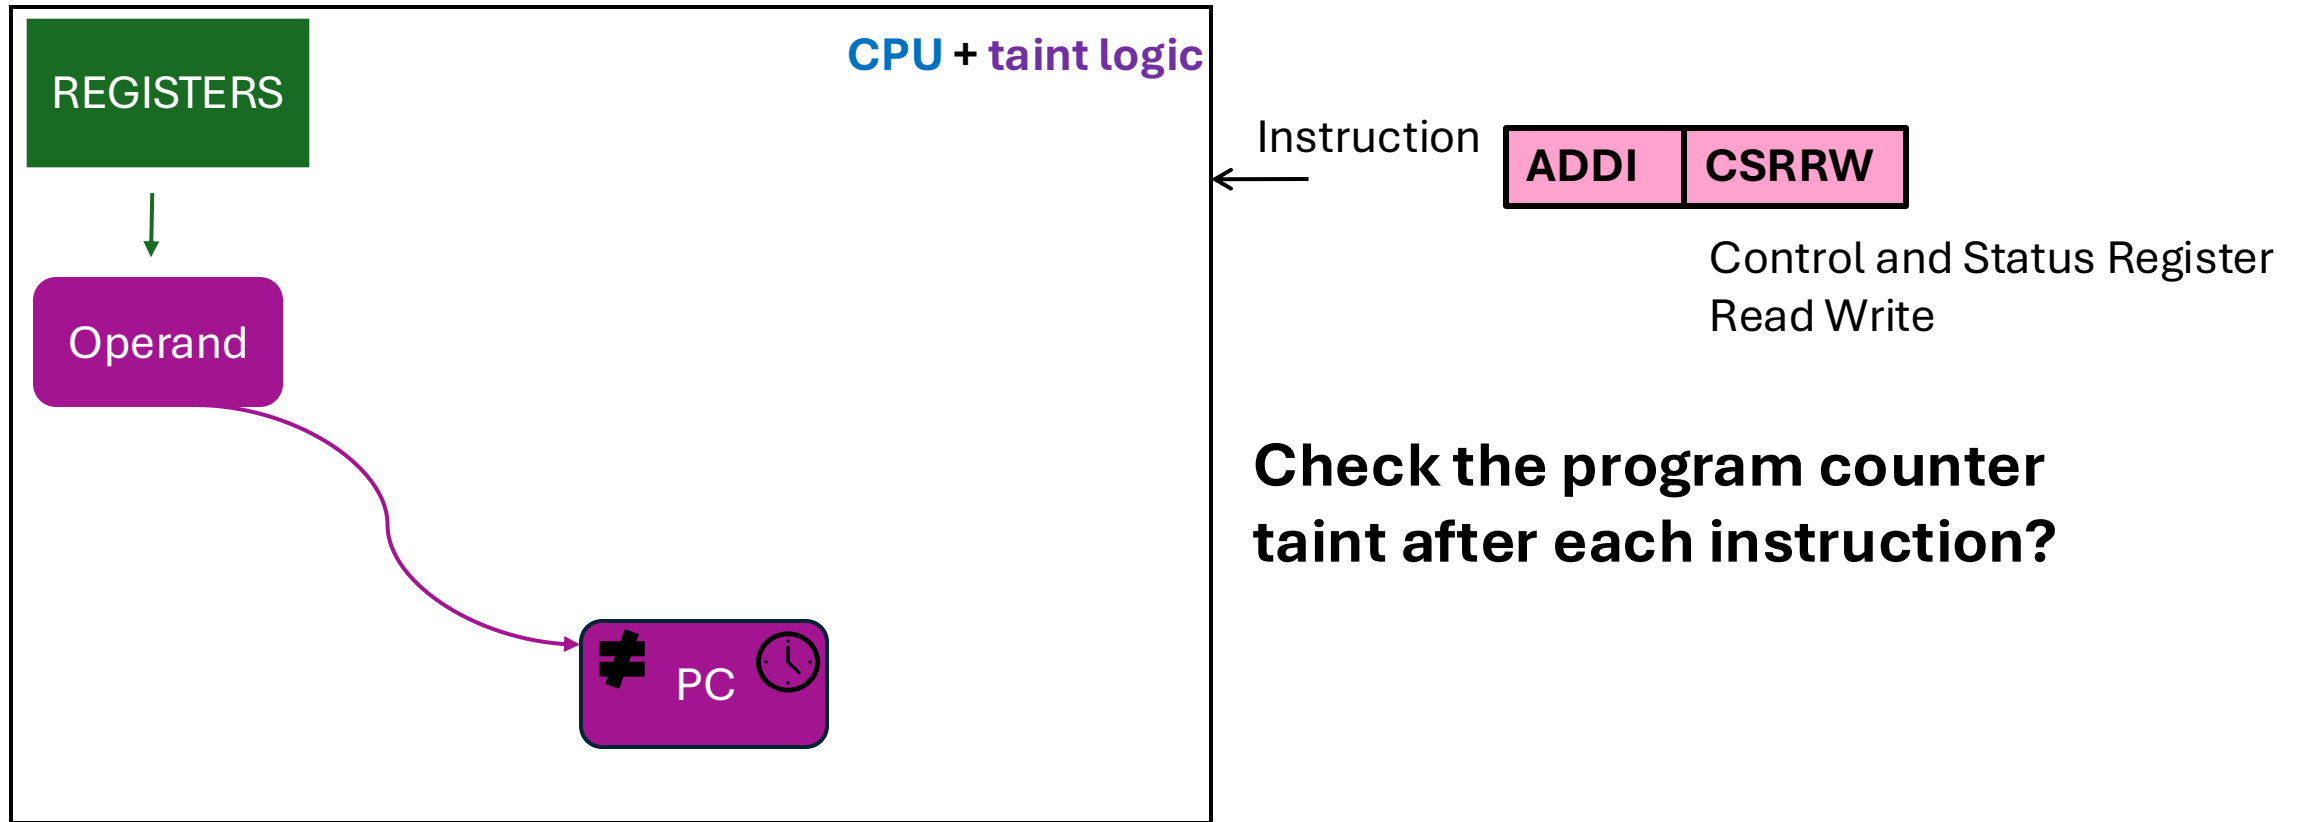
Identifying Insecure Instructions




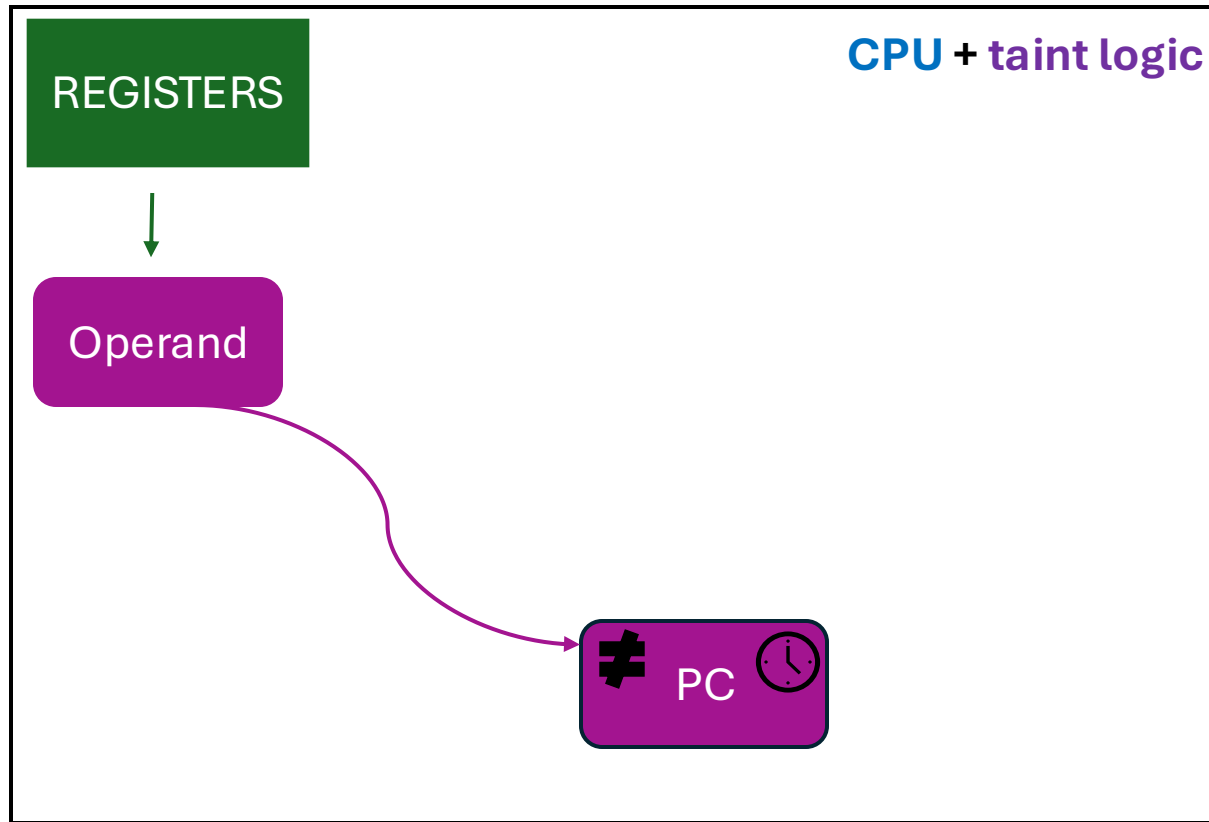
Identifying Insecure Instructions



Identifying Insecure Instructions


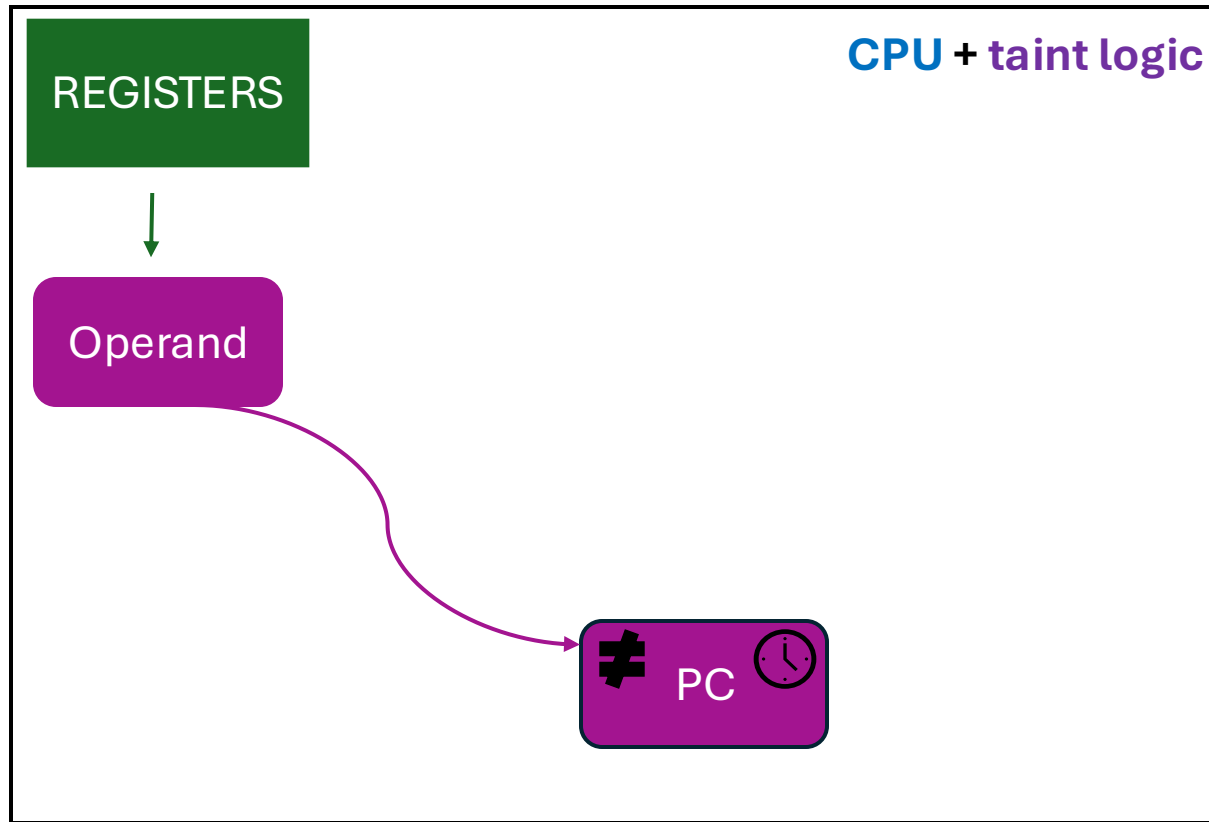


Identifying Insecure Instructions



- **addi's** operand leaks, but violation is associated with csrrw
- Checking program counter taint after each instruction is **imprecise**

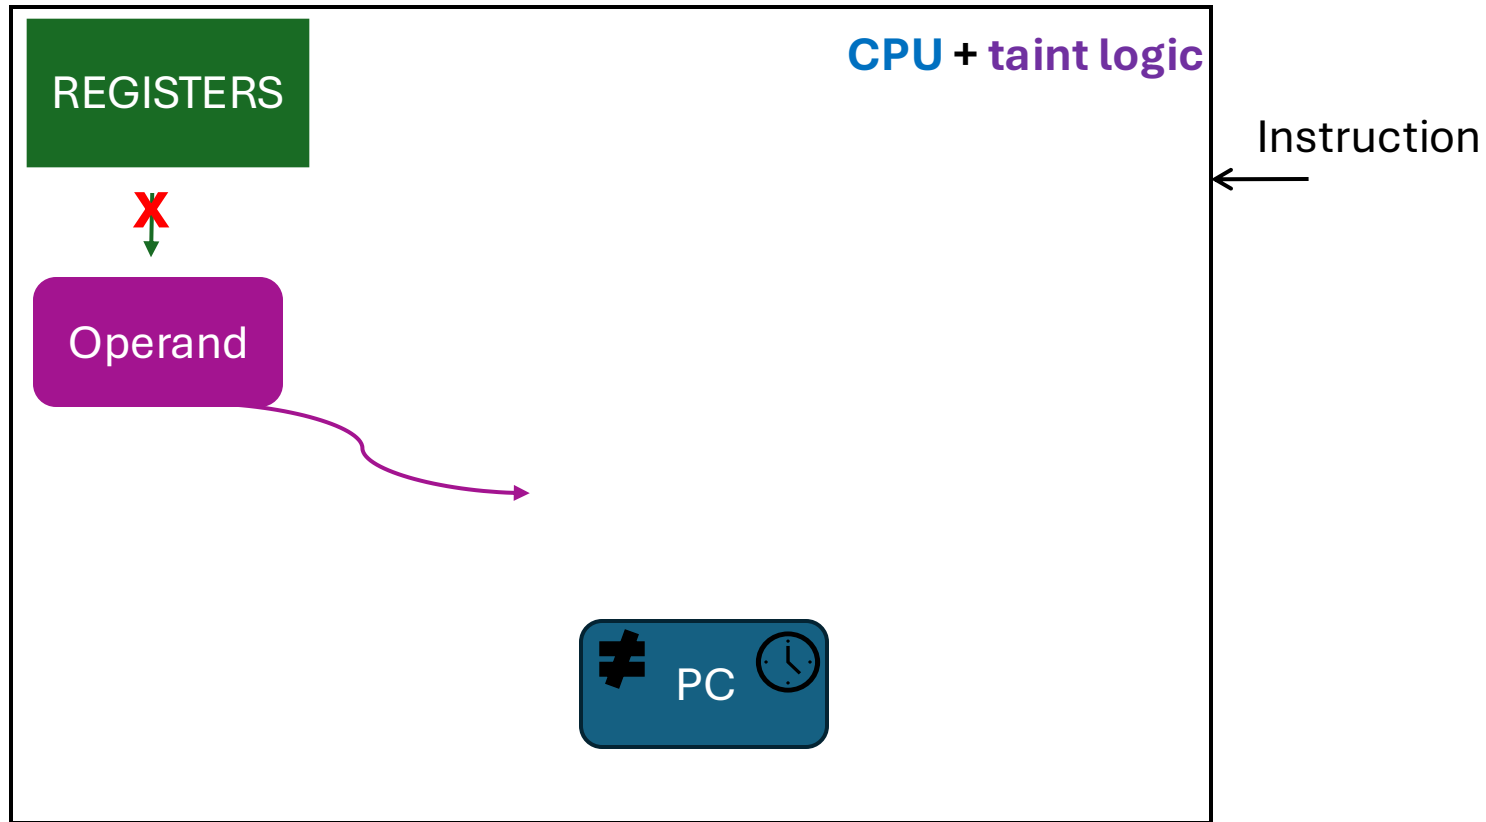
Identifying Insecure Instructions



- **addi's** operand leaks, but violation is associated with csrrw
- Checking program counter taint after each instruction is **imprecise**
- A bug may be hidden by csrrw's specified information flow

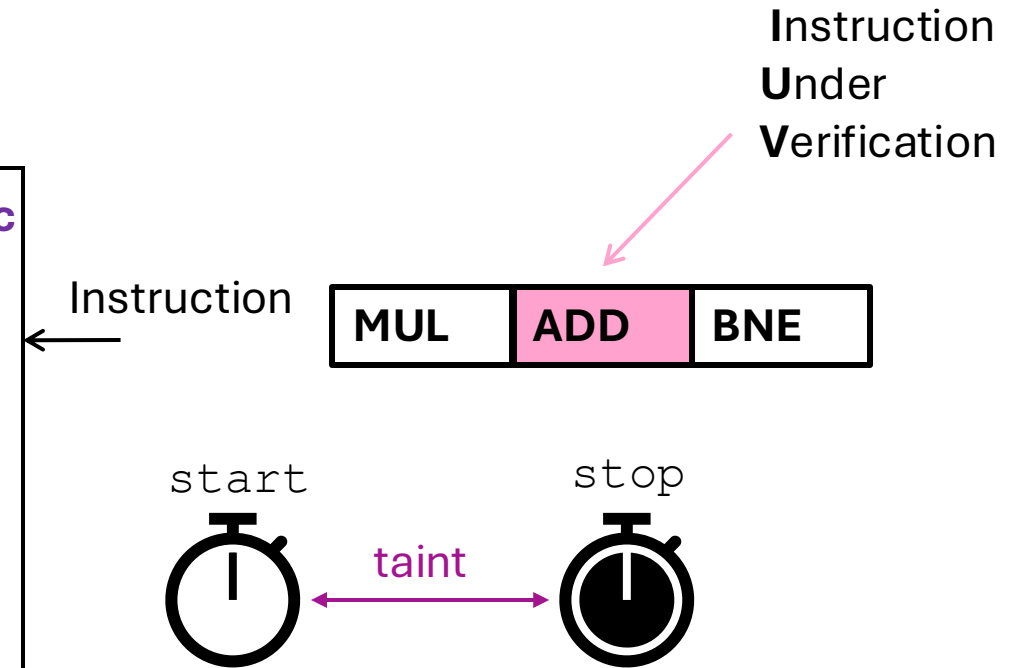
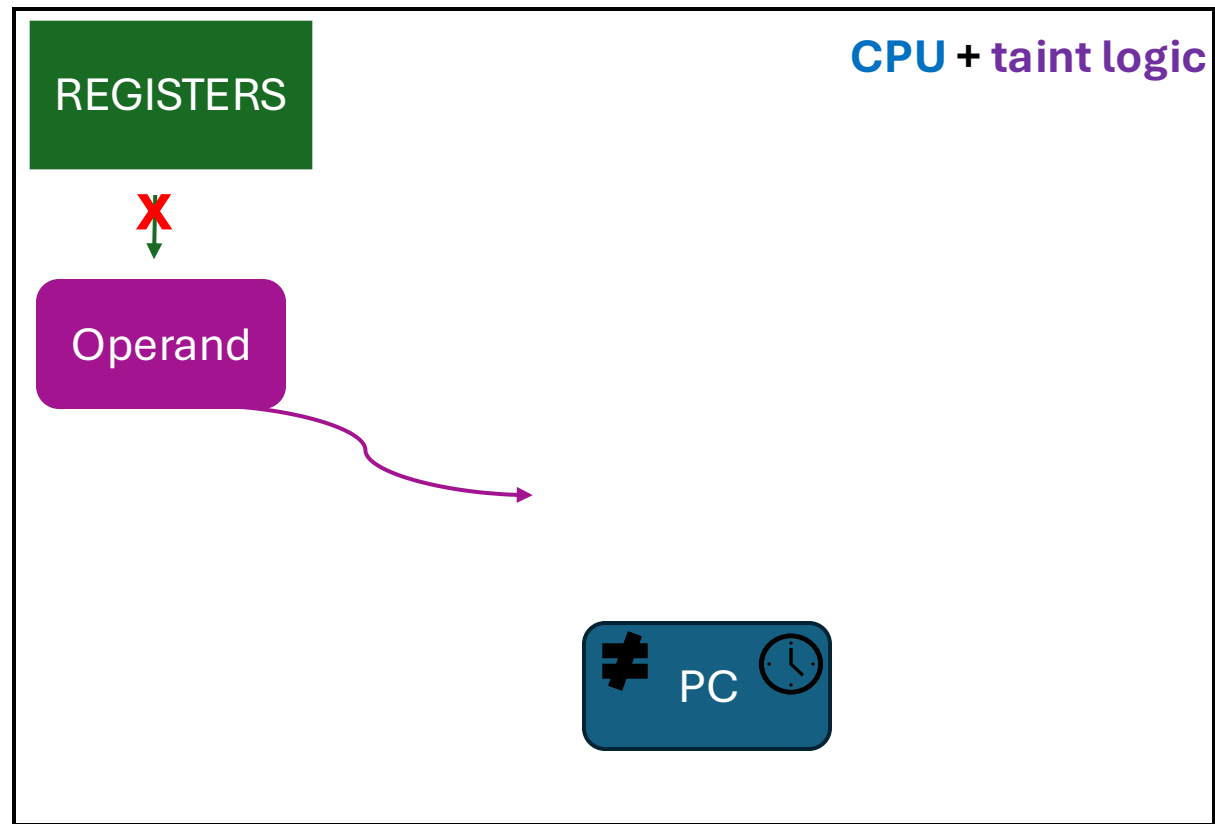
Precise Taint Injection

x = (taint) logic abstraction




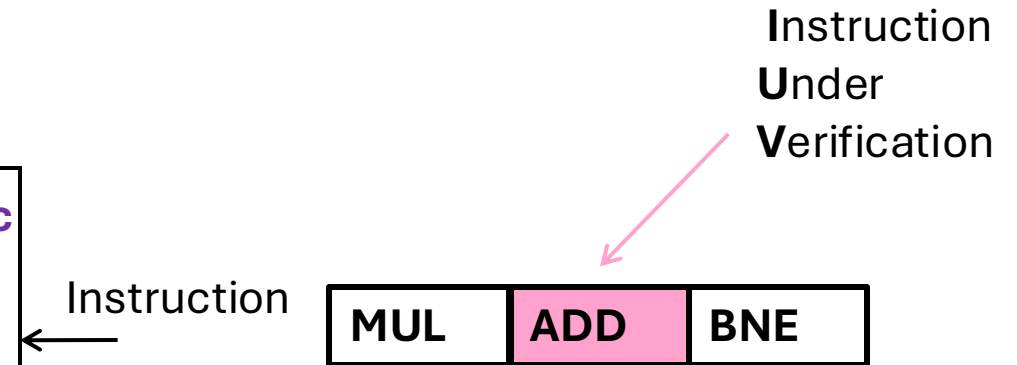
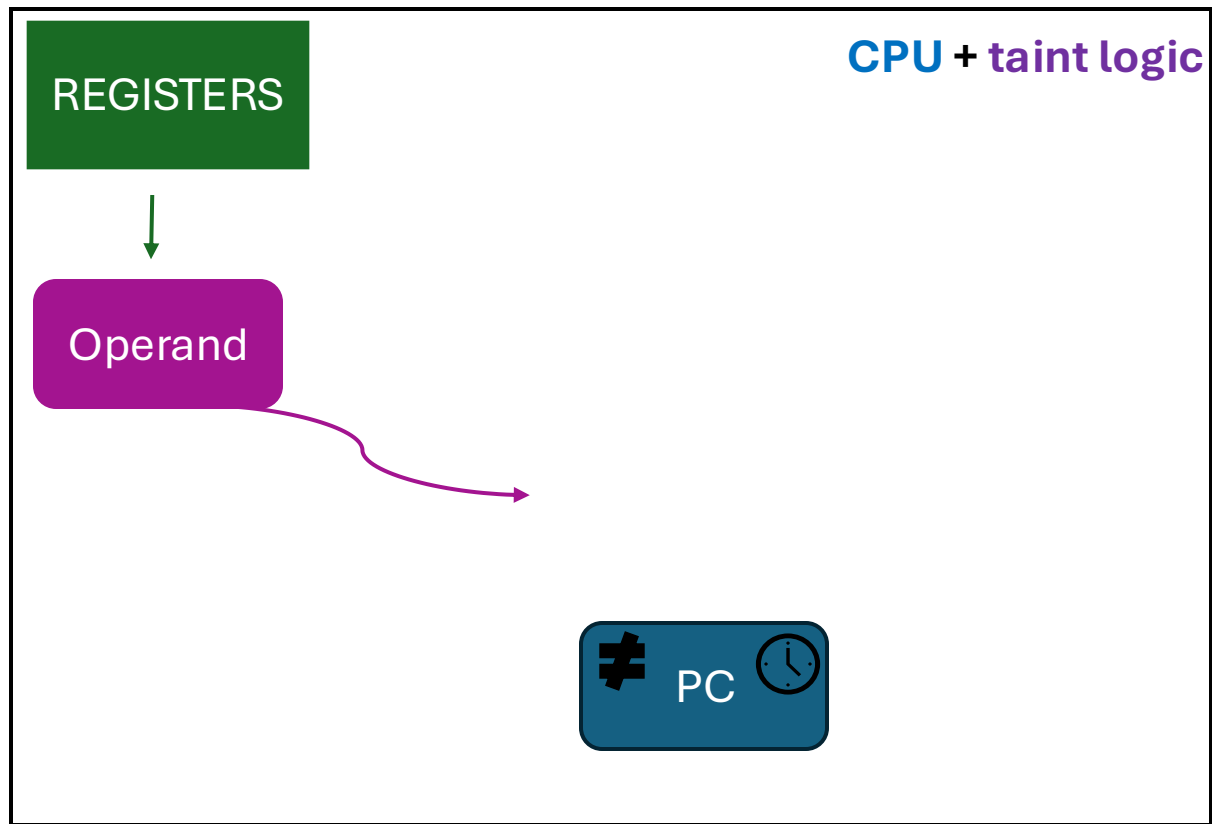
Precise Taint Injection

x = (taint) logic abstraction



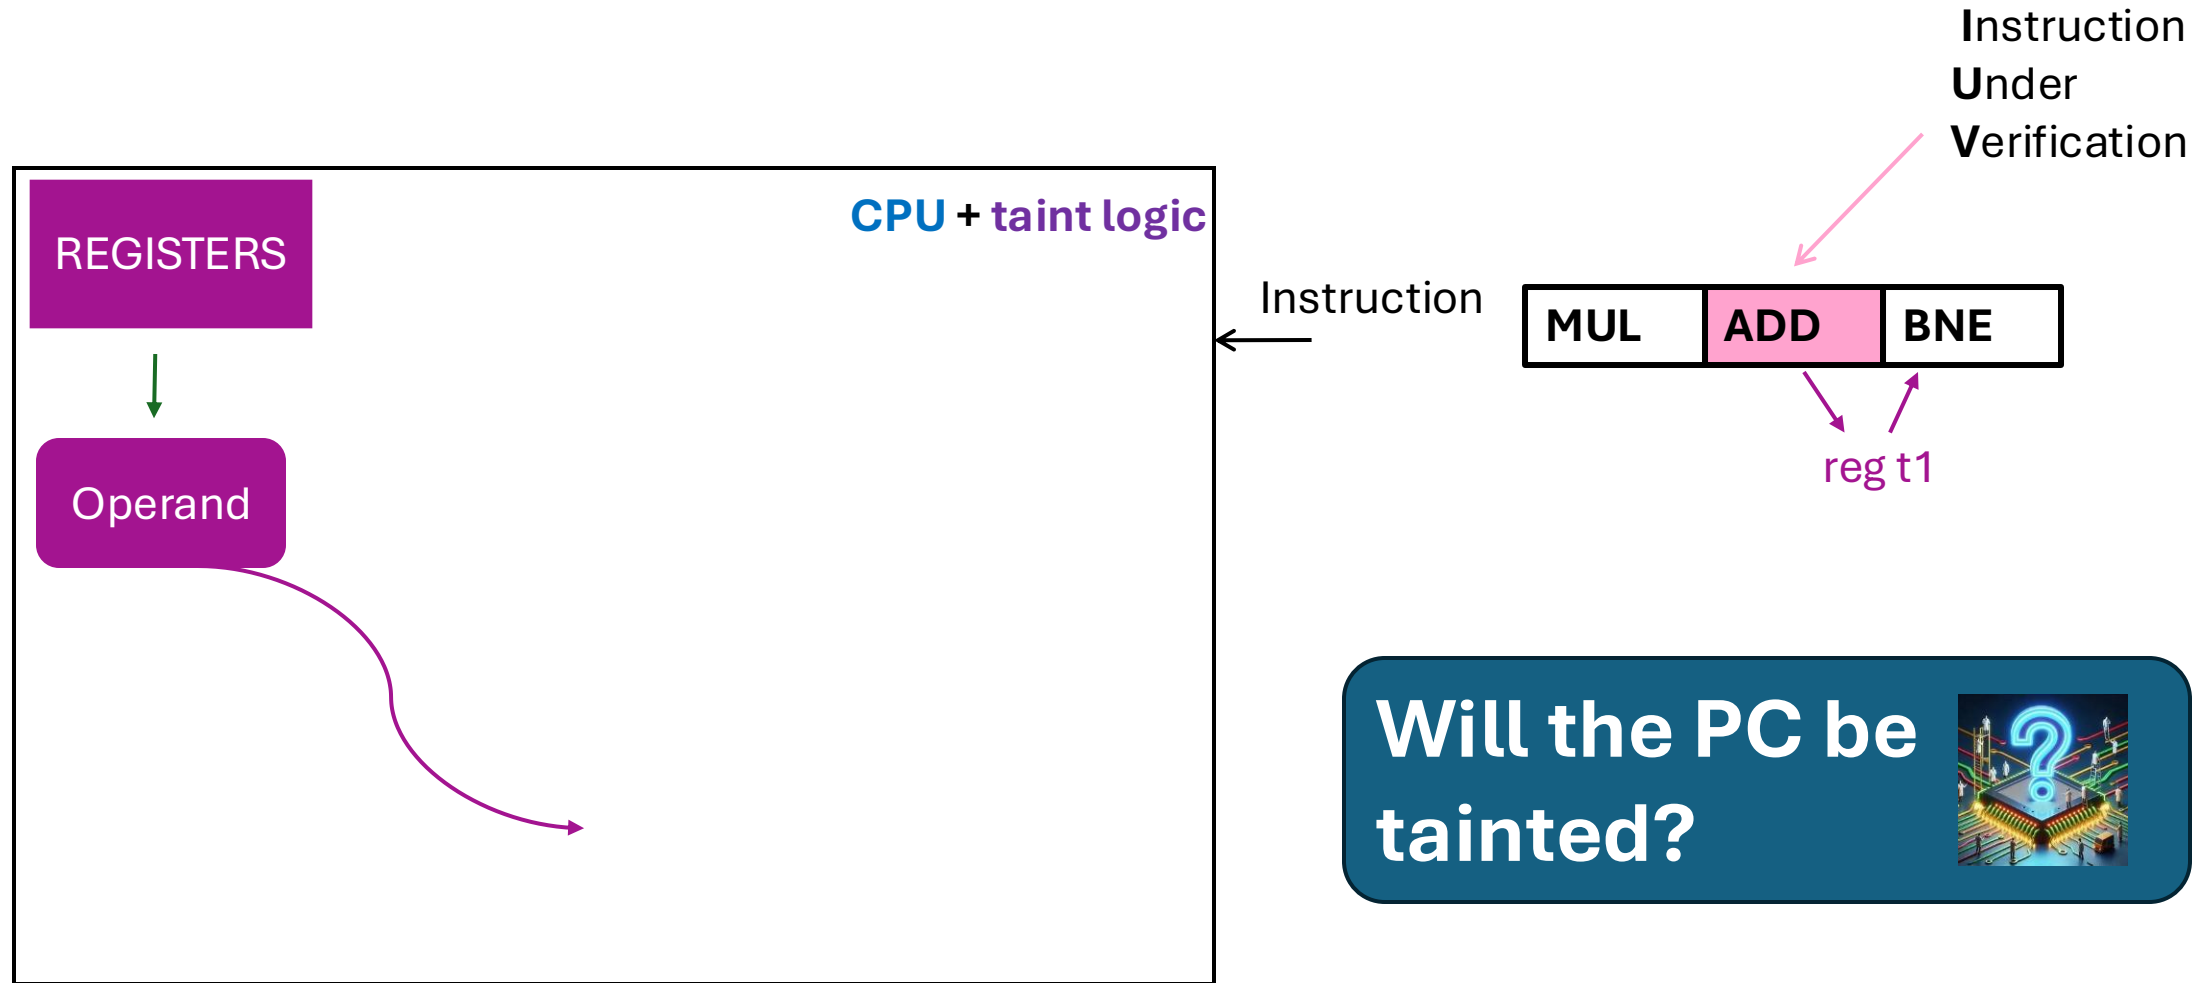
- Controlled **taint injection** per instruction
- Via SystemVerilog Assumptions

Precise Taint Injection

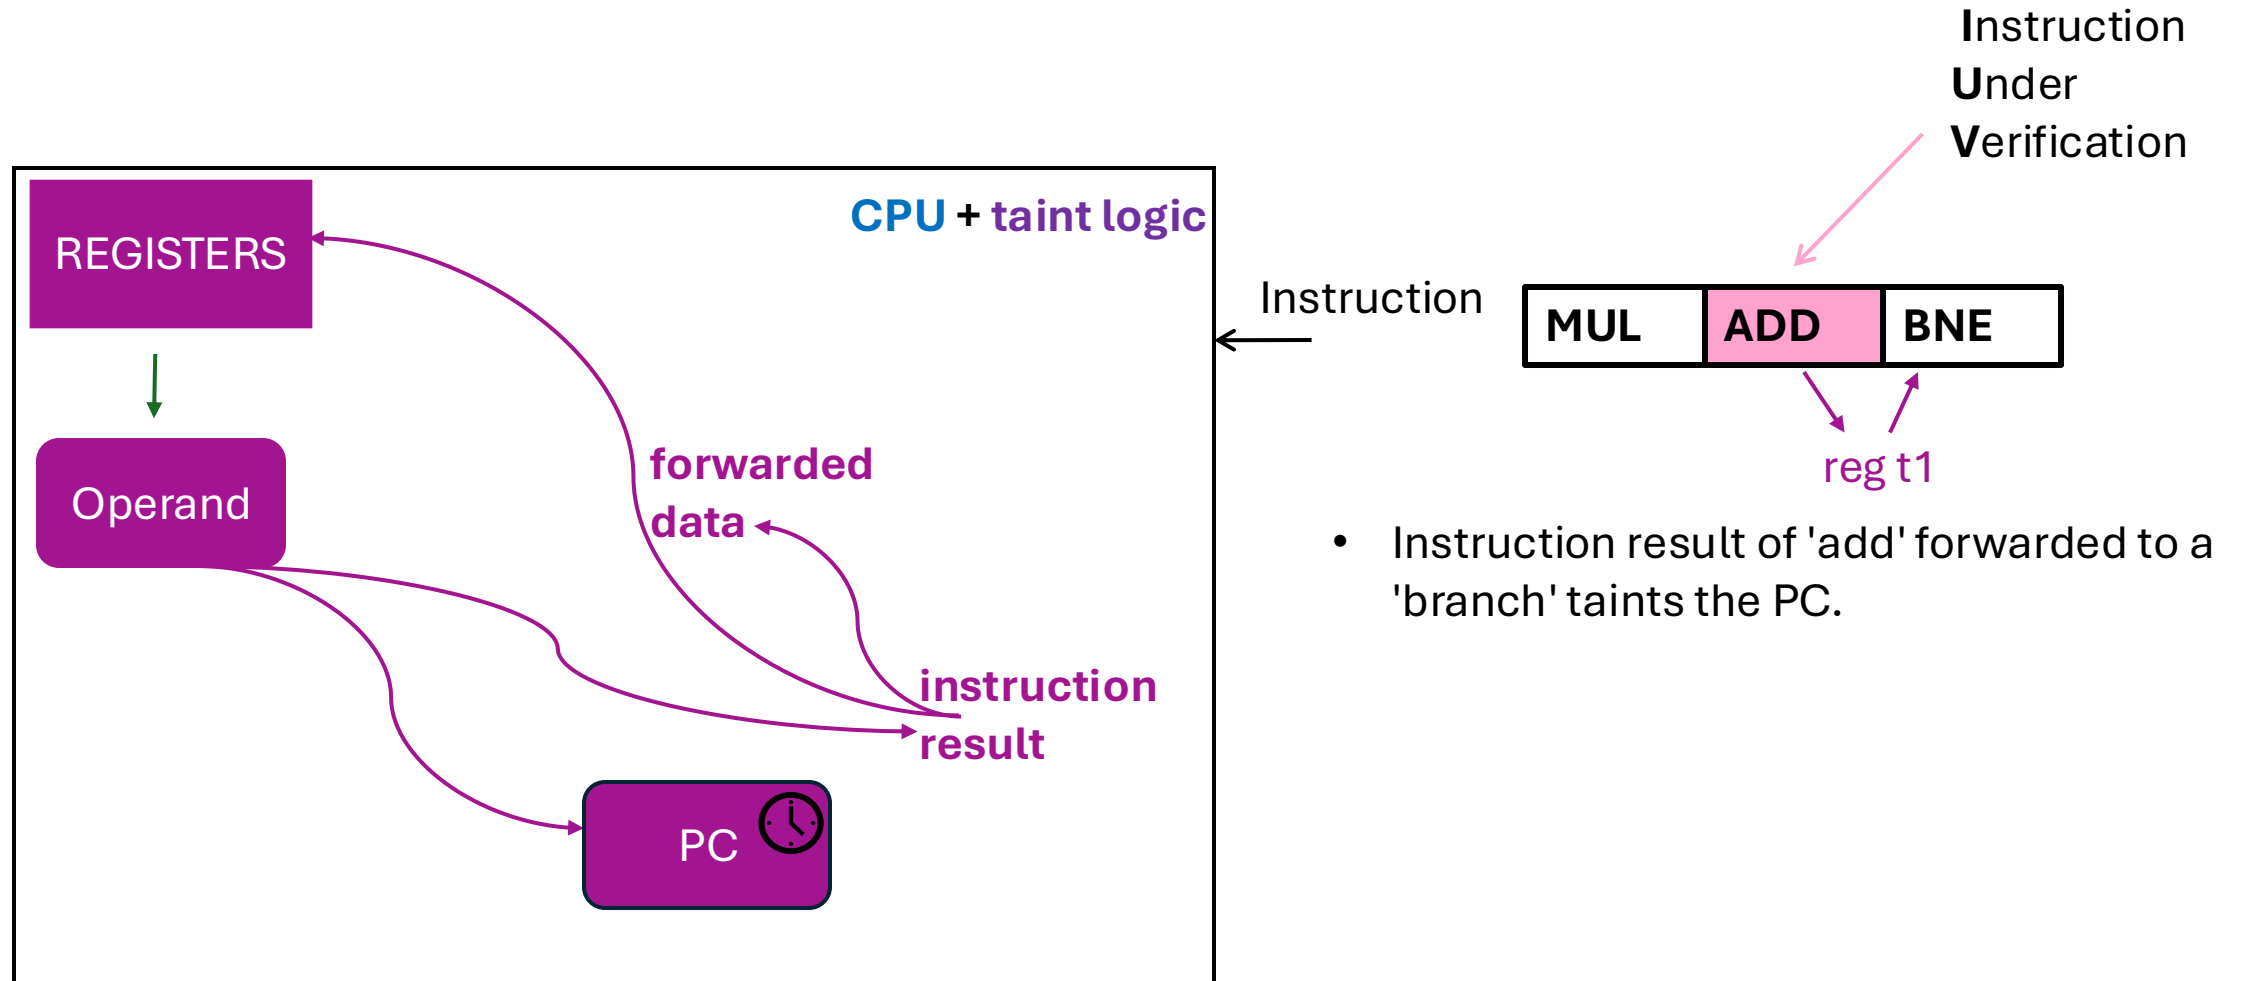


- Controlled **taint injection** per instruction
- Operand **reading conditions** automatically generated via static design analysis (custom Yosys[1] pass)

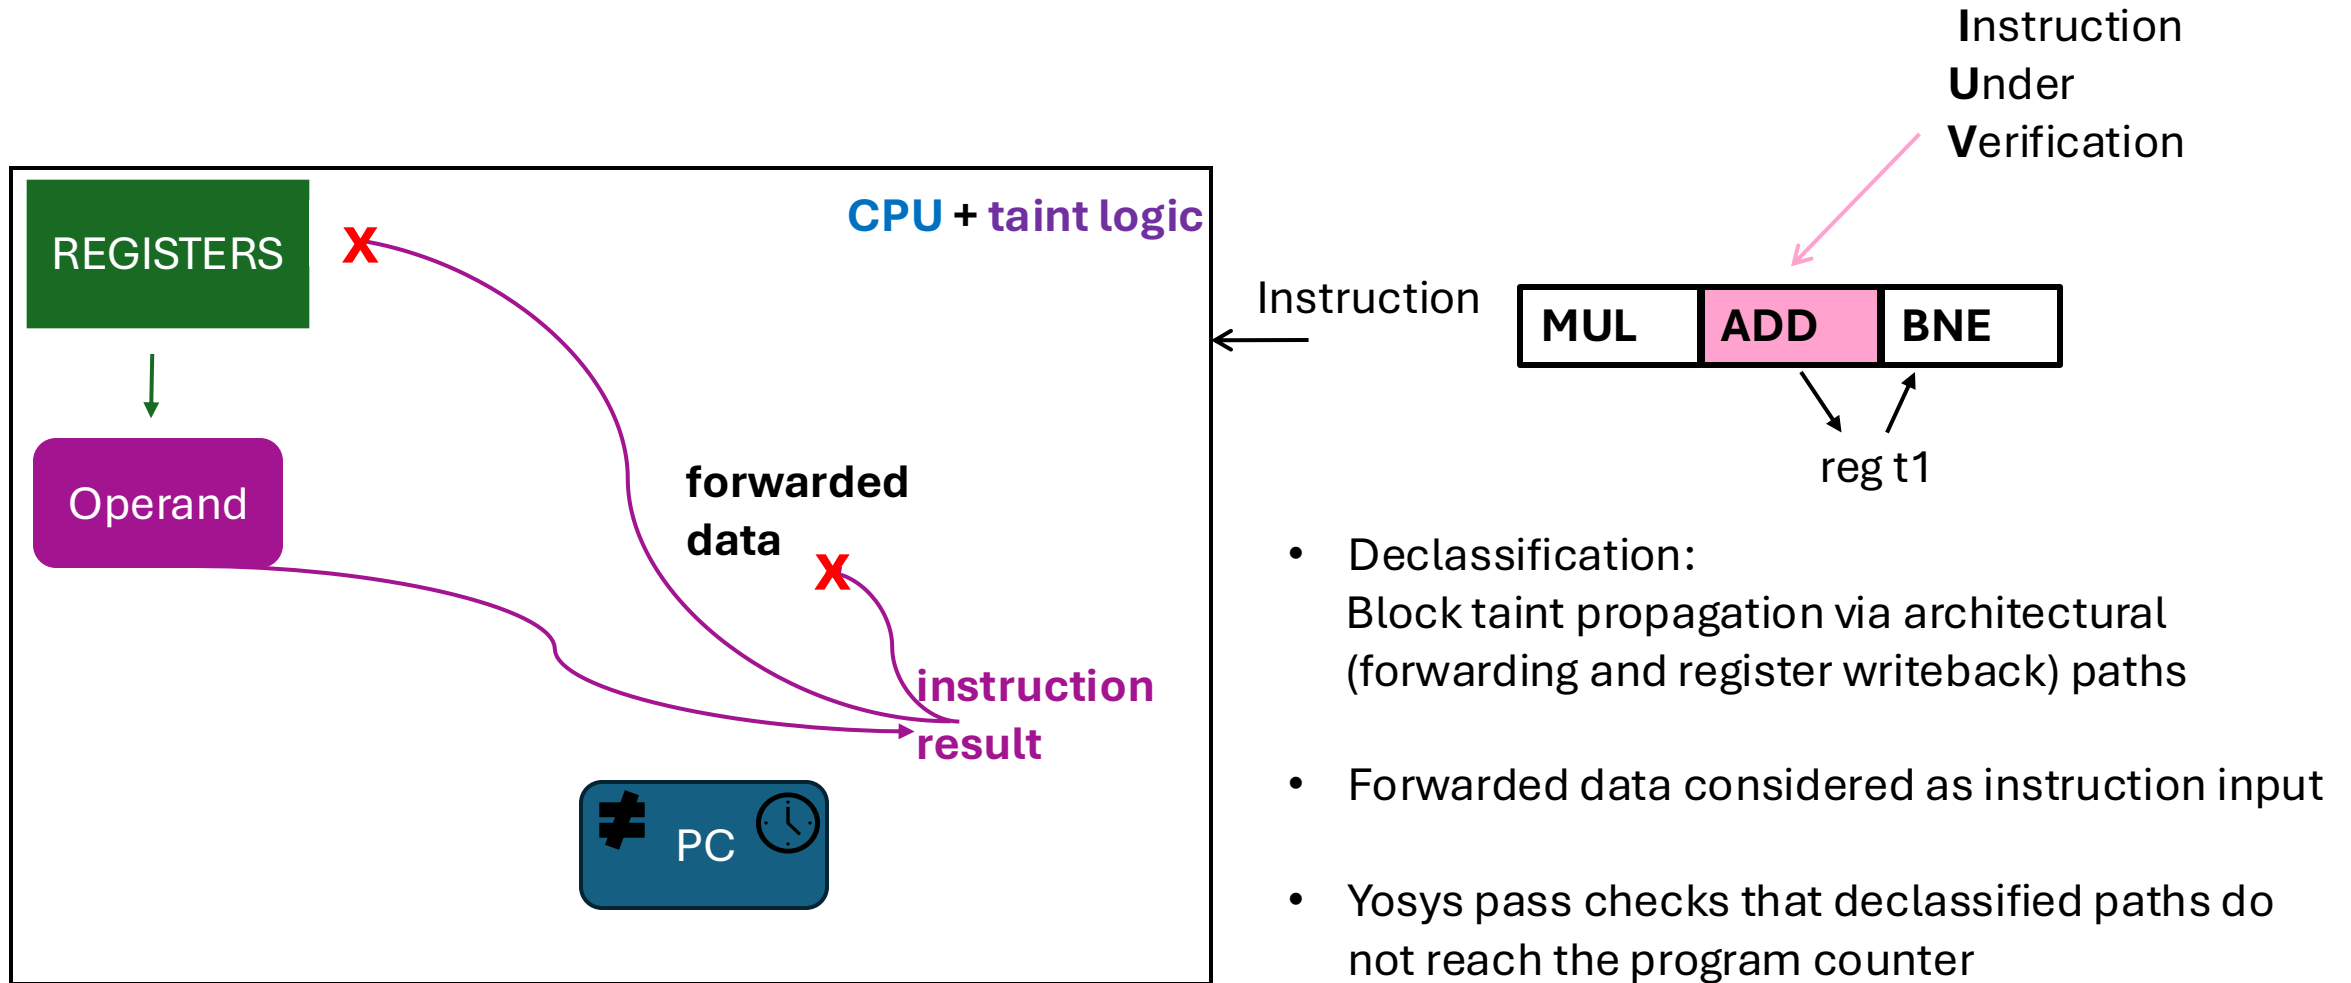
Declassification of Architectural Paths



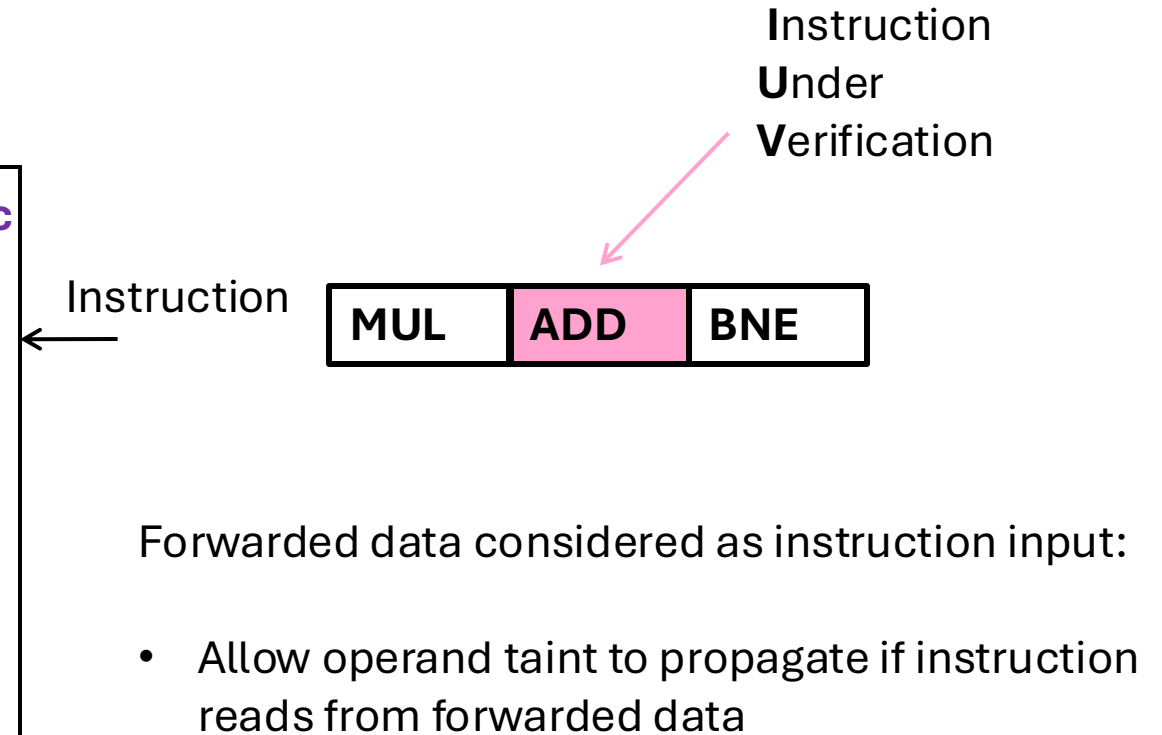
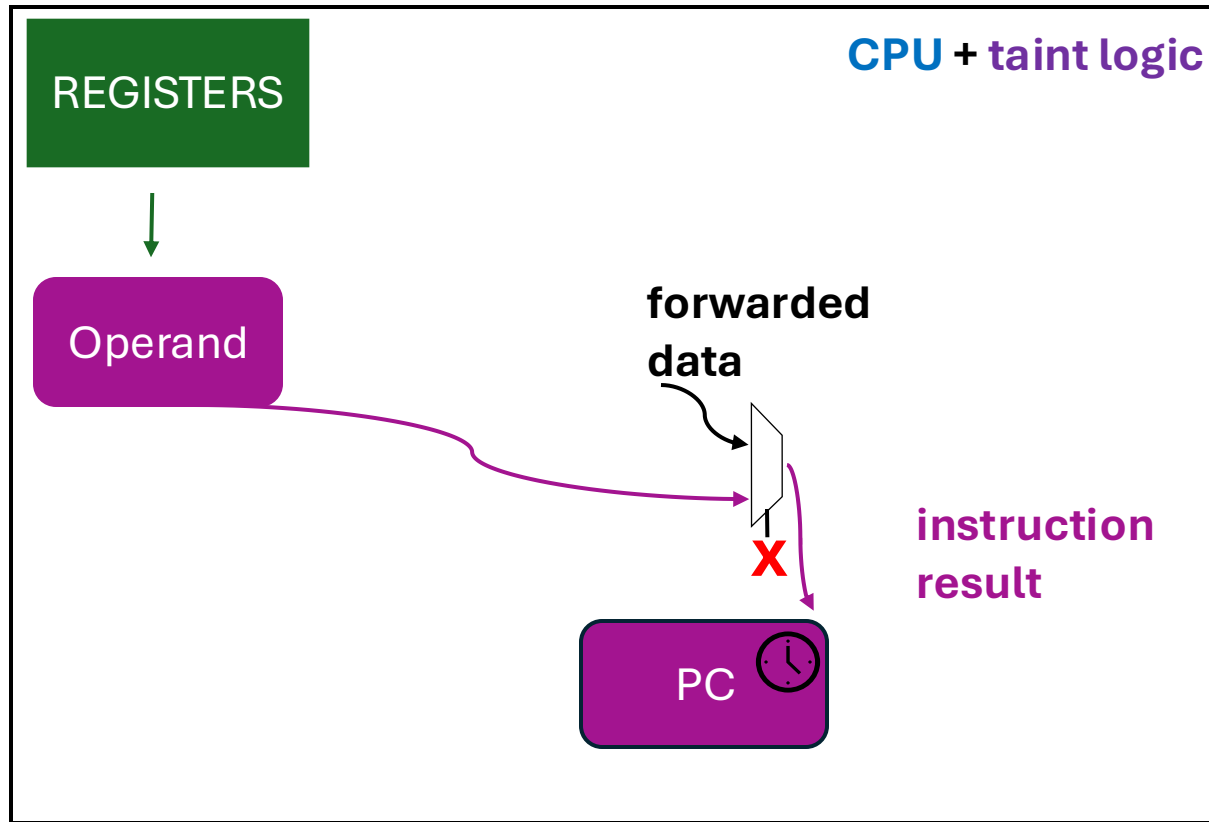
Declassification of Architectural Paths



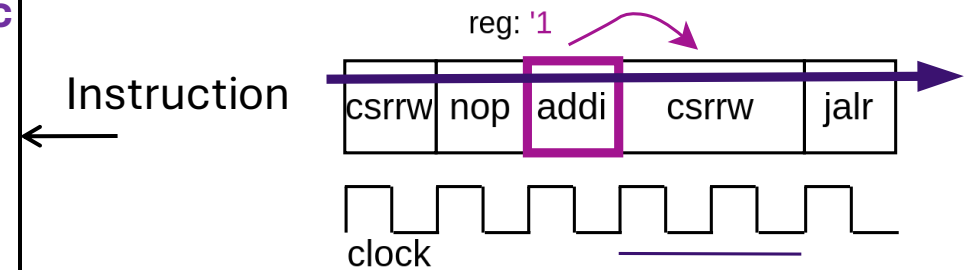
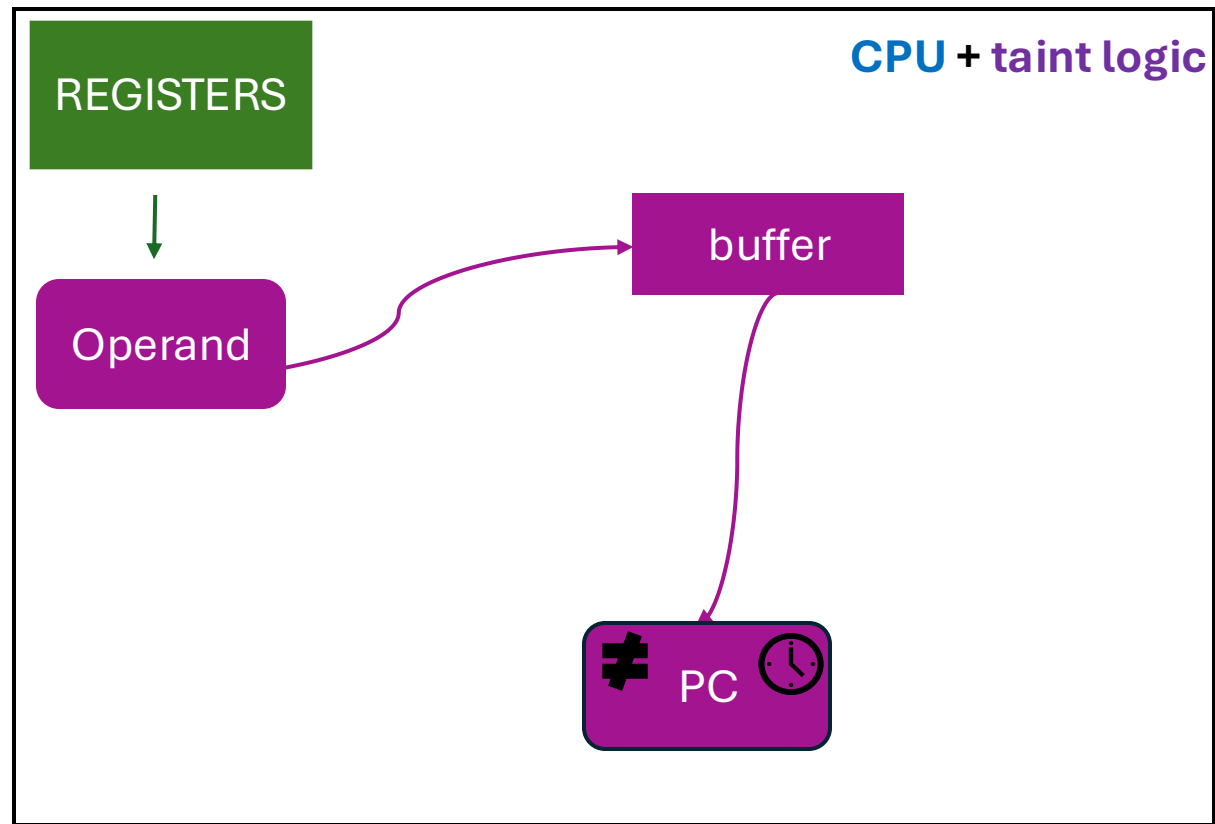
Declassification of Architectural Paths



Tracking Forwarded Data







Declassification of Architectural Paths



No other microarchitectural flows are blocked

Verified RISC-V CPUs





Microcontroller-class, in-order CPUs



State bits	3.2k	2.0k	2.5k	2.3k
Net bits	1.6k	1.4k	4.6k	6.7k

Verified RISC-V CPUs

Microcontroller-class, in-order CPUs







State bits	3.2k	2.0k	2.5k	2.3k
Net bits	1.6k	1.4k	4.6k	6.7k

Cell-	IFT / DFT	IFT / DFT	IFT / DFT	IFT / DFT
∅ time to PROVE	17 h / 8m	16m / 30s	9h / 10m	14.5h / 50 m
∅ time to FAIL	1h / 8m	37s / 15s	2h / 3m	11m / 34m

Model checker: Cadence Jasper Formal Property Verification App


Verified RISC-V CPUs

Microcontroller-class, in-order CPUs




State bits	3.2k	2.0k	2.5k	2.3k
Net bits	1.6k	1.4k	4.6k	6.7k

Cell-	IFT / DFT	IFT / DFT	IFT / DFT	IFT / DFT
∅ time to PROVE	17 h / 8m	16m / 30s	9h / 10m	14.5h / 50 m
∅ time to FAIL	1h / 8m	37s / 15s	2h / 3m	11m / 34m


PROVEN instructions	38	25	27	38 + all crypto instr.
VULNERABLE instructions 	3 (documented)	8	14	3 (known)

New Discovered Security Vulnerabilities





Constant time violation:
CVE-2023-51974


Architectural control flow




Microarchitectural control flow



Two control-flow hijacks:
CVE-2023-51973
CVE-2024-44927




New Discovered Security Vulnerabilities





Constant time violation:
CVE-2023-51974


Architectural control flow



Microarchitectural control flow



Two control-flow hijacks:
CVE-2023-51973
CVE-2024-44927



Constant time violation + data leakage:
CVE-2024-28365



Control-flow hijack



Conclusion

- Introduced and formalized a generalized CPU security property



μCFI - Microarchitectural Control-flow Integrity

Conclusion

- Introduced and formalized a generalized CPU security property



μCFI - Microarchitectural Control-flow Integrity

- Automated verification method & implementation
- 4 open-source RISC-V CPUs verified
- Discovered 5 new vulnerabilities - 4 CVEs



Conclusion

- Introduced and formalized a generalized CPU security property



μCFI - Microarchitectural Control-flow Integrity

- Automated verification method & implementation
- 4 open-source RISC-V CPUs verified
- Discovered 5 new vulnerabilities - 4 CVEs



Video:

<https://www.youtube.com/watch?v=Kxp-5kNMt40&t>

Information:



<https://comsec.ethz.ch/>

Code:



[comsec-group/mucfi](https://github.com/comsec-group/mucfi)

Contact:



<https://www.linkedin.com/in/katharina-ceesay-seitz-ba521087/>



@K_Ceesay-Seitz, @FlavienSolt



kceesay@ethz.ch, flavien.solt@eecs.berkeley.edu

Thank you! Questions?


BACKUP

Taint Start Condition

Update Condition Yosys Pass

Read-from Condition = the condition in which a signal is updated with a chosen signal's value.

Taint start



```
yosys update_condition -read-from-signals "cpuregs" -signal_name "cpuregs_rs1"
```

CPU code (PicoRV32):

```
1349     always @* begin
1350         decoded_rs = 'bx;
1351         if (ENABLE_REGS_DUALPORT) begin
1352             `ifndef RISC_V_FORMAL_BLACKBOX_REGS
1353                 cpuregs_rs1 = decoded_rs1 ? cpuregs[decoded_rs1] : 0;
1354                 cpuregs_rs2 = decoded_rs2 ? cpuregs[decoded_rs2] : 0;
1355             `else ...
```

Generated Read-from Condition:

```
1  bit gen_regrd_rs1;
2  assign gen_regrd_rs1 =
3  ((| decoded_rs1));
4  ;
```



Taint Stop Condition

Update Condition Yosys Pass

Update Condition (UC) = the condition in which a signal is updated with another value than its own previous value.

Taint stop

Operand's register
data read signal



Condition in which
the read data MAY be new

For example:

- enable condition of a flip flop
- '1' (True) for continuous assignments

Precise Taint Injection Conditions



Sample Instruction Word (IW) in formal setup

start



IW == IUV and taint start condition

stop



Taint stop condition

Potentially
taint
multiple
times per
instruction




Simple & precise
counter examples

Instruction
Under
Verification




start



taint

stop

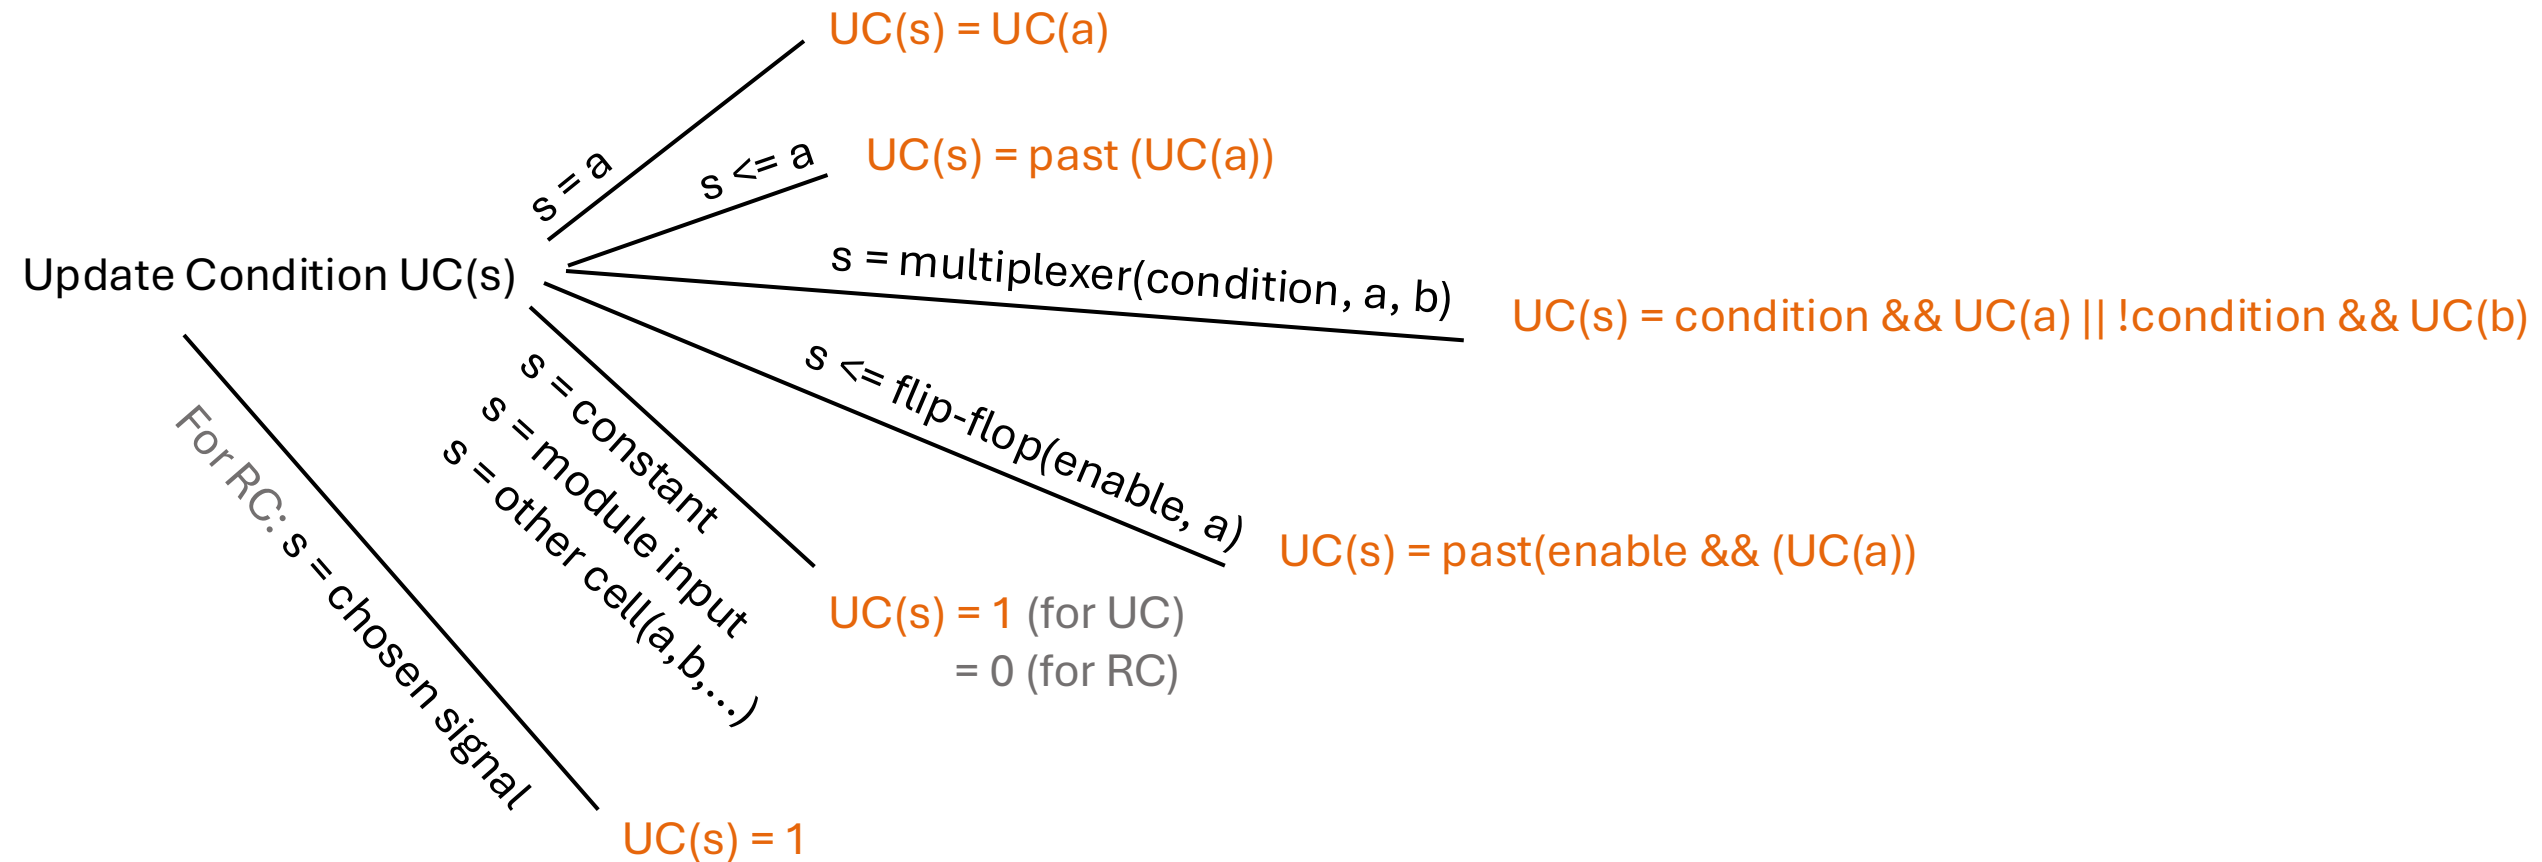


- Controlled **taint injection** per instruction

Update Condition (UC) / Read-from Condition (RC) Yosys Pass


s ... signal

a,b ... other internal signals
'past' = custom attribute



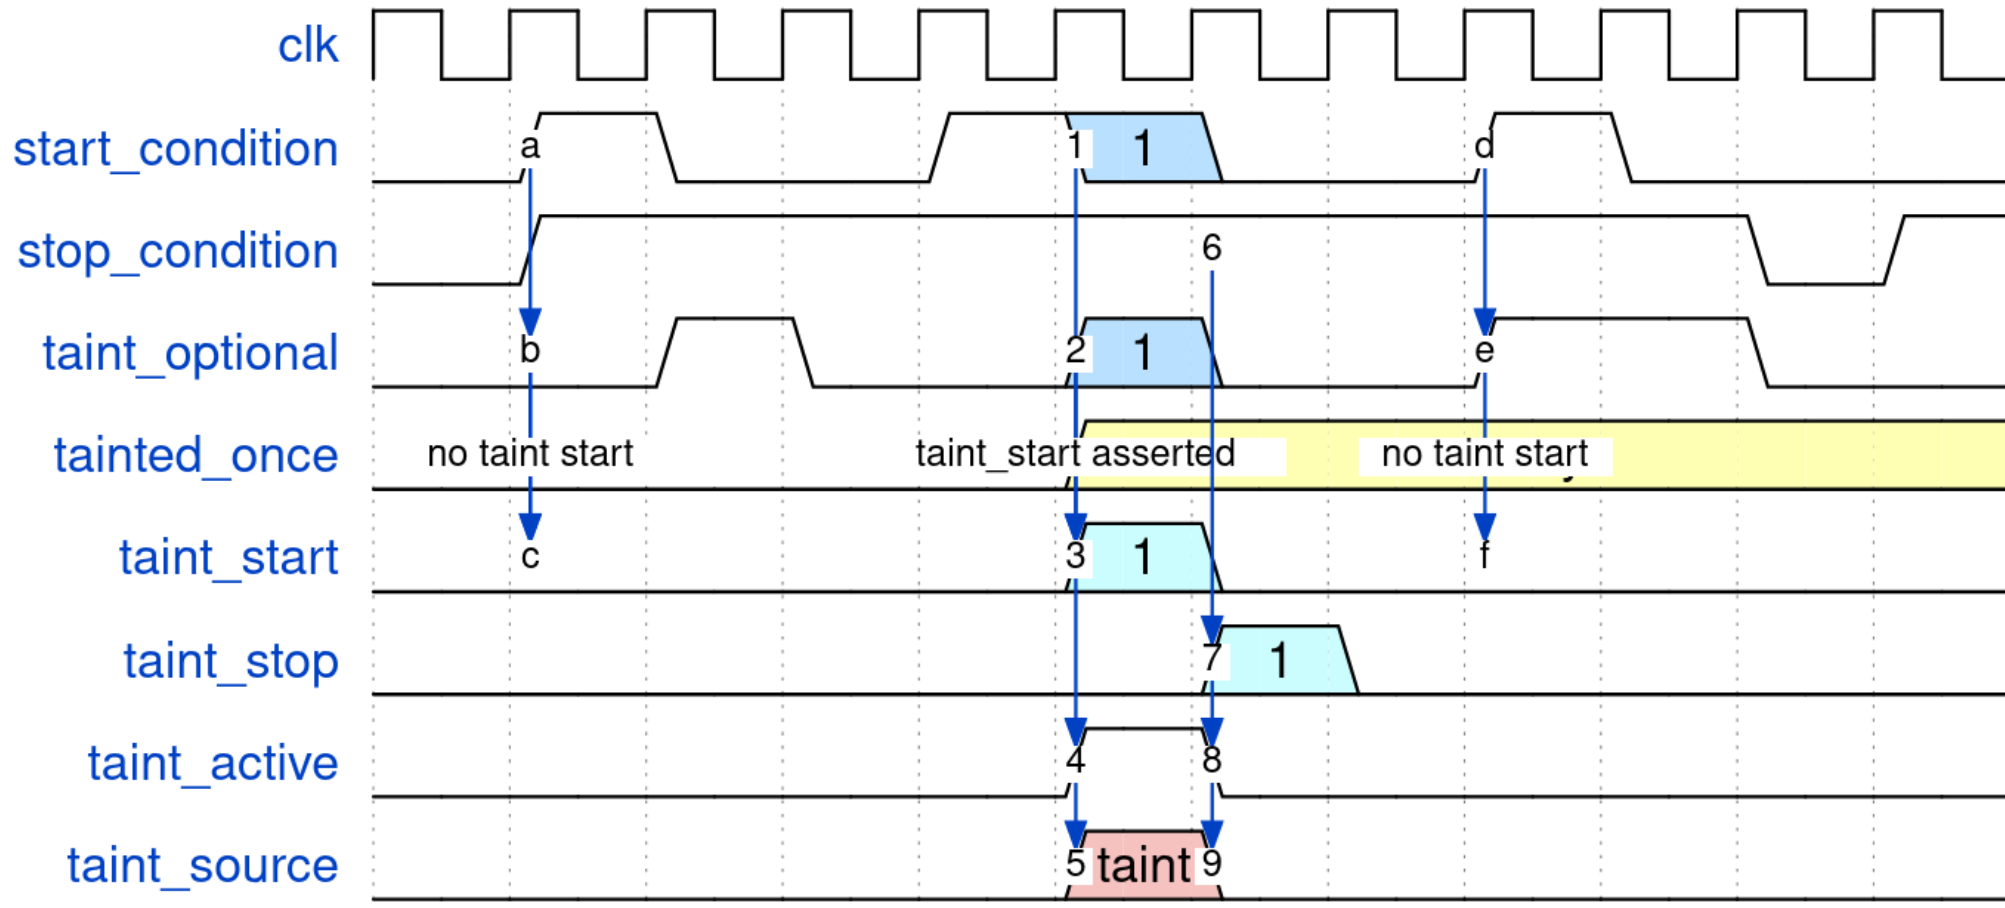
Find Forwarding Multiplexer Yosys Pass

- Automatically identifies forwarding multiplexers
- Checks declassification precondition: all outgoing paths of declassified signals reach another declassified signal or data source without passing PC




1. Traverse outgoing paths of **forwarded data output** and check declassification precondition
2. If a mux uses **forwarded data output**, back-traverse multiplexers' other input's driving logic.
3. Is it directly assigned with **operand's register data read** signal?
 - No: continue at mux output
 - Yes: Forwarding mux found **X** --> return mux select signal

Taint Injection Assumptions



Introducing μCFI - Microarchitectural Control-flow Integrity

Microarchitectural control flow (μCF)



μCFI only allows explicitly ISA specified data dependencies of the μCF

