Google Cloud

NetInfra

P4-BAR

Scaling P4-Based Automated Reasoning (Performance and Coverage)

P4 Workshop, Oct. 3, 2024

Ali Kheradmand (Google), Meghana Sistla (UT Austin*)

* Work done while at Google

Highlights of Our Journey Developing P4-Symbolic

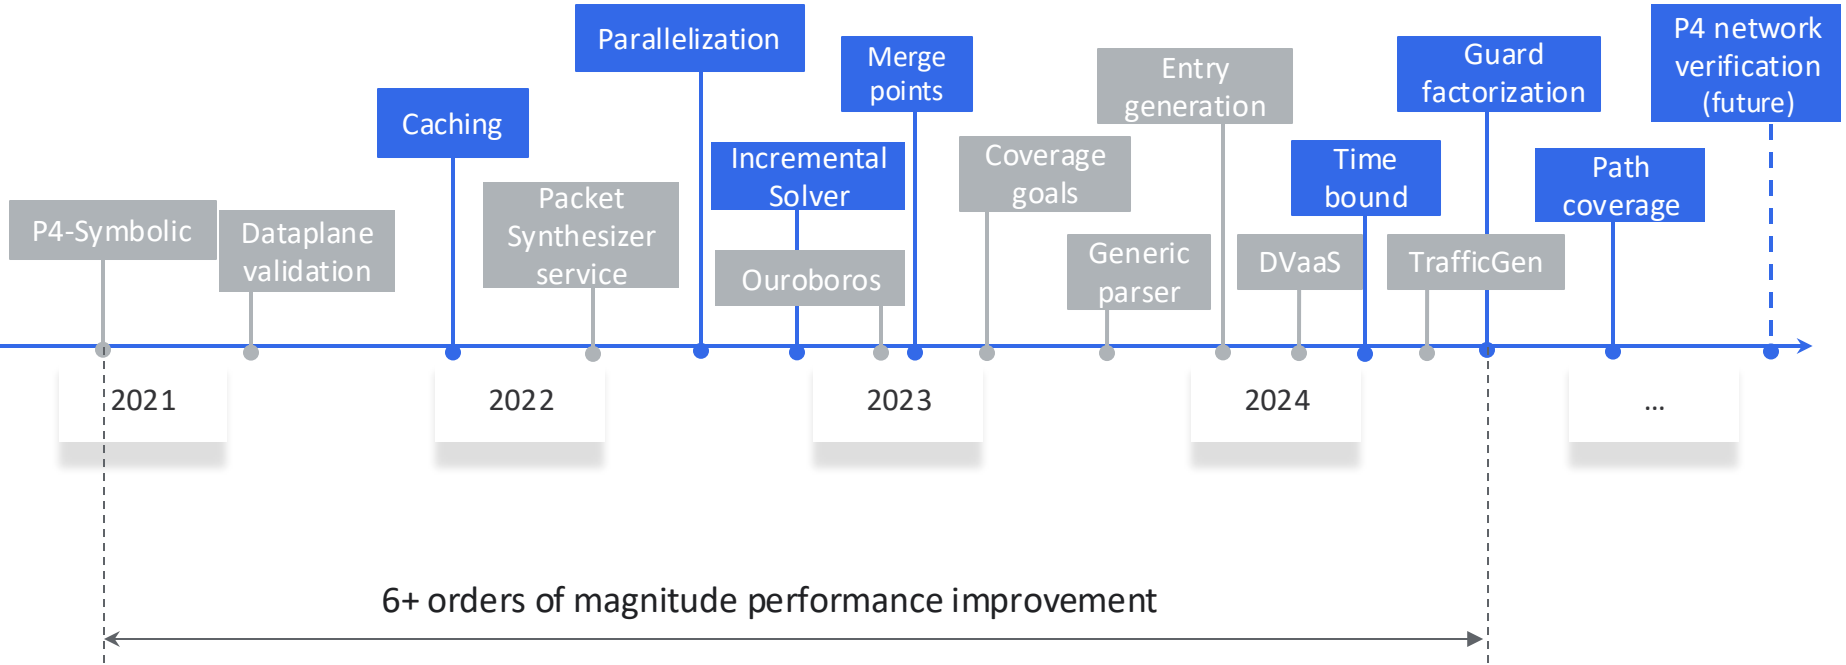

Overall goal: Ensure network works as intended

Subgoal 1: Ensure controller produces correct table entries (according to intents)

Subgoal 2: Ensure switch works as expected (according to table entries)

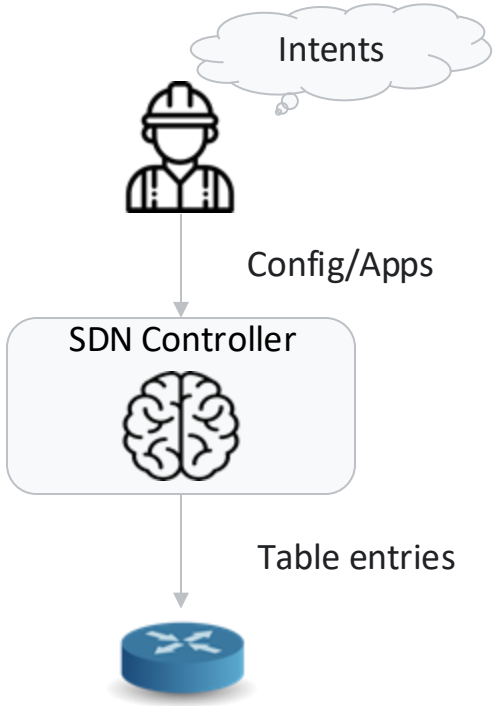

Focus: Ensure switch stack works as expected

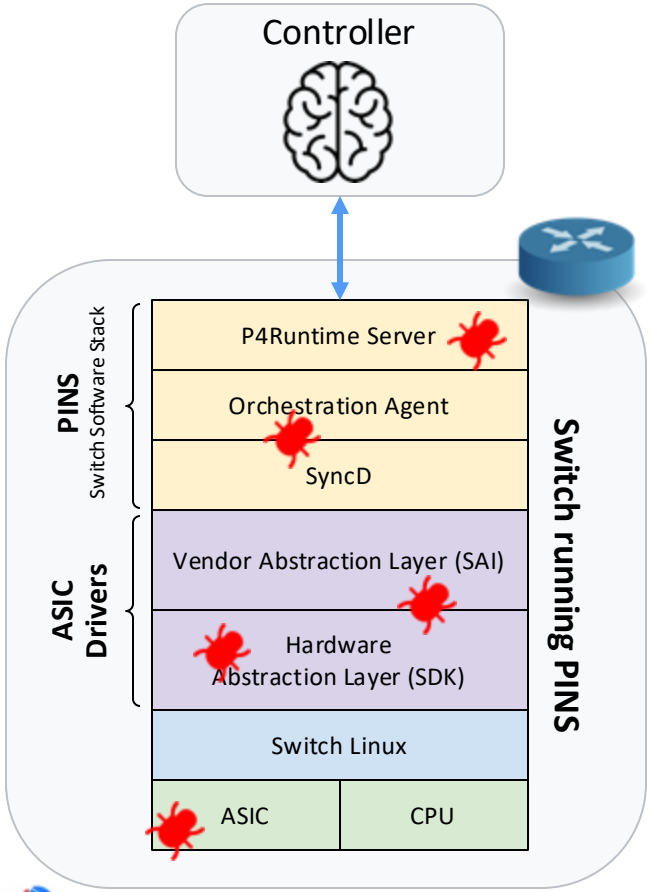

Traditional: manually write tests

- Exponentially large space to cover
  - Labor intensive
- Hard to evolve

Our way: **Automatically** derive tests from a **formal specification** of how the switch should work

- Comprehensive coverage
- Effortless evolution

Need:

1. Specification language
2. Test generation tools

P4-Based Automated Reasoning (P4-BAR)

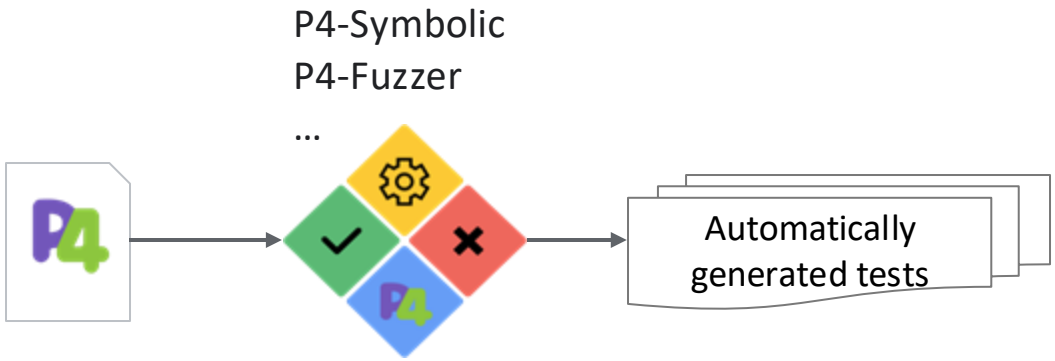

300+ (unique) bugs found so far

(and many more bugs prevented)

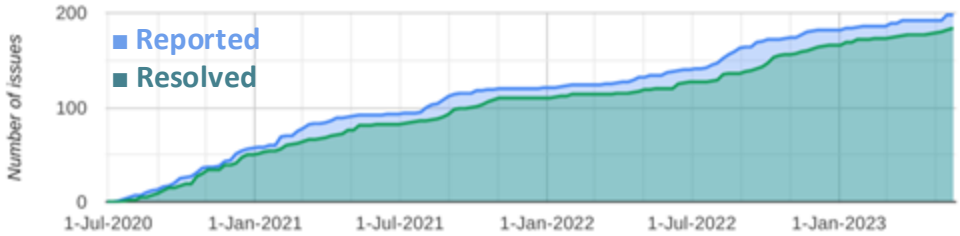

SwitchV: Automated SDN Switch Validation with P4 Models (SIGCOMM'22)

Kinan Dak Albab, Jonathan Dilorenzo, Stefan Heule, Ali Kheradmand, Steffen Smolka, Konstantin Weitz, Muhammad Tirmazi, Jiaqi Gao, Minlan Yu

Dataplane Validation

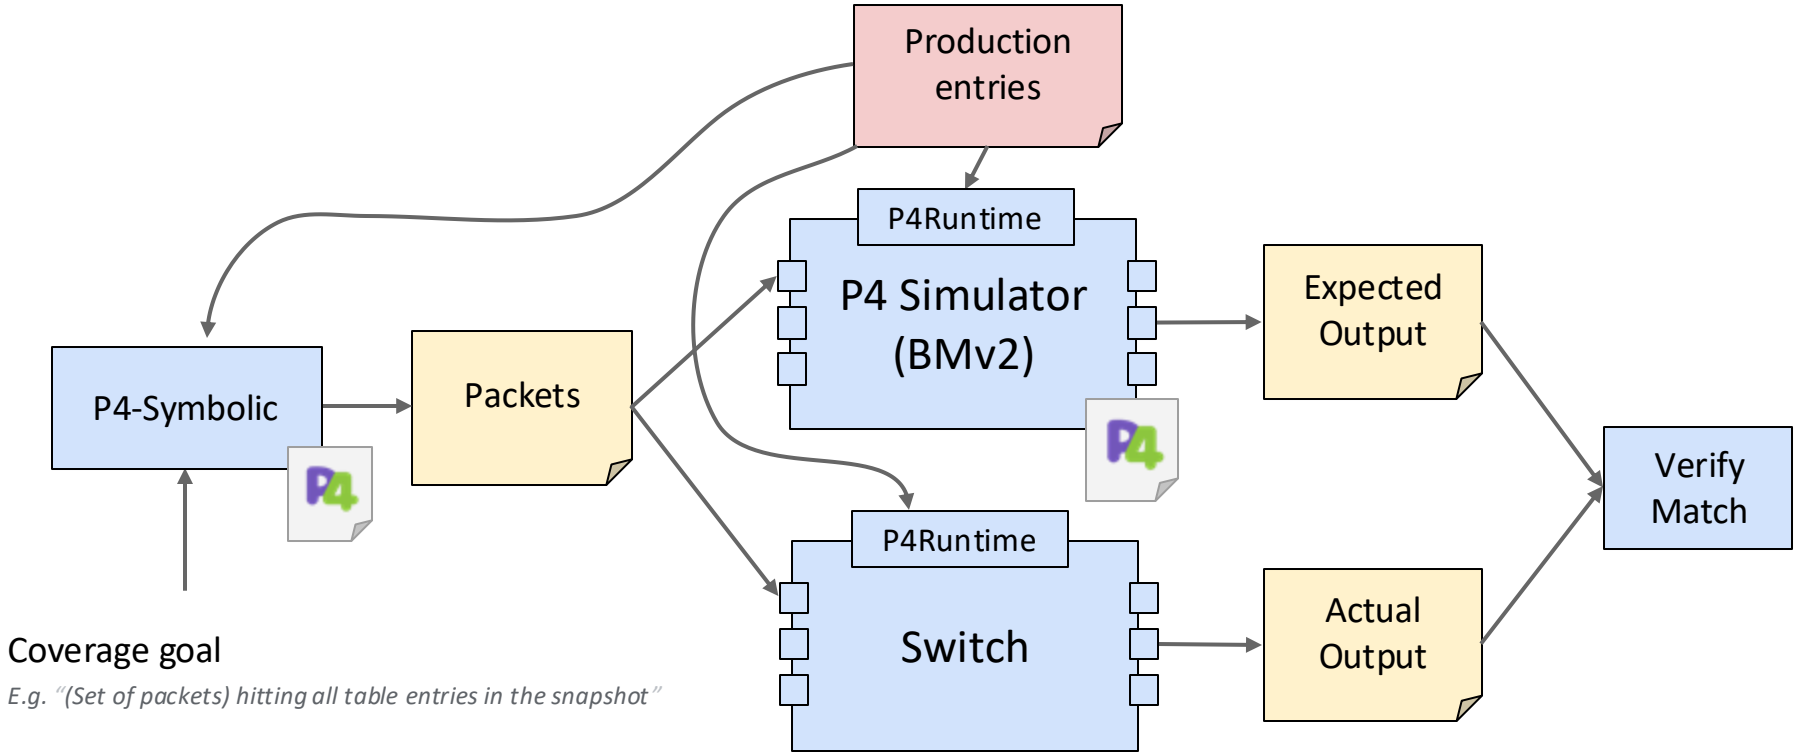

Symbolic Execution

"Give me the input that triggers "do_B()."

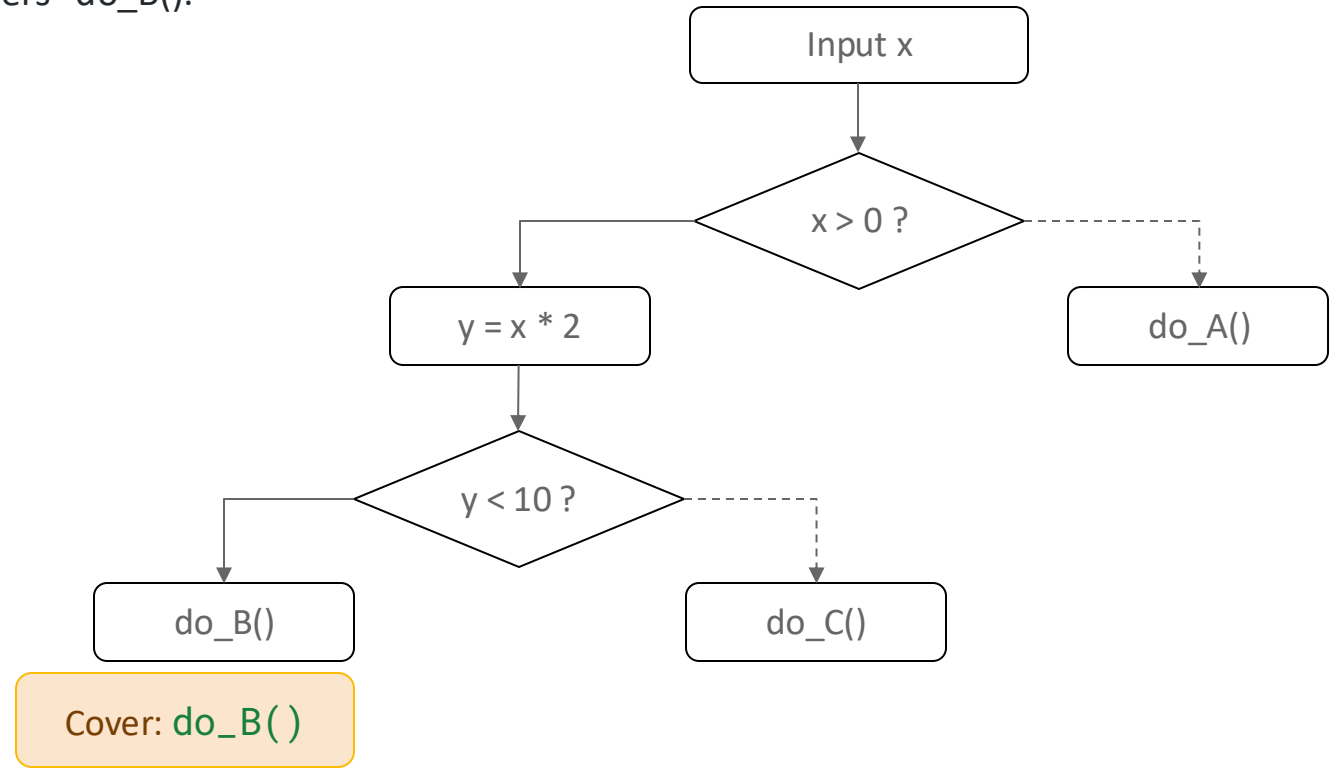

Symbolic Execution

"Give me the input that triggers "do_B()."

⇒ Solve("do_B() is reached")

⇒ Solve((x > 0) ∧ (y = x * 2) ∧ (y < 10))

⇒ A solution: {**x = 1**, y = 2}
(with Z3 solver)

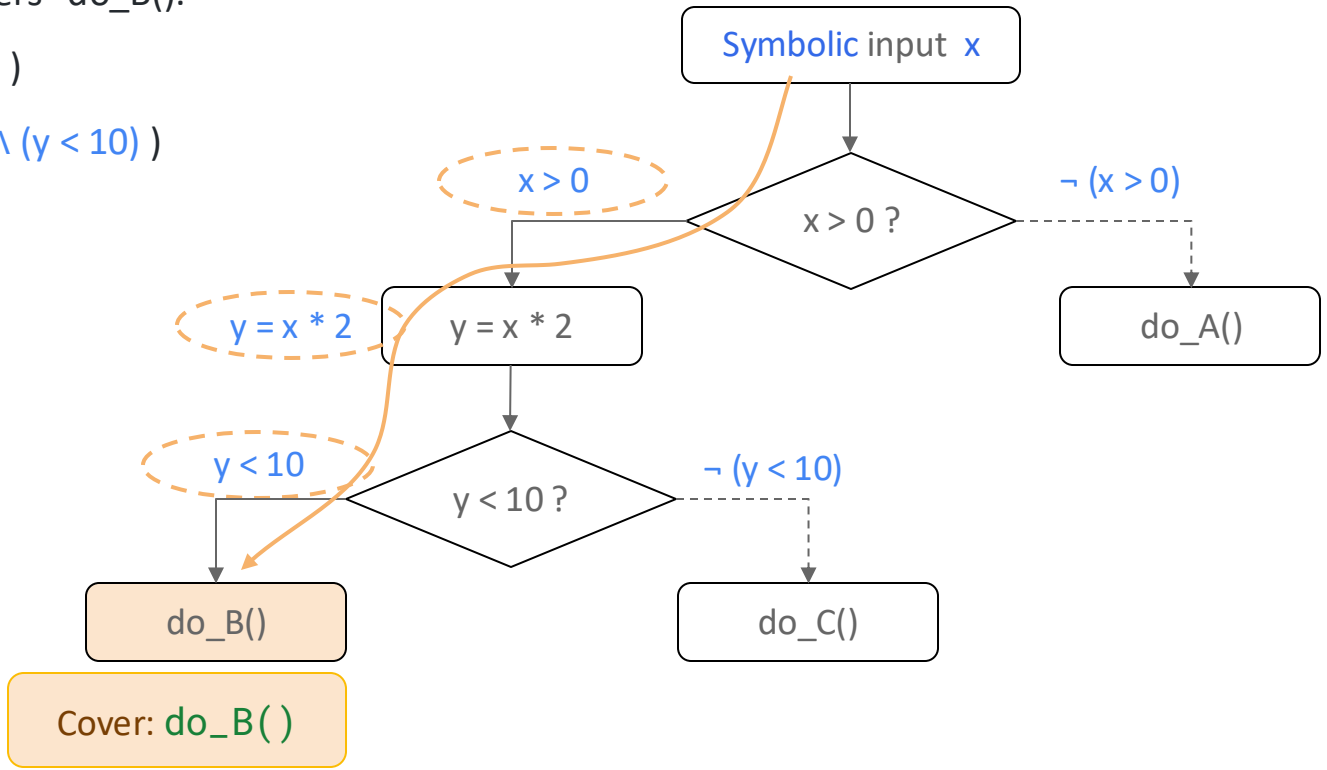

Symbolic Execution in P4

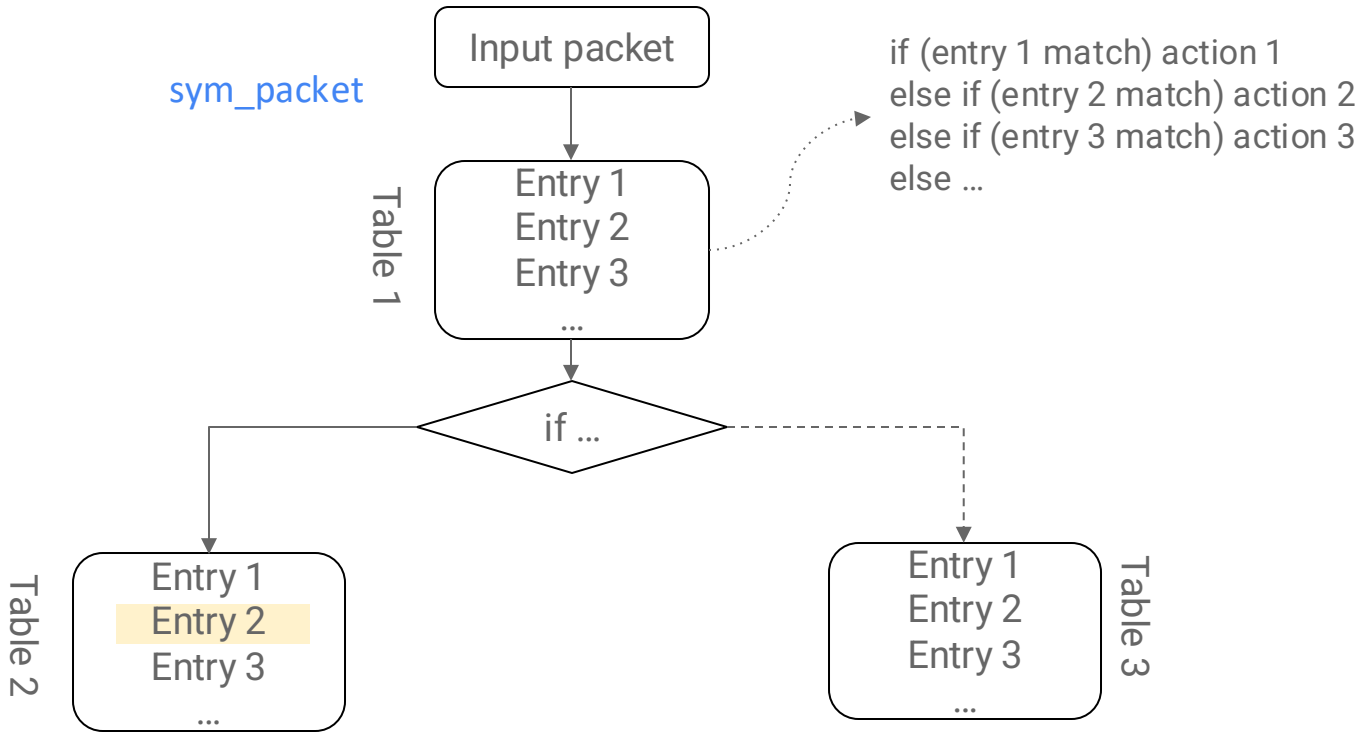

P4-Symbolic

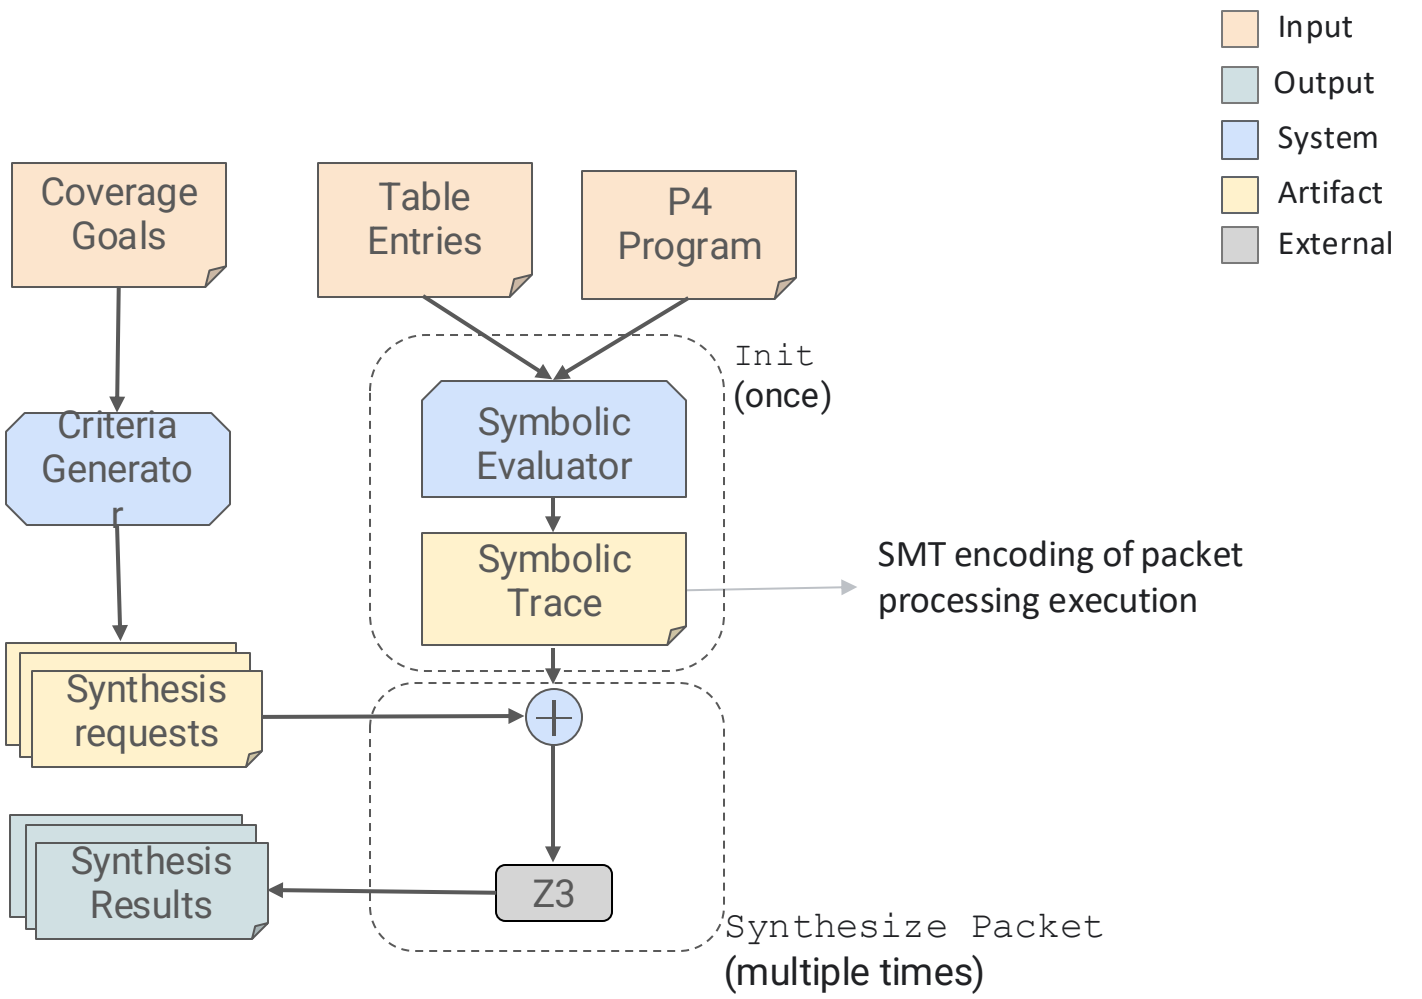
Input
Output
System
Artifact
External
e.g.
“Hitting all table entries”
Coverage Goals
Table Entries
P4 Program
Init
(once)
Criteria Generato r
Symbolic Evaluator
Symbolic Trace
SMT encoding of packet processing execution
e.g.
Request 1: packet must hit 12th entry in table acl_ingress
Request 2: …
…
Synthesis requests
e.g.
Result 1: packet = …
Result 2: unsatisfiable
…
Synthesis Results
Z3
Synthesize Packet
(multiple times)

P4-Symbolic in use

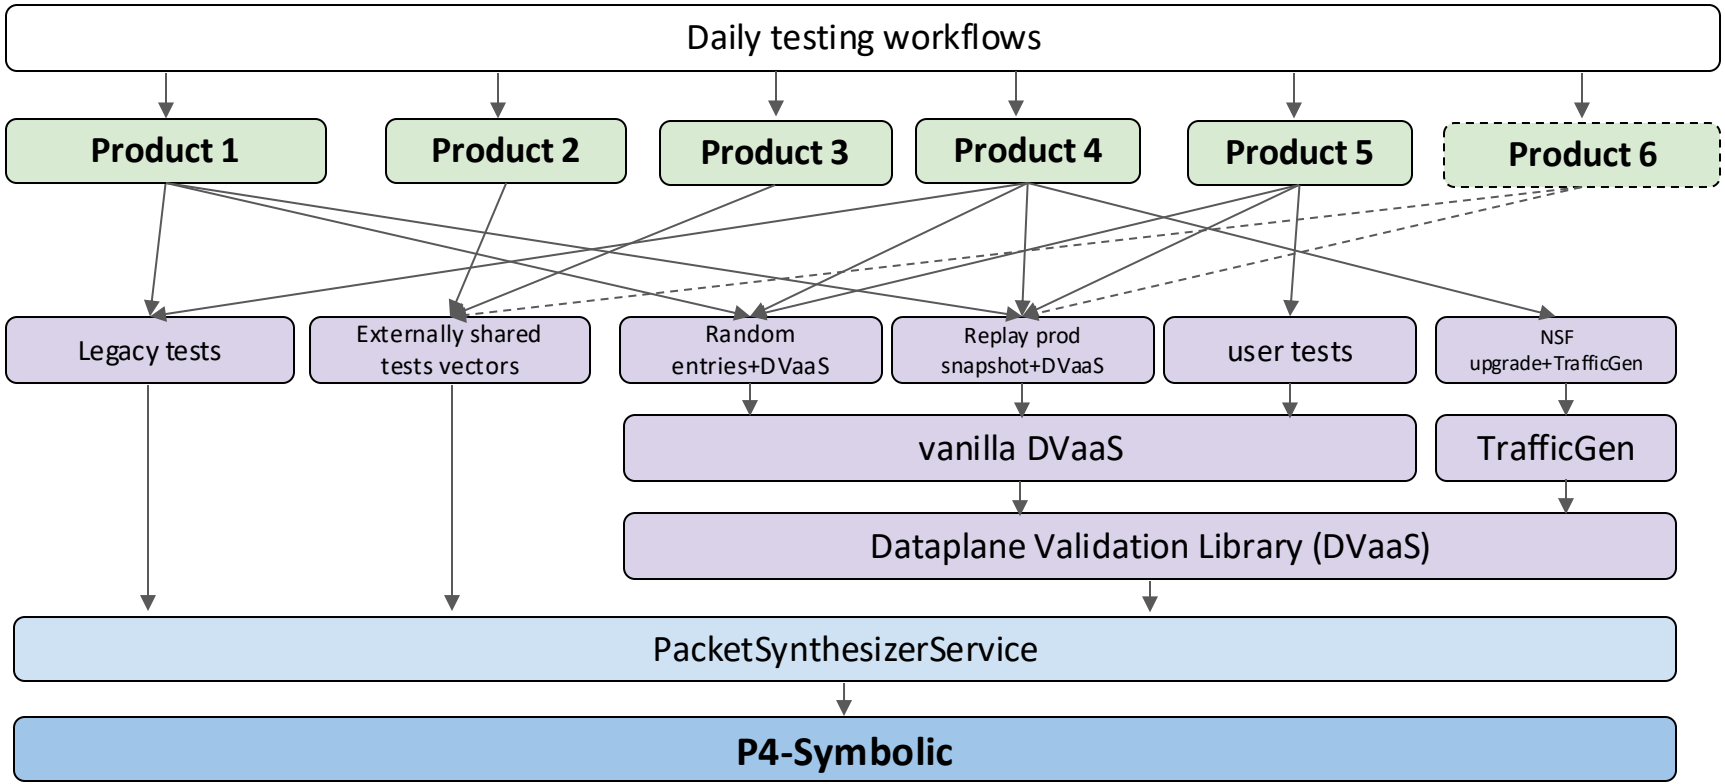

Problem

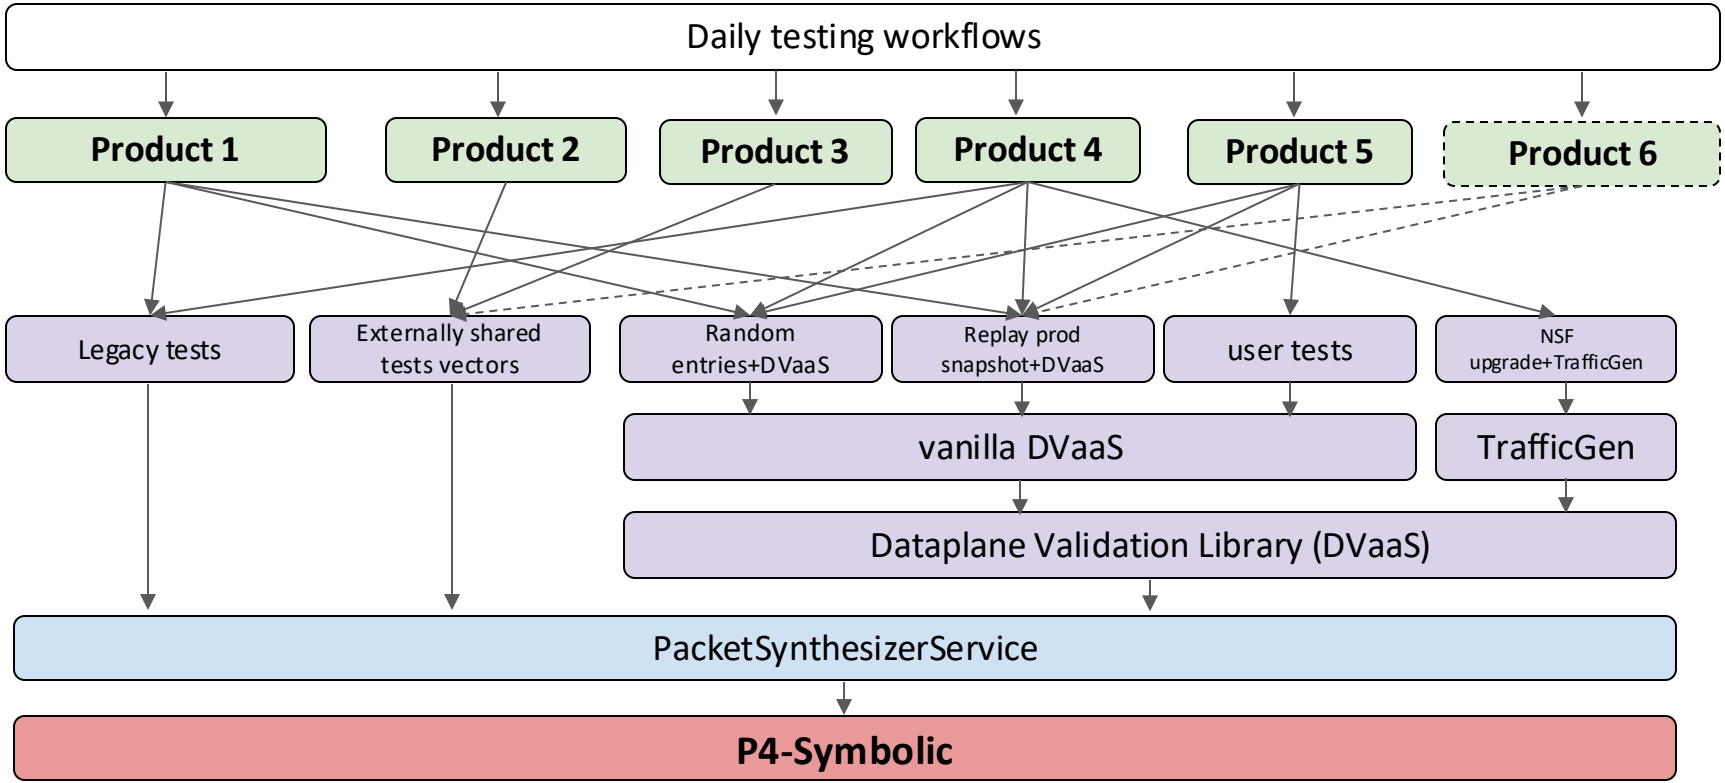

NP-hard problem => computationally expensive => bottleneck

Outline

- Background and Context
- P4-Symbolic
- Performance Improvements
- Coverage Improvements
- Future

Undesirable “Solution”

Reduce coverage

- Smaller coverage goals
  - E.g. Ignore expensive tables, entries
- Time bound coverage:
  - Stop execution after a certain “time limit” (even if coverage goal not achieved)

Undesirable, but at times necessary as a last resort

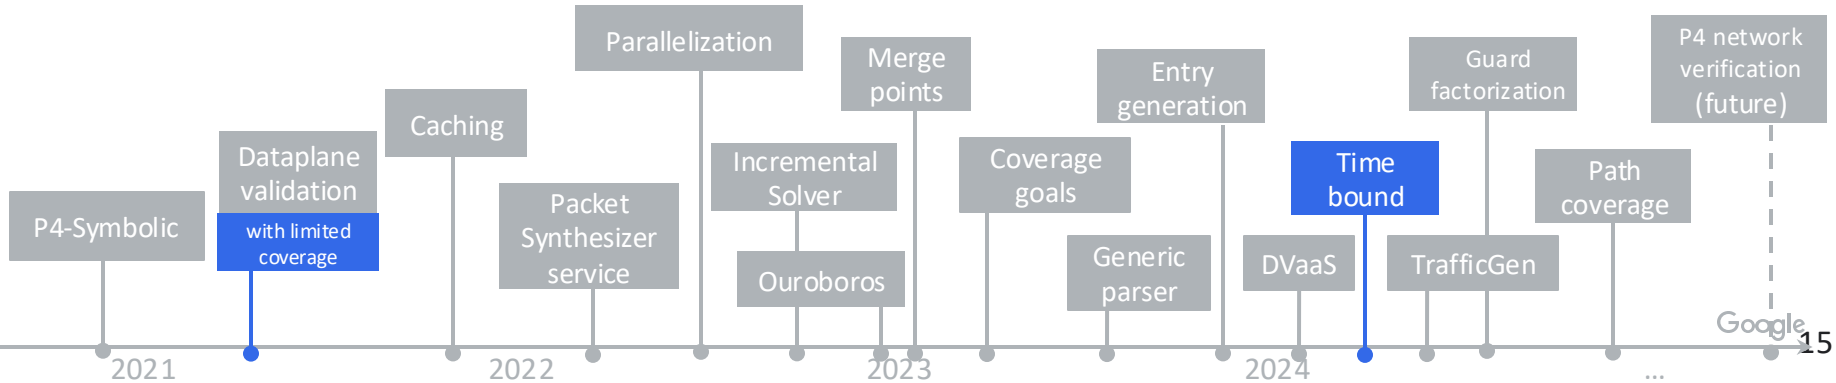

1. Offline packet synthesis (caching)

No need to regenerate packets unless inputs (P4 model, entries, goals) change

Do not allow code merge until cache is populated

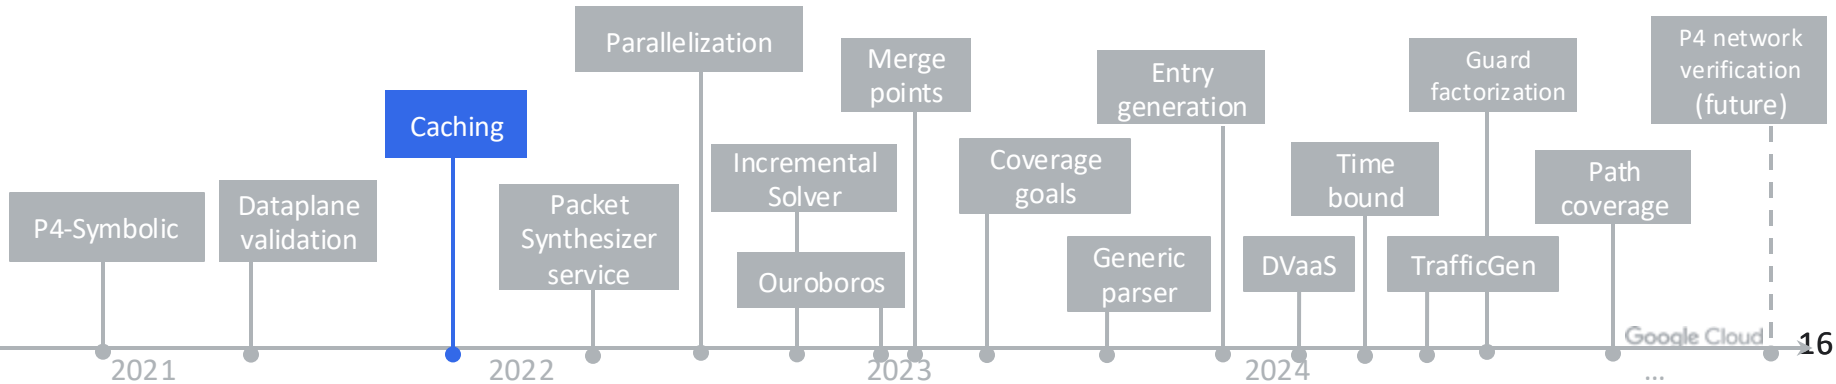

1. Offline packet synthesis (caching)

Caveats

- Time to populate cache
- Frequent P4 model updates
  - Headache with concurrent development
- Ineffective in tests that frequently update entries,

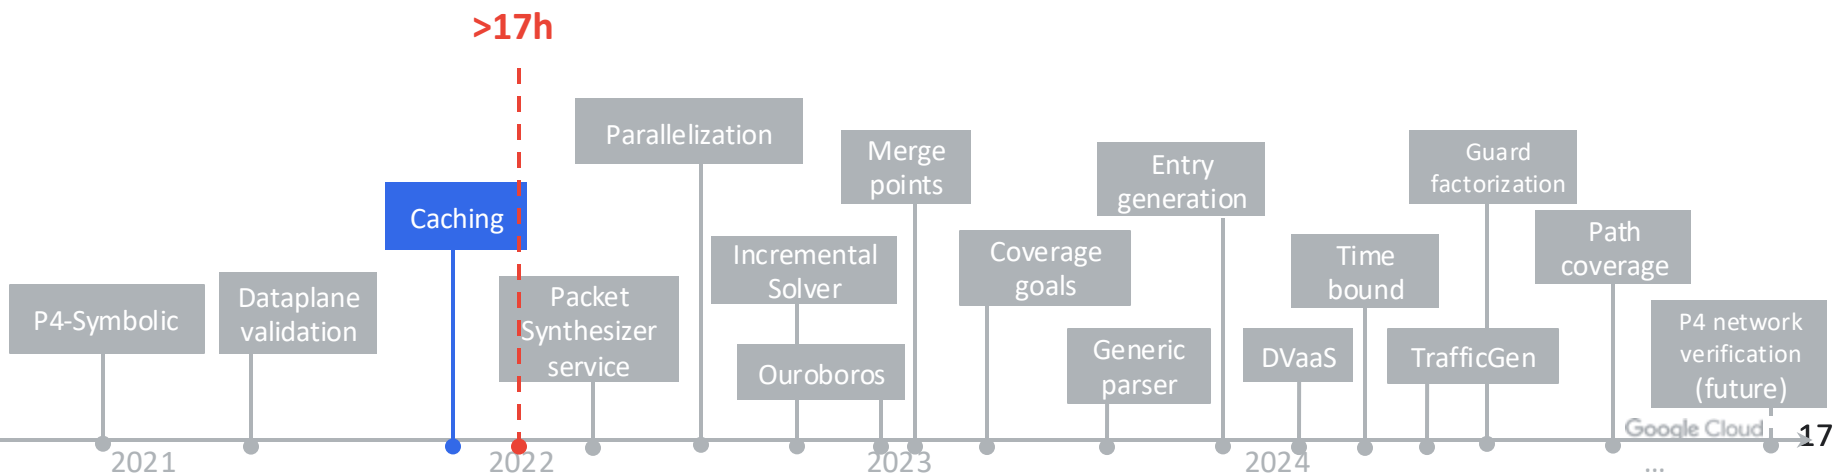

2. Parallelization

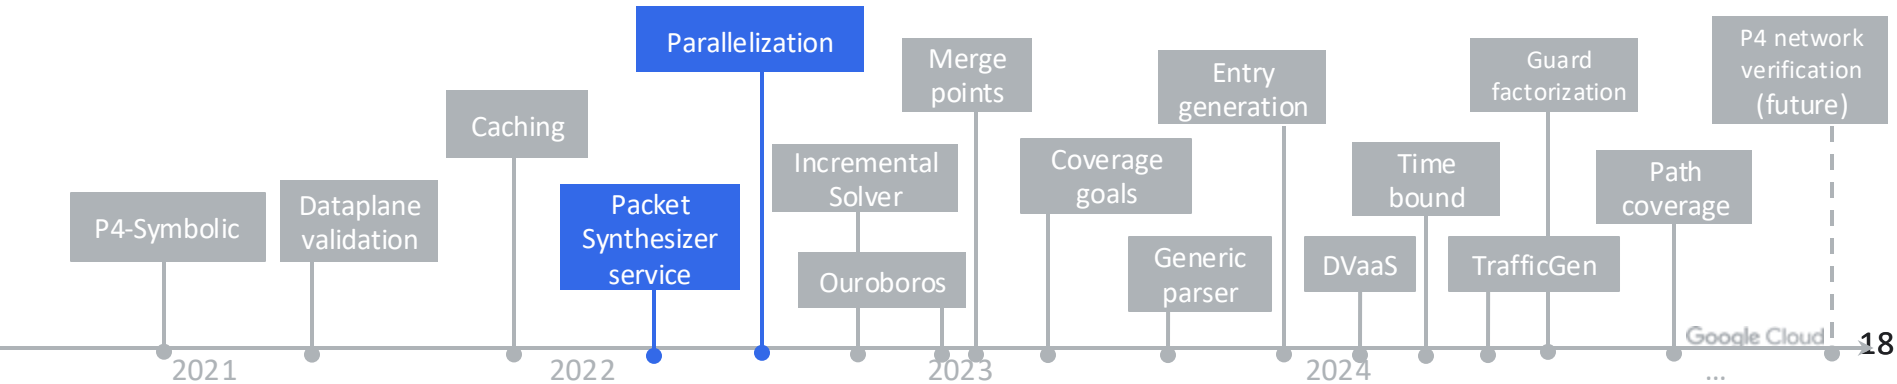

2. Parallelization

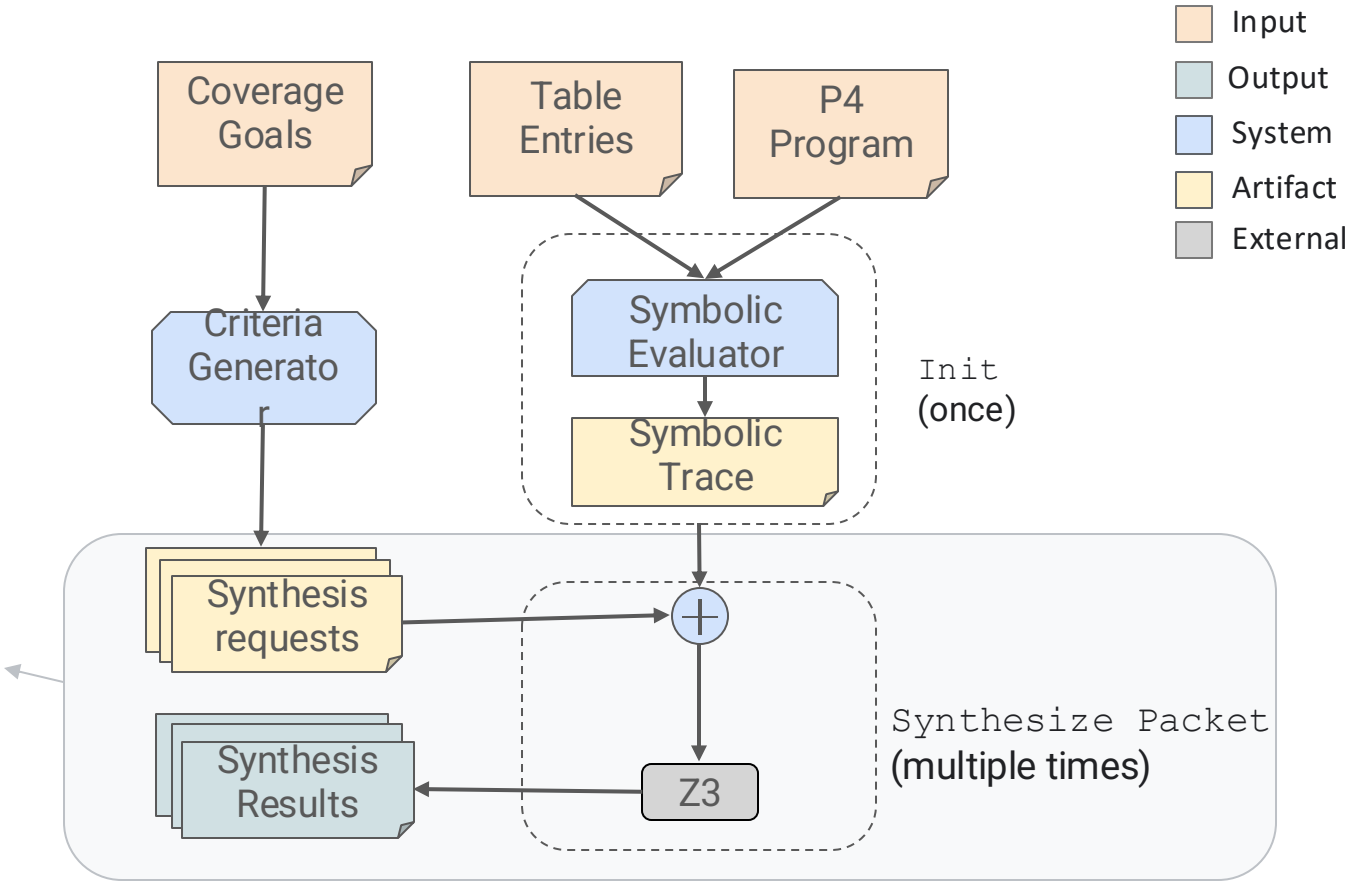

Independent of each other

⇓

embarrassingly parallelizable

2. Parallelization

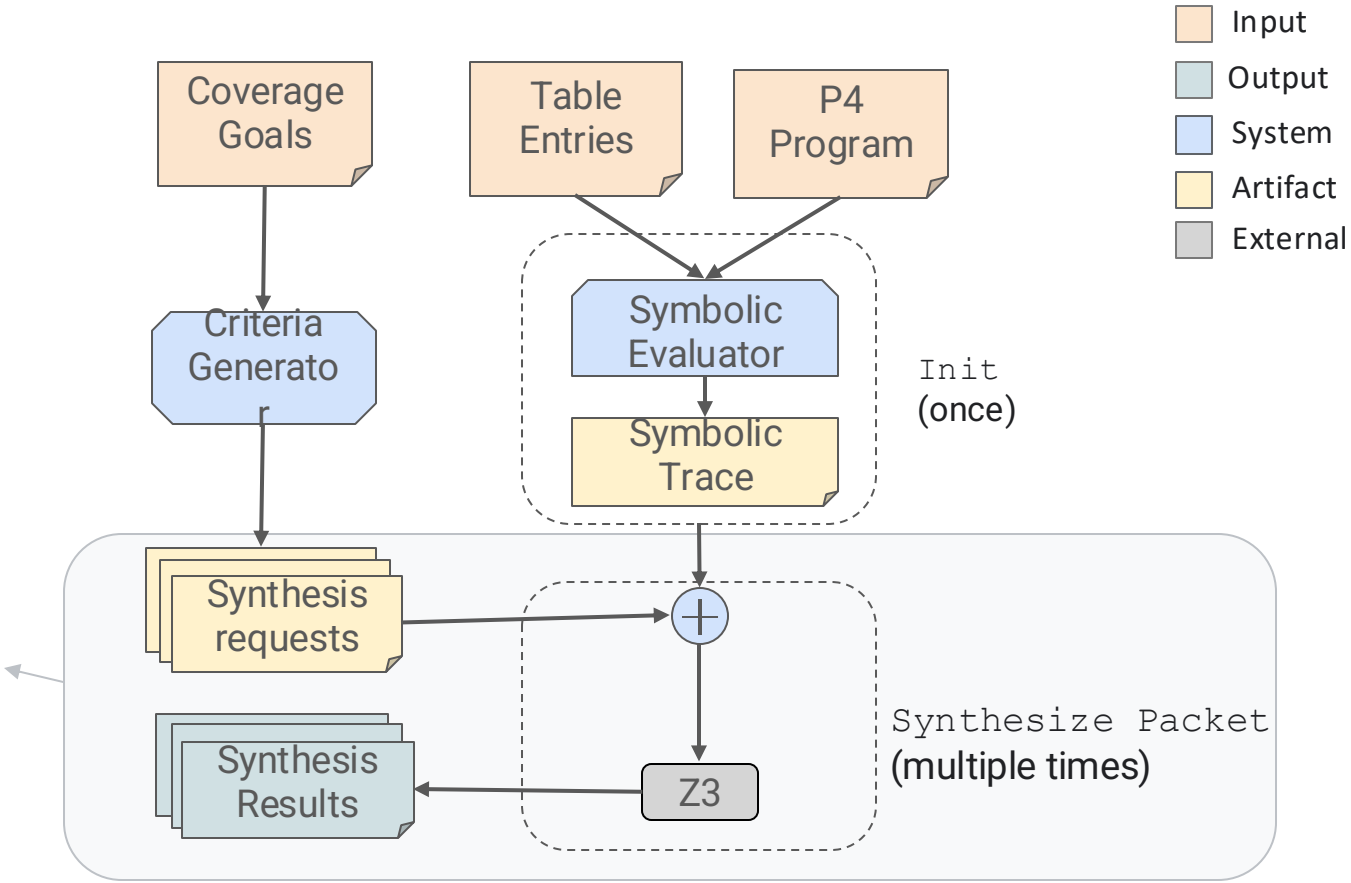

Independent of each other

⇓

embarrassingly parallelizable

2. Parallelization

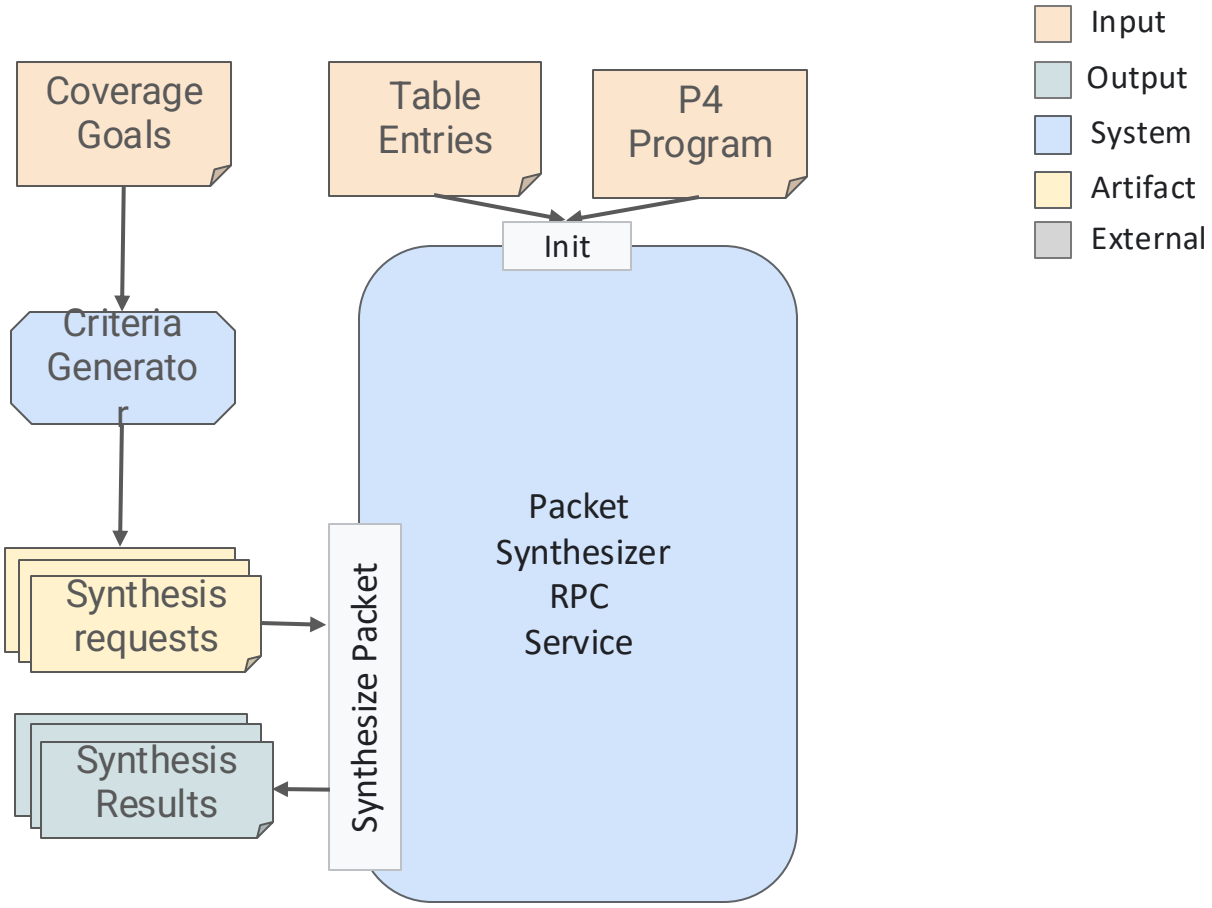

Independent of each other

=

embarrassingly parallelizable

2. Parallelization

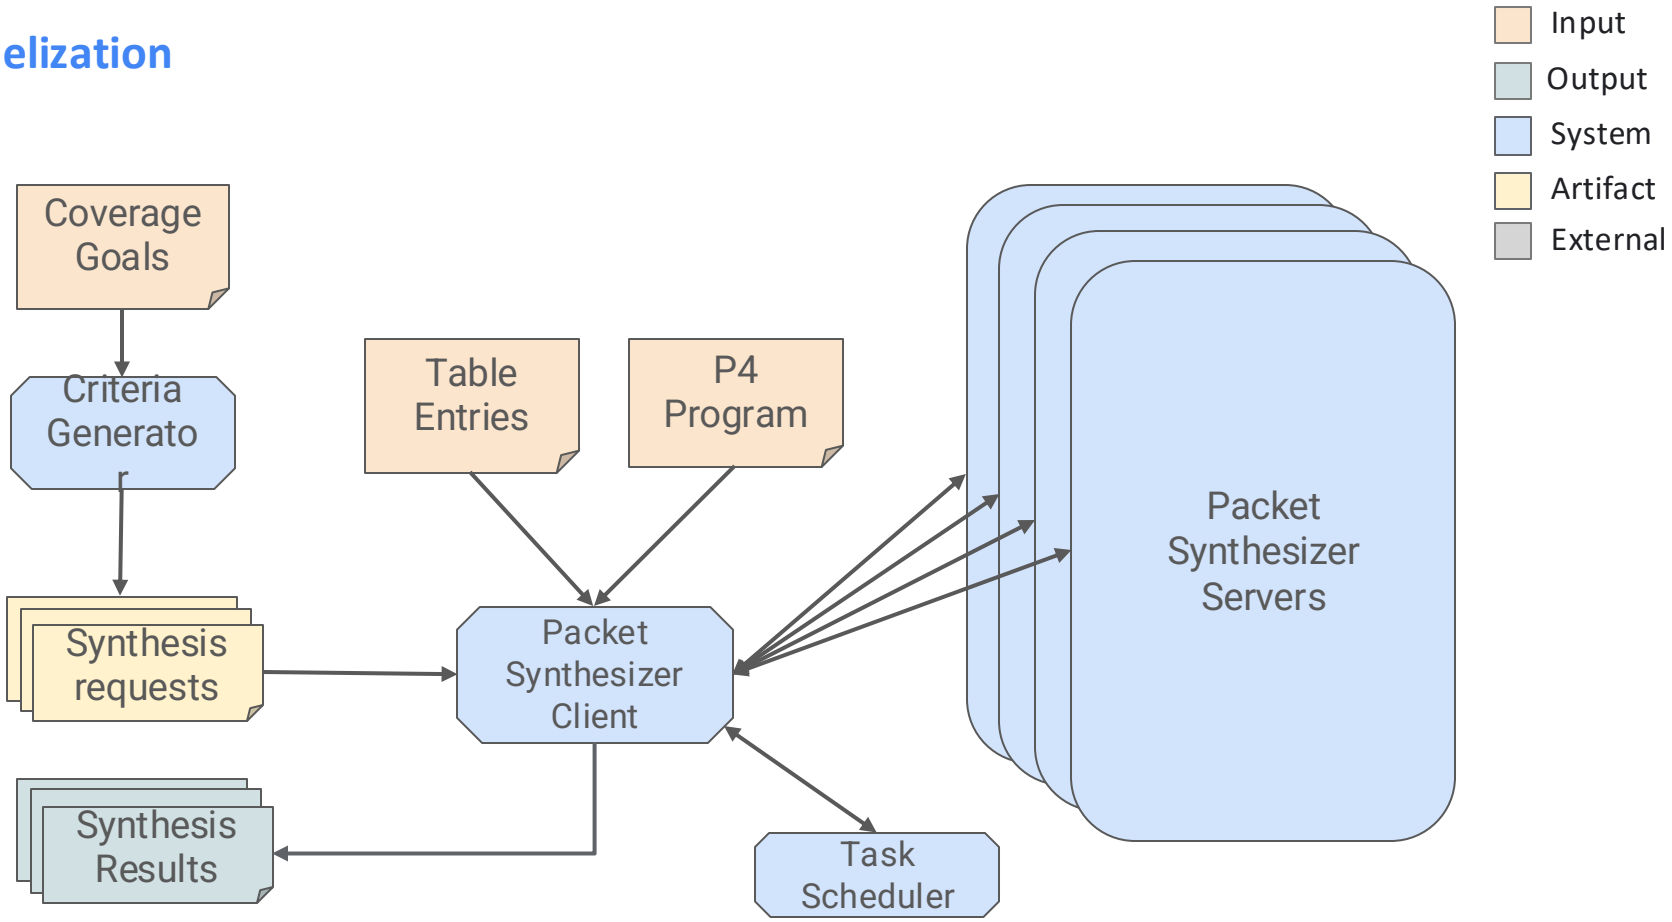

3. Symbolic Execution Merge Points

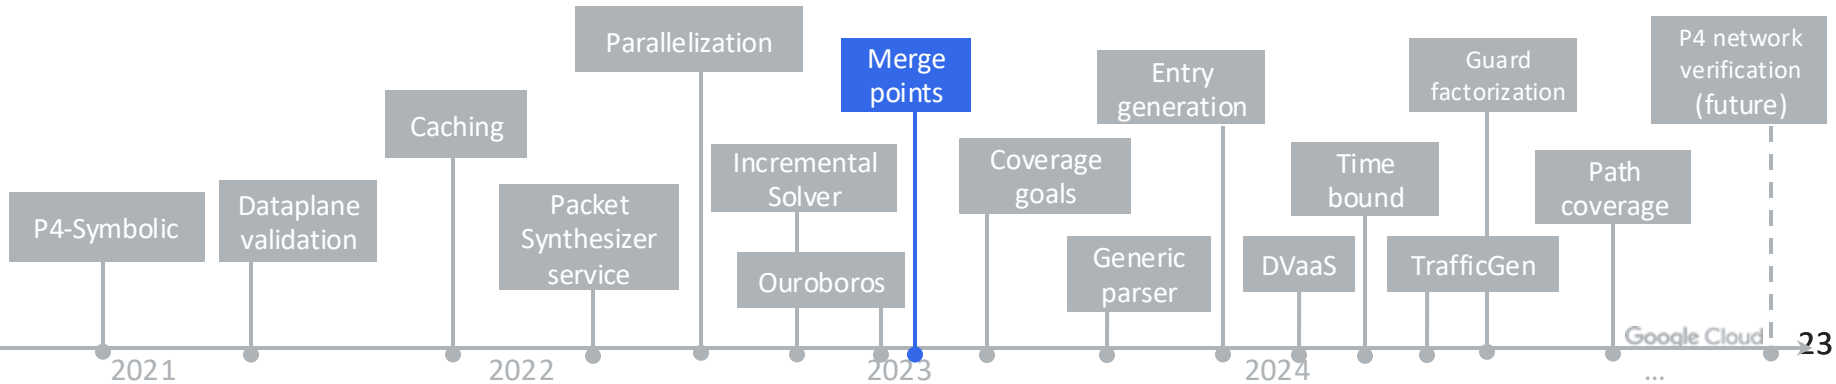

3. Symbolic Execution Merge Points

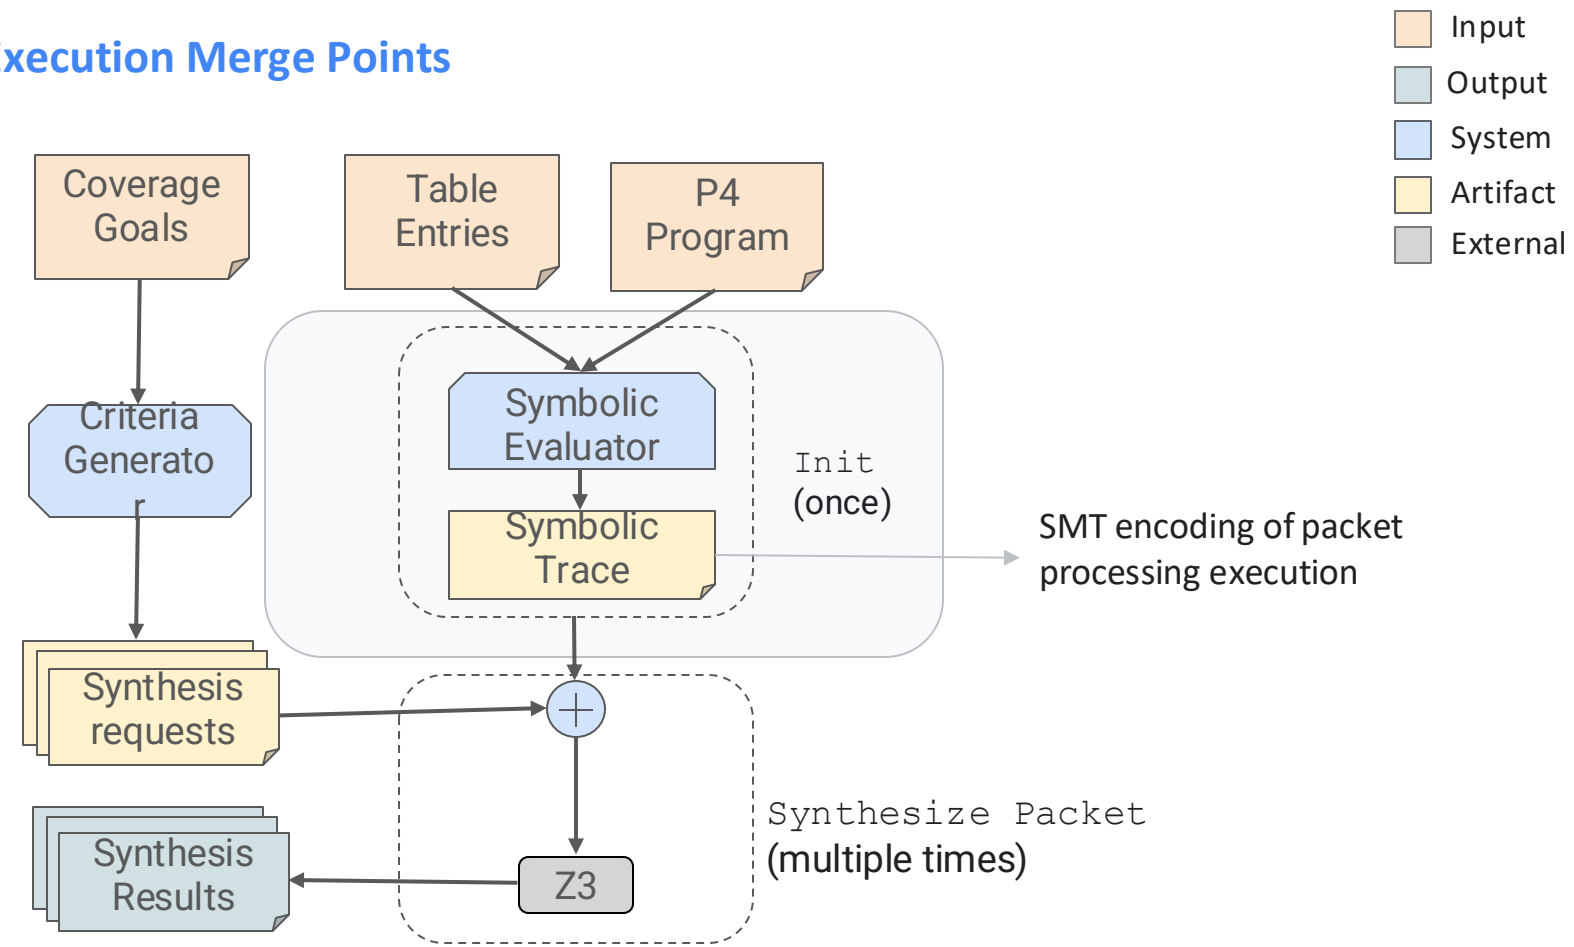

3. Symbolic Execution Merge Points

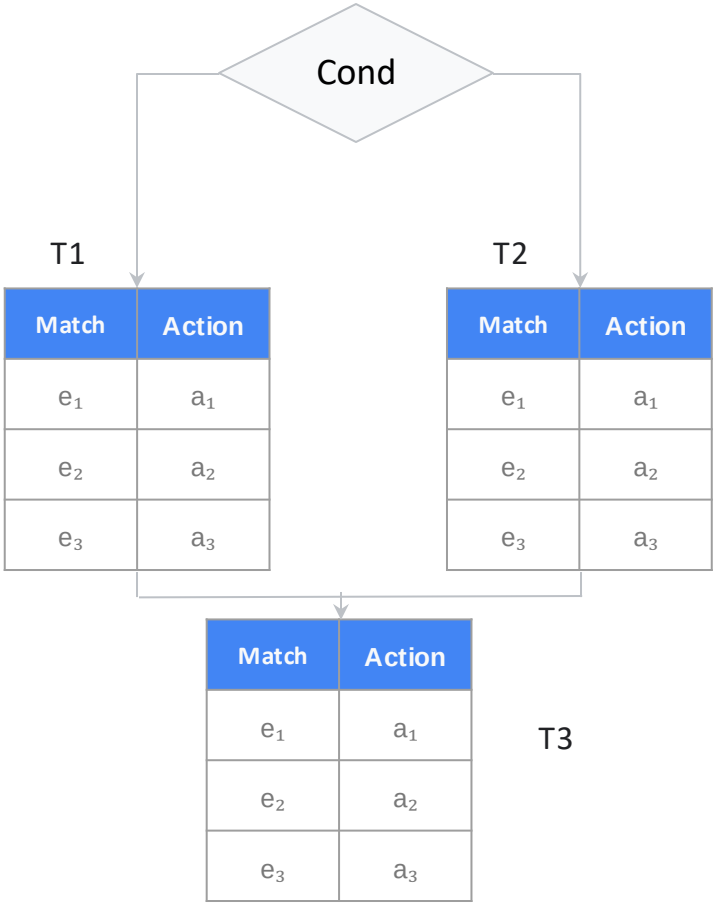

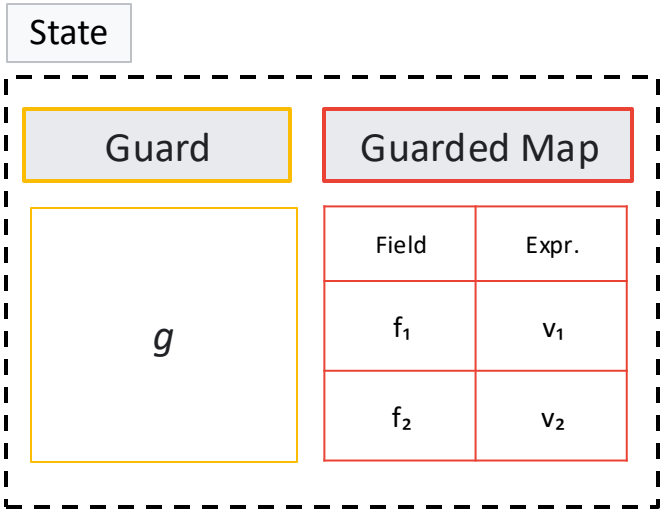

3. Symbolic Execution Merge Points

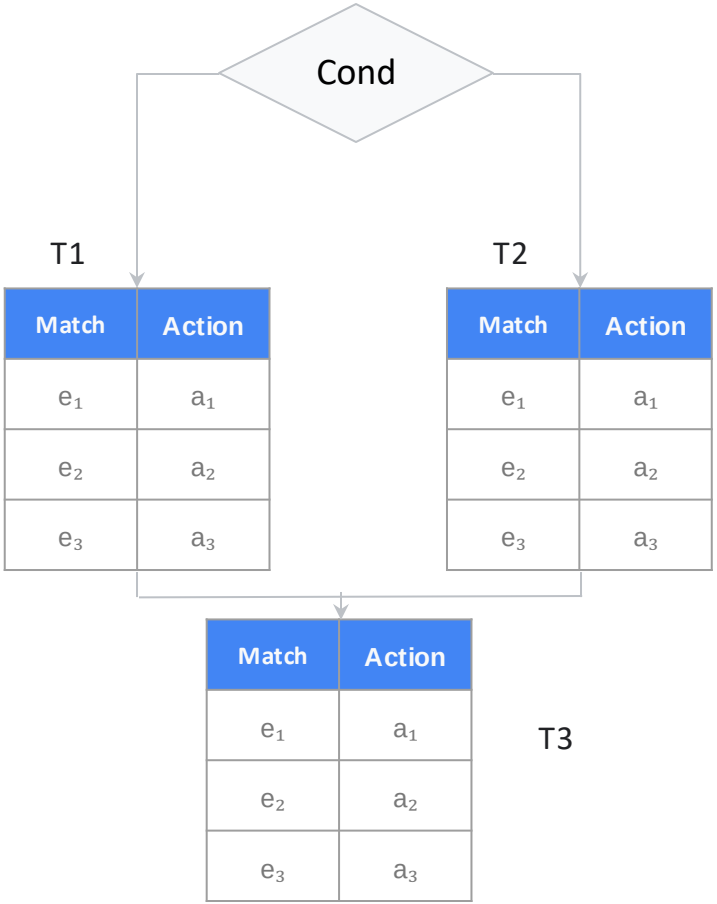

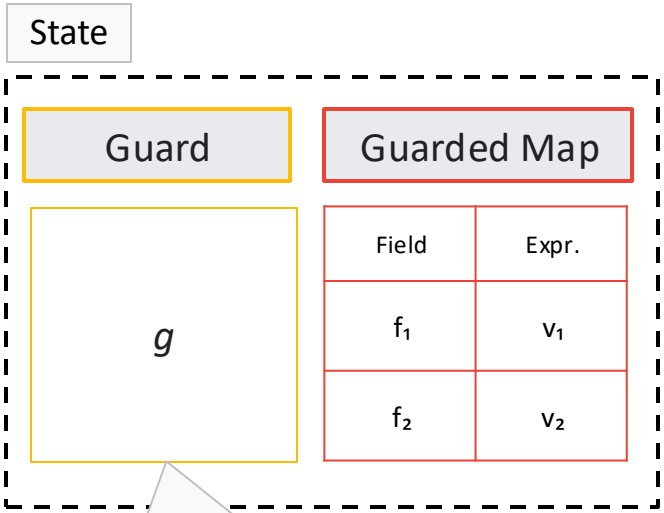

Guard: conditions that allow execution to reach the current point

When is it updated: (1) Conditionals, (2) Table entry **match** condition

3. Symbolic Execution Merge Points

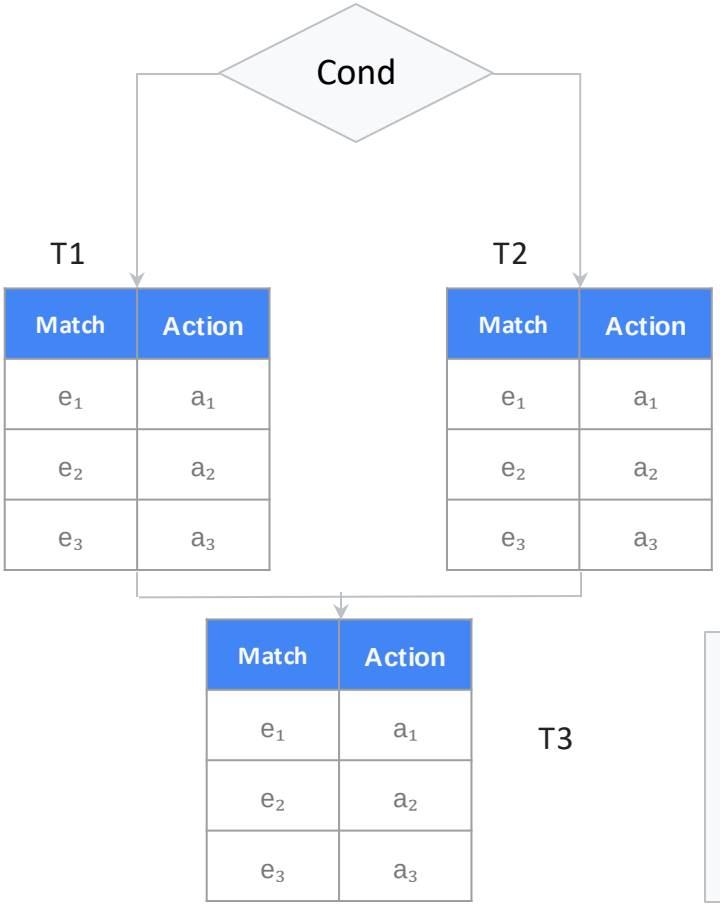

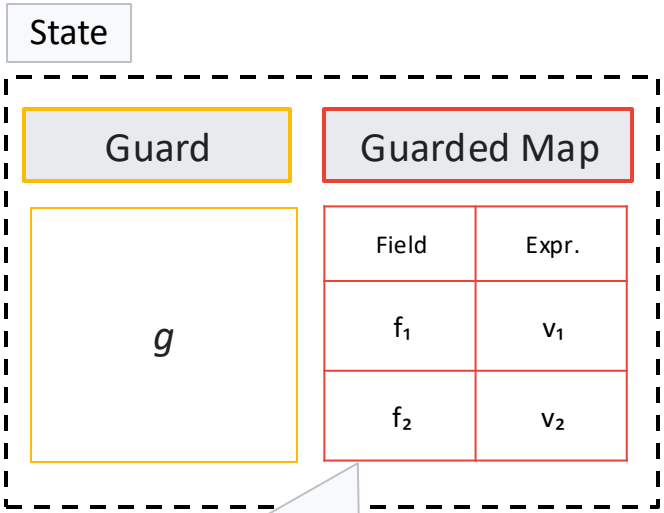

Guarded Map: *fields in header + metadata -> SMT expression* (value of the field at the current execution point)

When is it updated: Table entry **action**

3. Symbolic Execution Merge Points

No Merge Points

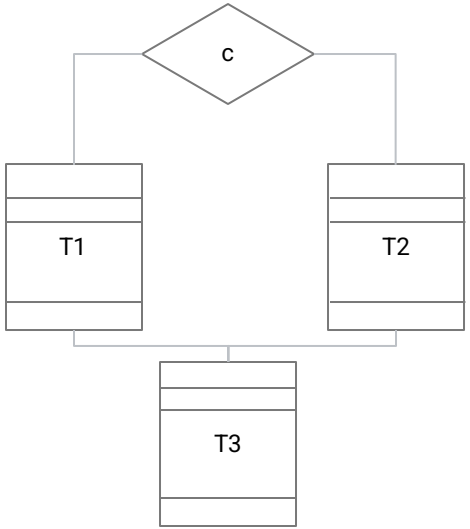

With Merge Points

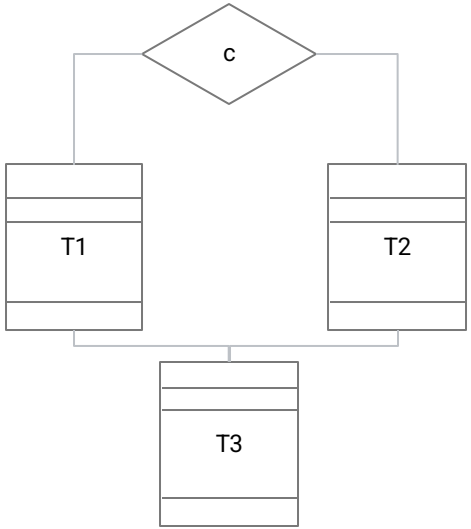

3. Symbolic Execution Merge Points

No Merge Points

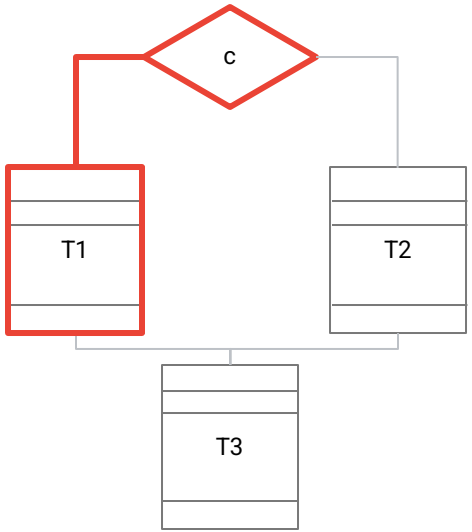

With Merge Points

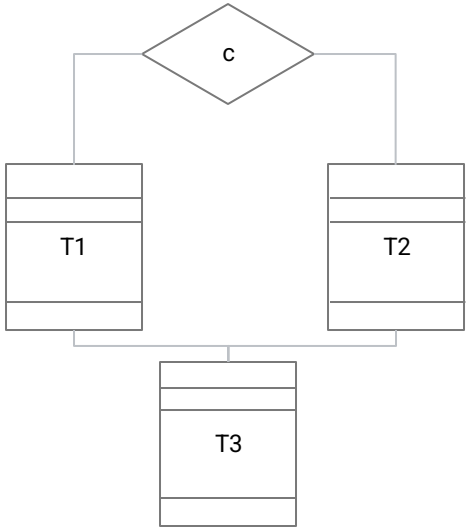

3. Symbolic Execution Merge Points

No Merge Points

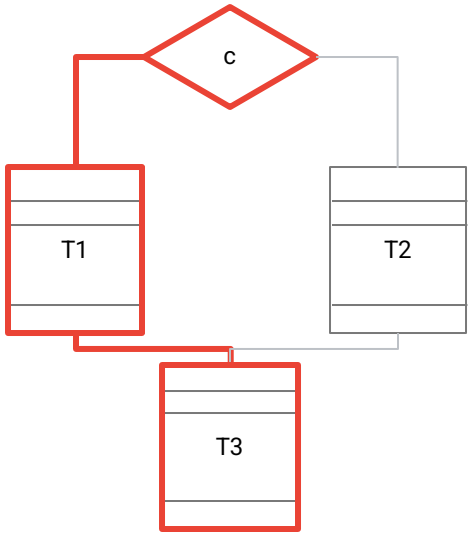

With Merge Points

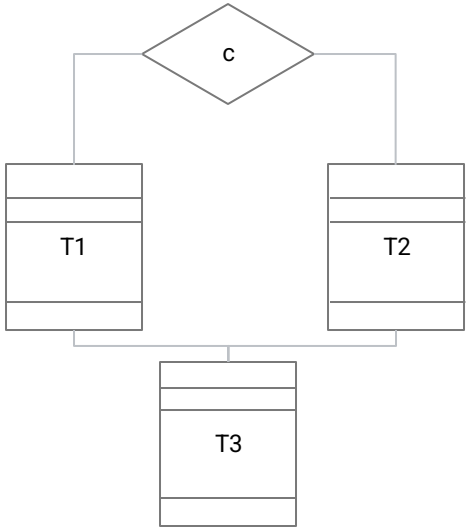

3. Symbolic Execution Merge Points

No Merge Points

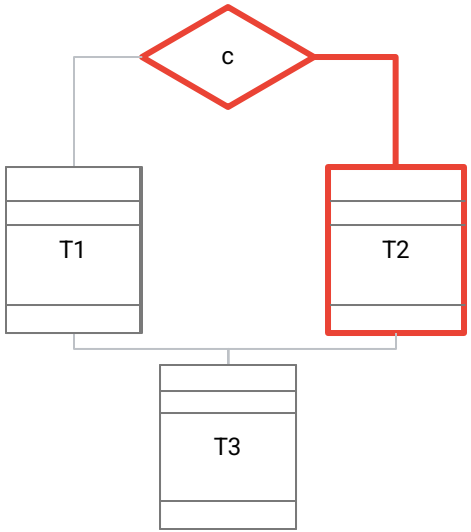

With Merge Points

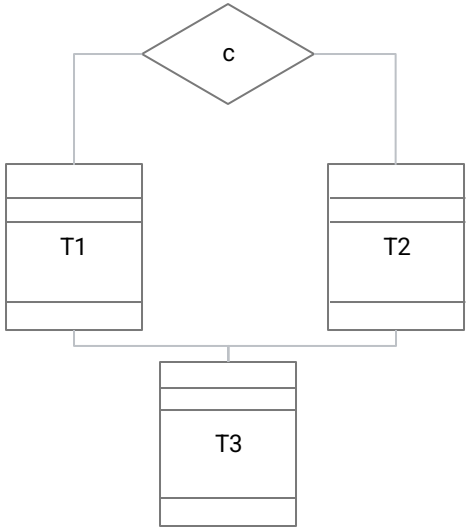

3. Symbolic Execution Merge Points

No Merge Points

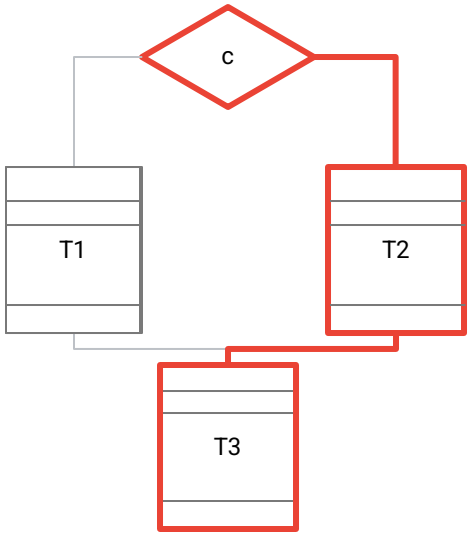

With Merge Points

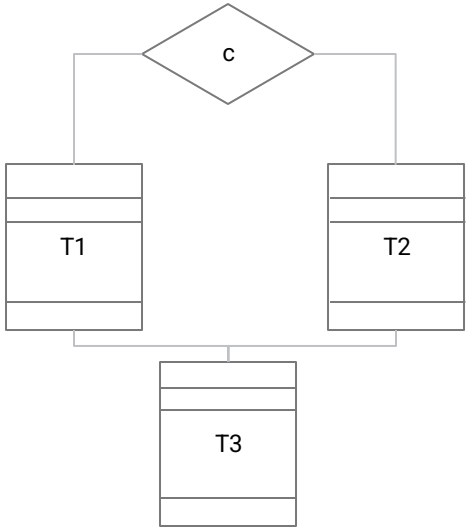

3. Symbolic Execution Merge Points

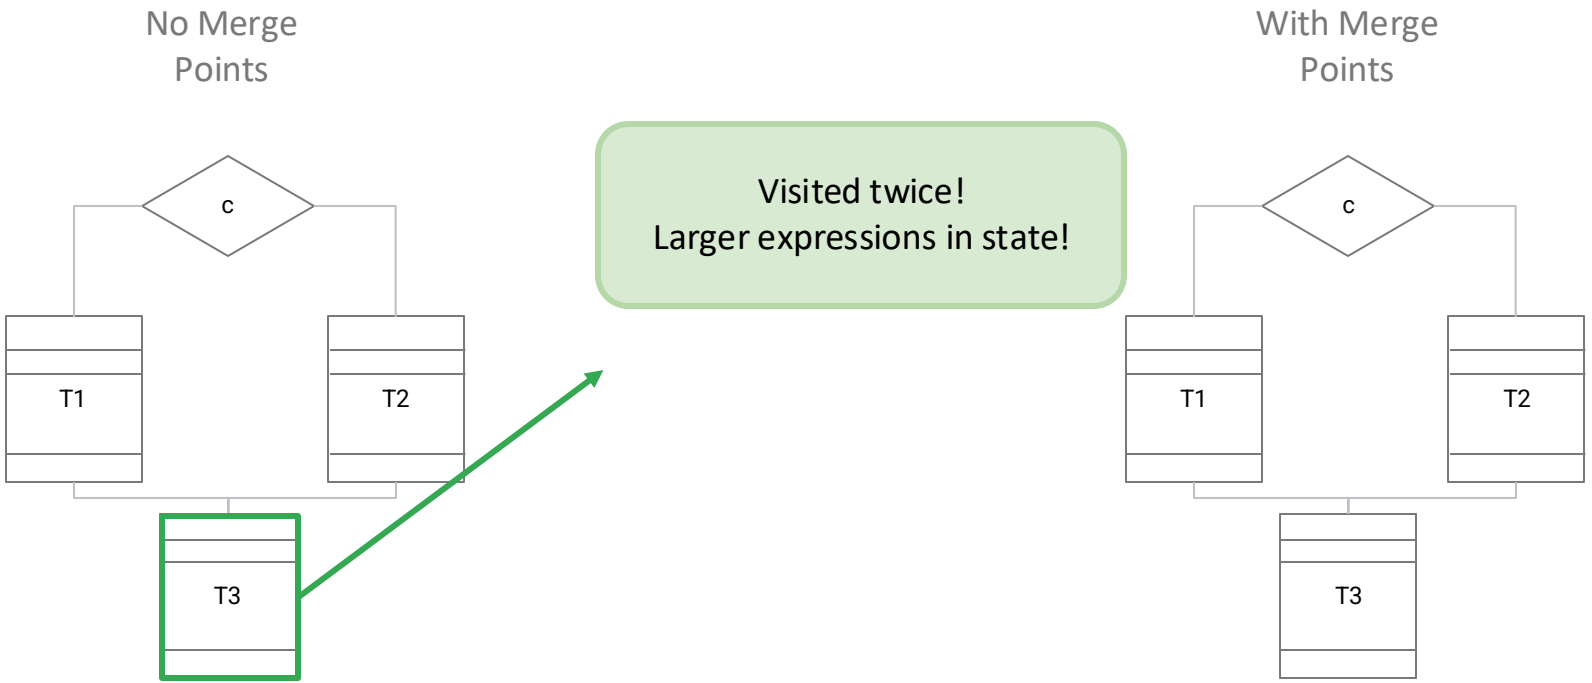

3. Symbolic Execution Merge Points

No Merge Points

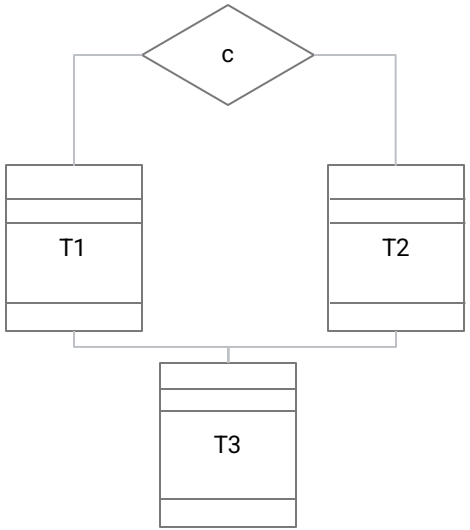

With Merge Points

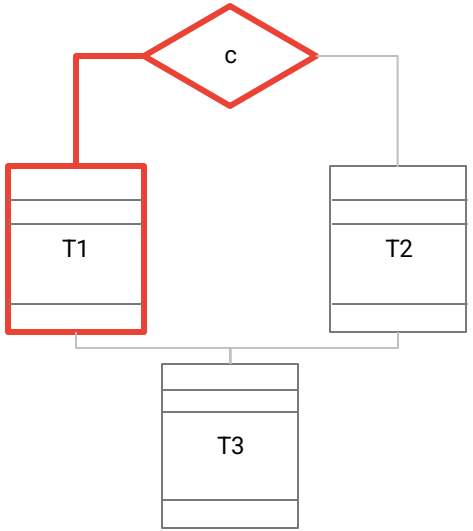

3. Symbolic Execution Merge Points

No Merge Points

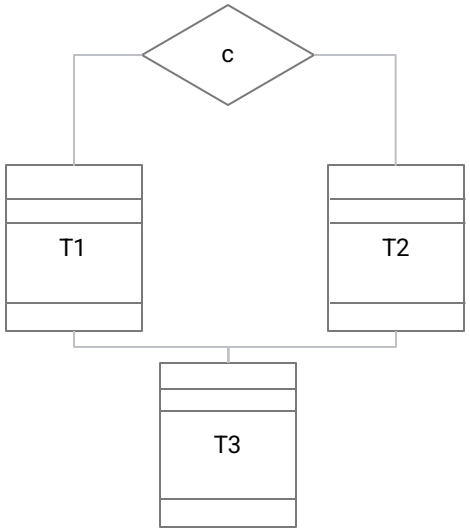

With Merge Points

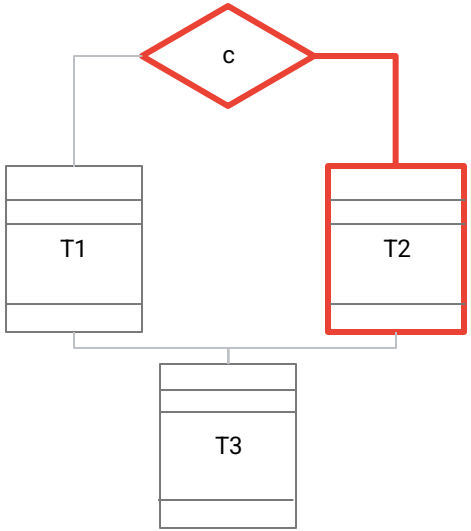

3. Symbolic Execution Merge Points

No Merge Points

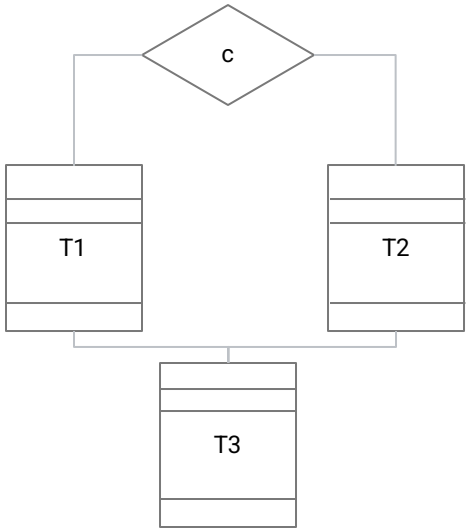

With Merge Points

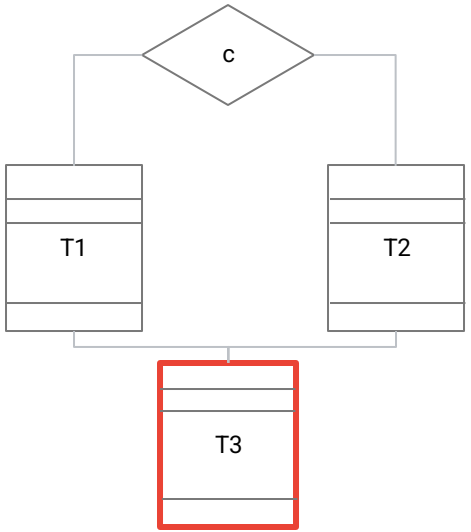

3. Symbolic Execution Merge Points

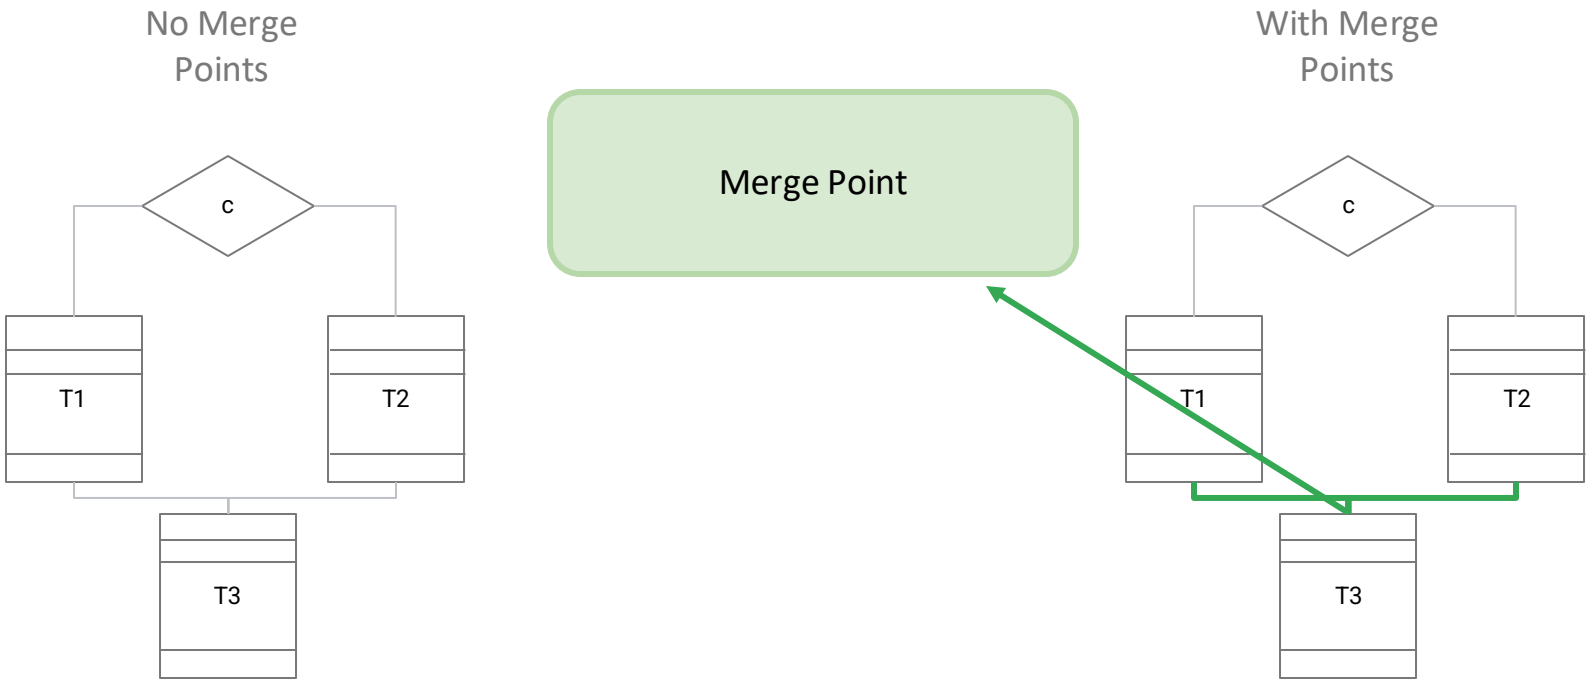

3. Symbolic Execution Merge Points

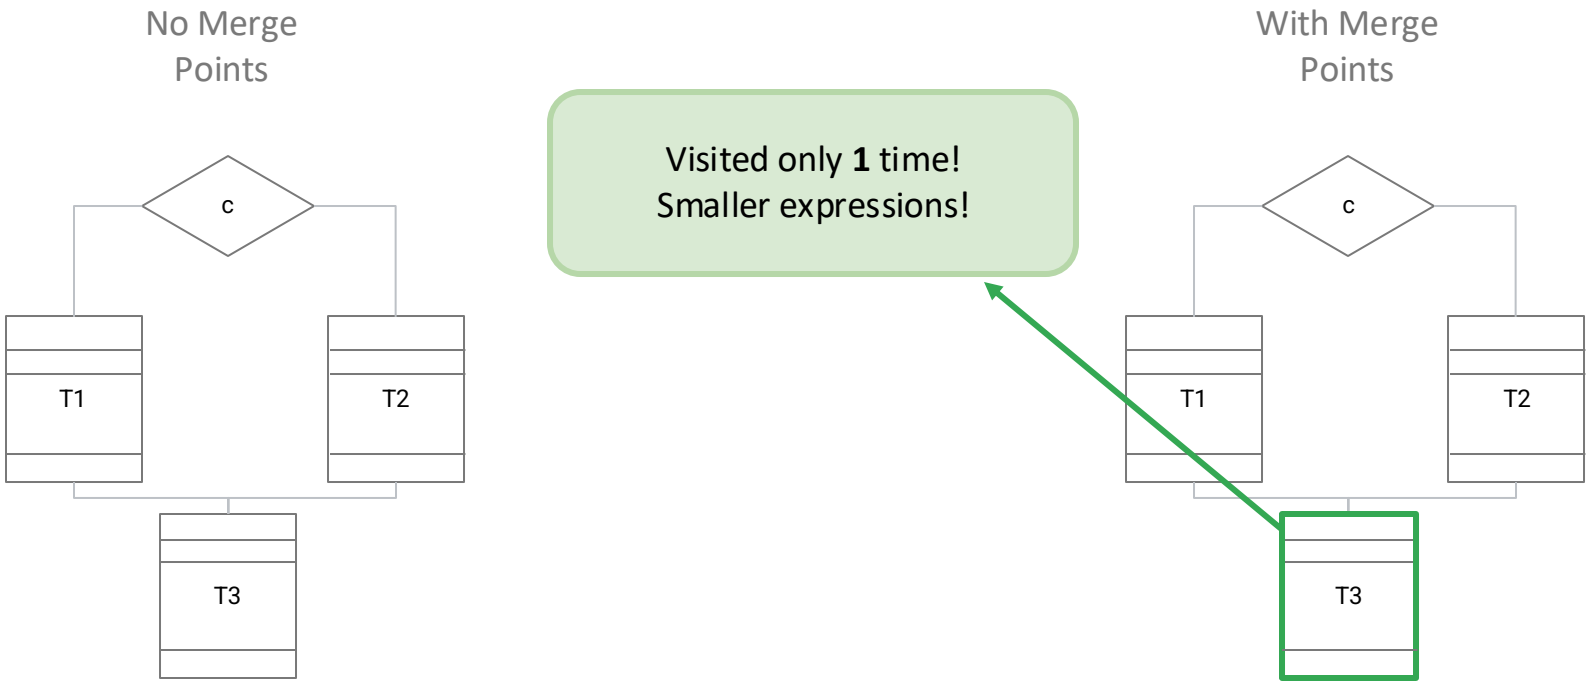

Results so far

	Product 1		Product 2	
	Clos stage 2	Clos stage 3	Clos stage 2	Clos stage 3
# Packet synthesis requests ≈ {entries}x{packet fate}	~1000	~1000	~1500	~3500
Runtime (before improvements)	~10 mins	~10 mins	~40 mins	**~7 hours ↑** **(17 hours at some point)**
Runtime **(with improvements)** parallelization, merge points	**<5s**	**<5s**	**<30s**	**~1m**

Did we solve the problem?

YES!

but only temporarily! :(

Last resorts
- Reduced coverage
- Re-enabled time-bound synthesis
- Relied on offline synthesis

New products and use cases
- More complex pipelines
- Significantly larger (5x) snapshots

More compute (servers) did not help

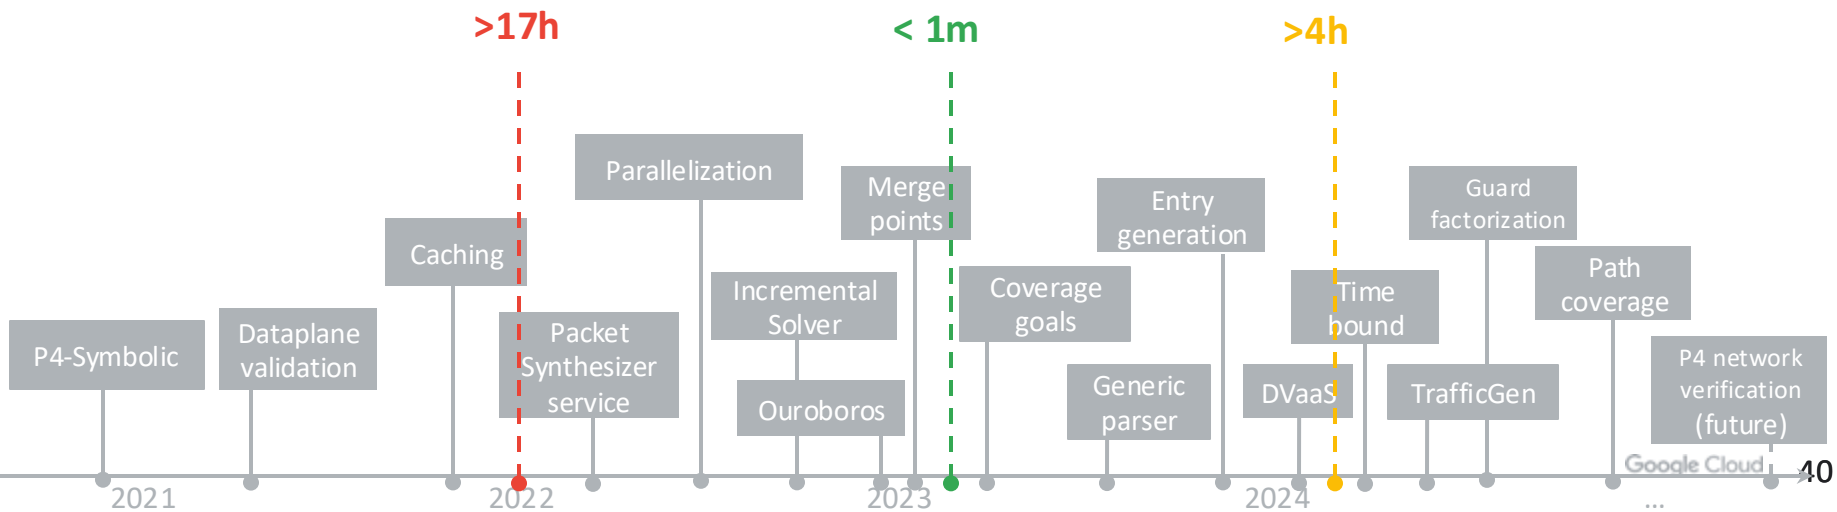

4. Guard Factorization

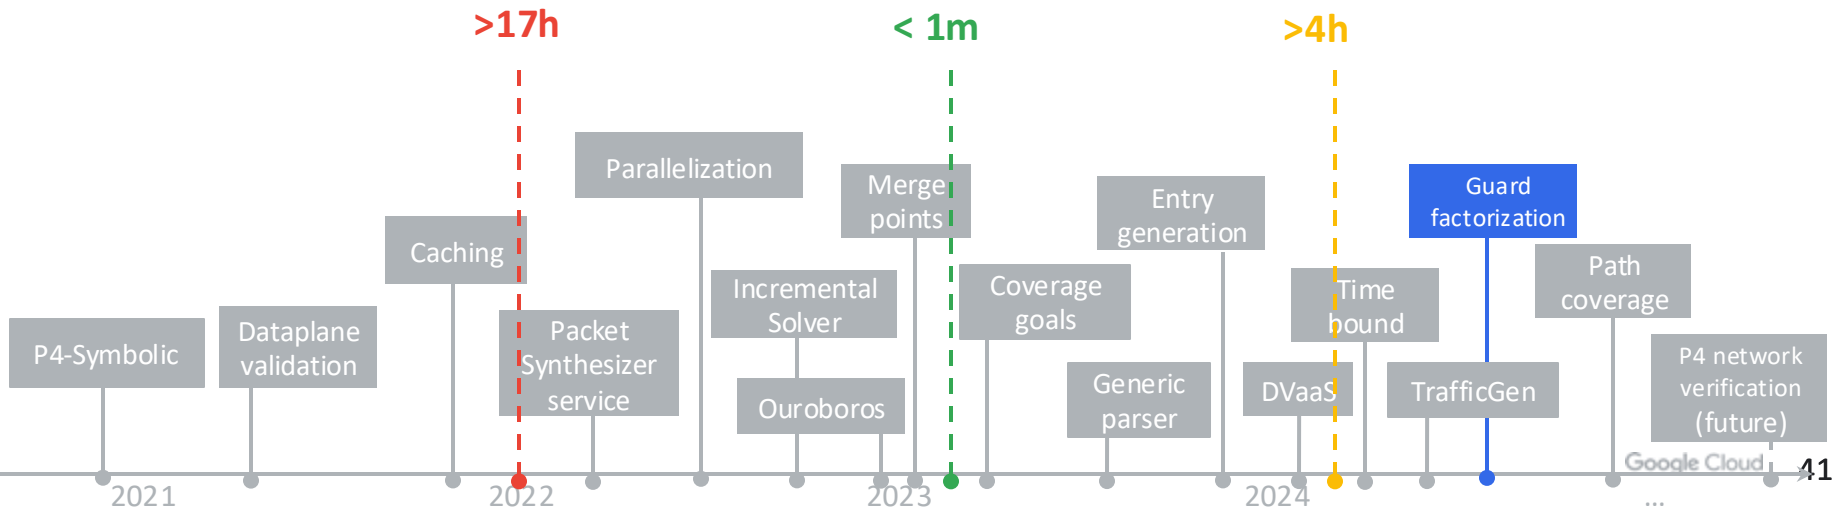

4. Guard Factorization

Current Symbolic Execution Flow

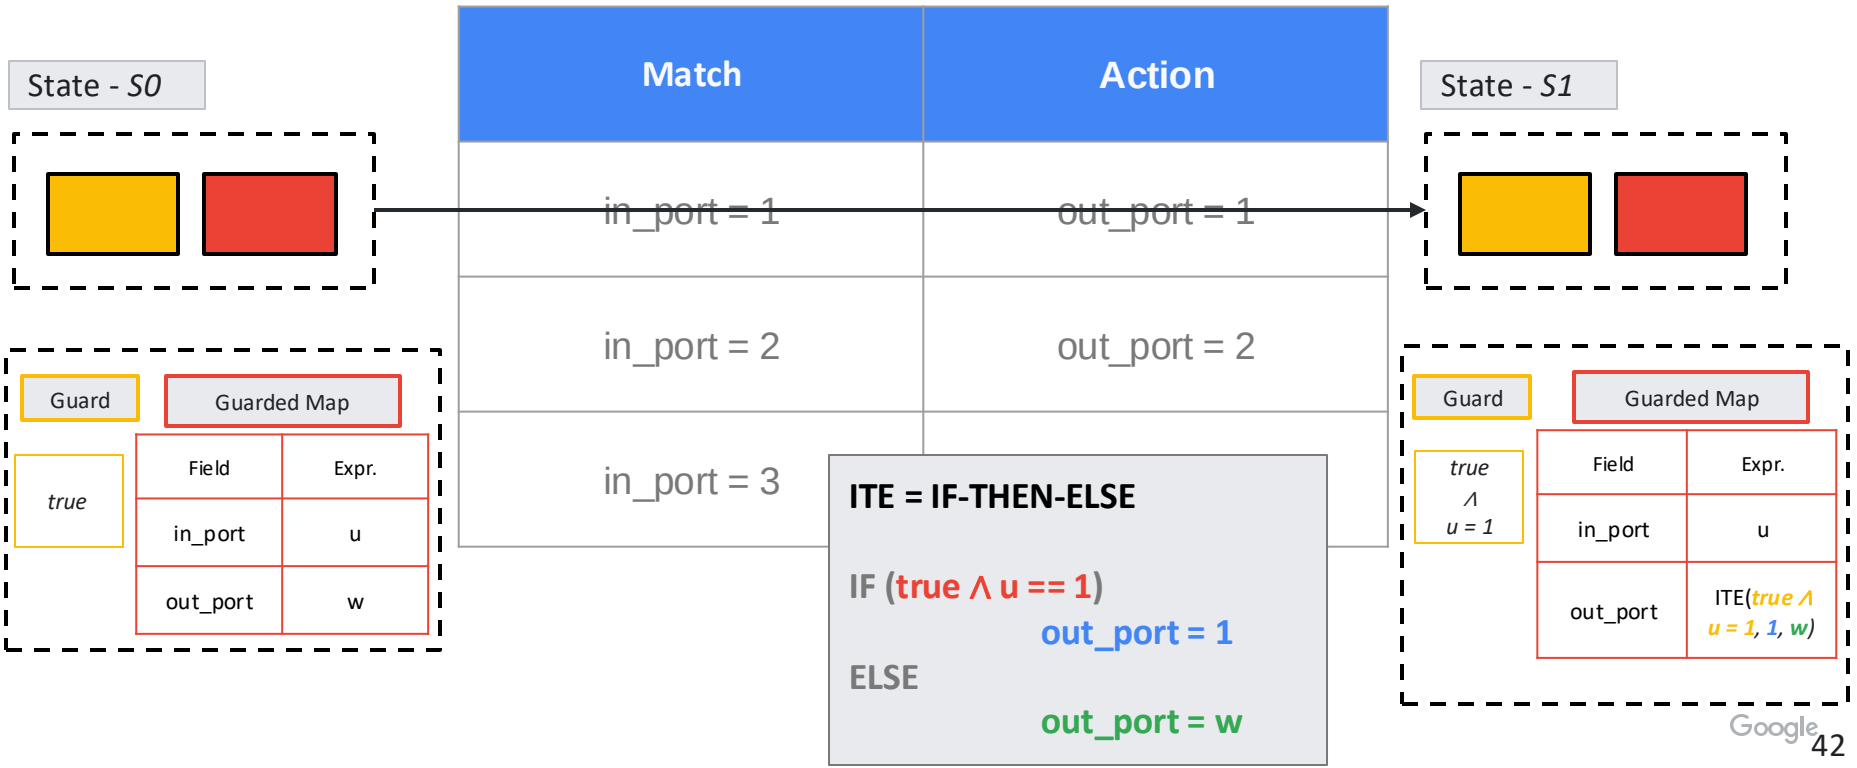

4. Guard Factorization

Current Symbolic Execution Flow

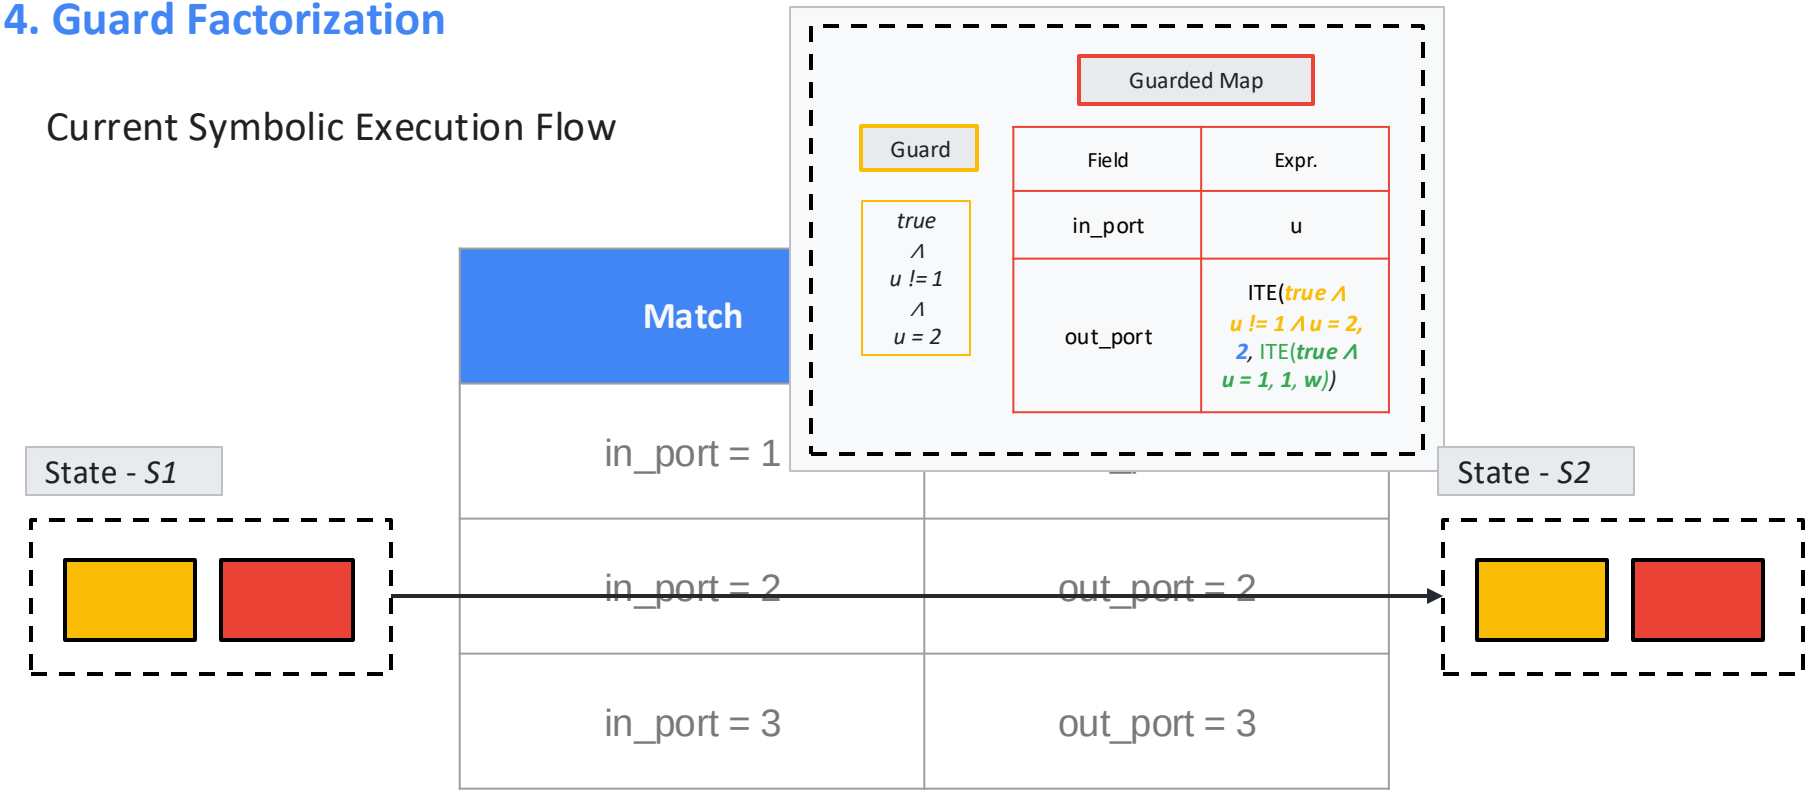

4. Guard Factorization

Current Symbolic Execution Flow

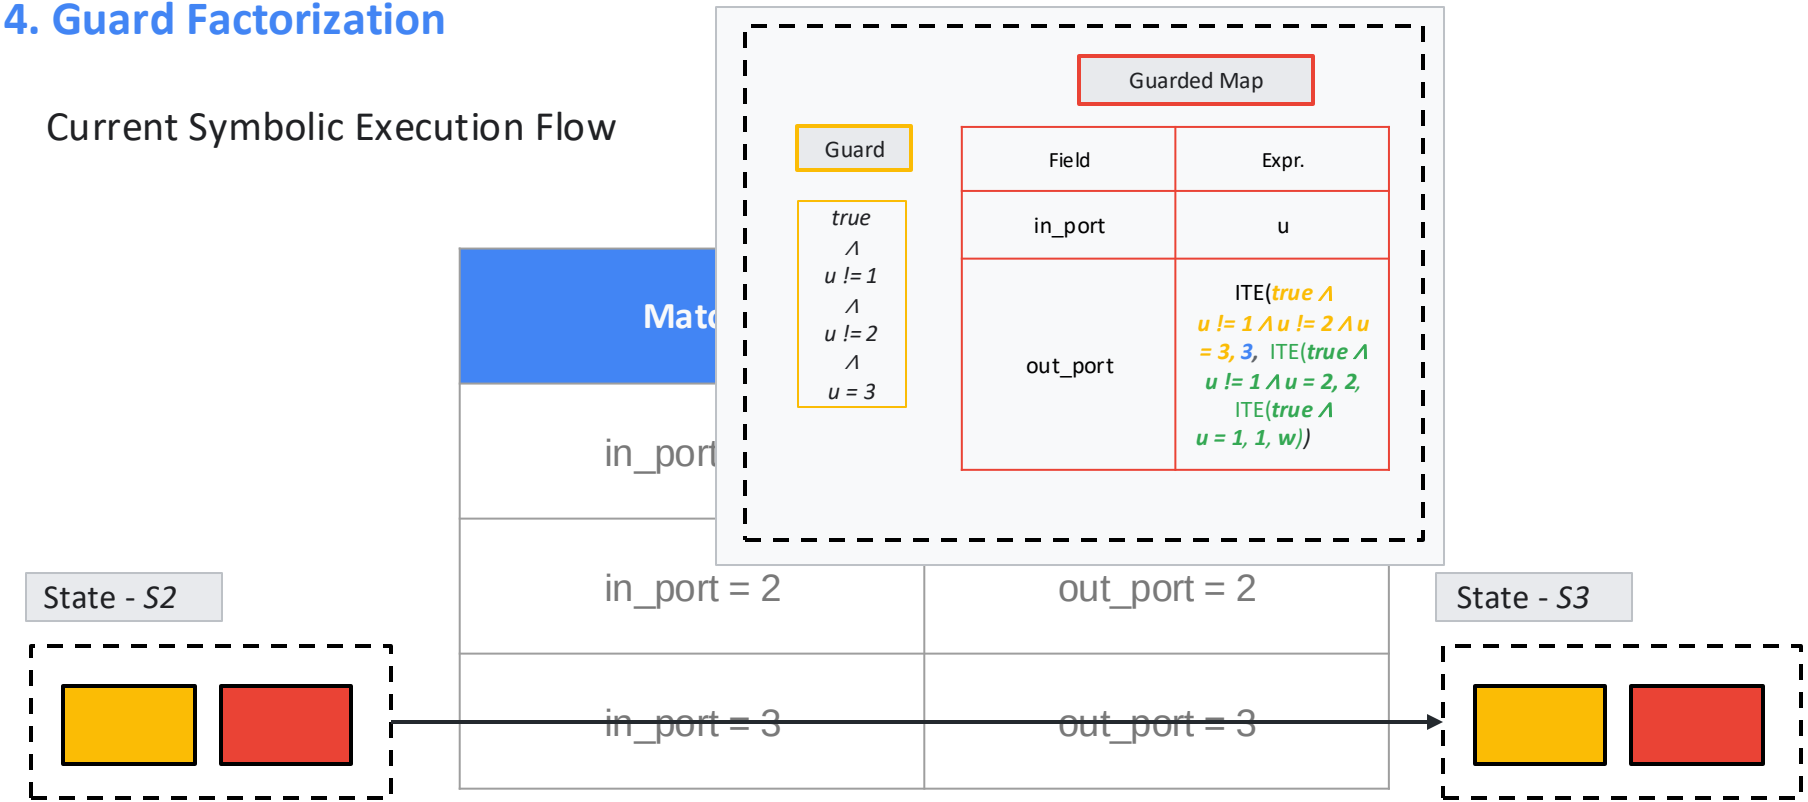

4. Guard Factorization

Optimized Symbolic Execution Flow

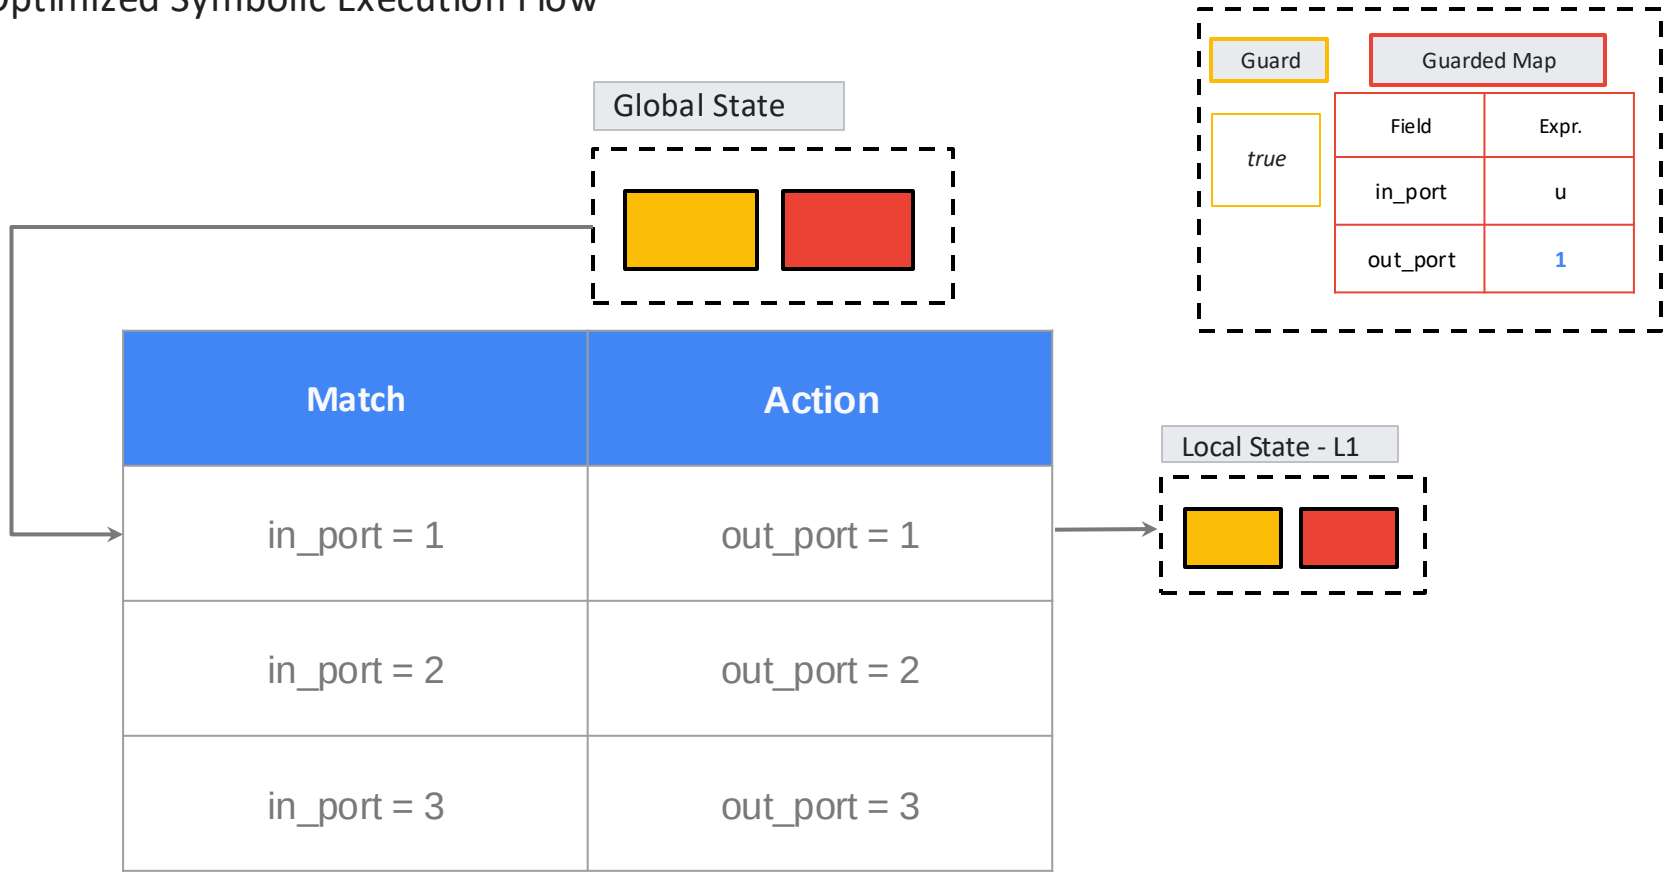

4. Guard Factorization

Optimized Symbolic Execution Flow

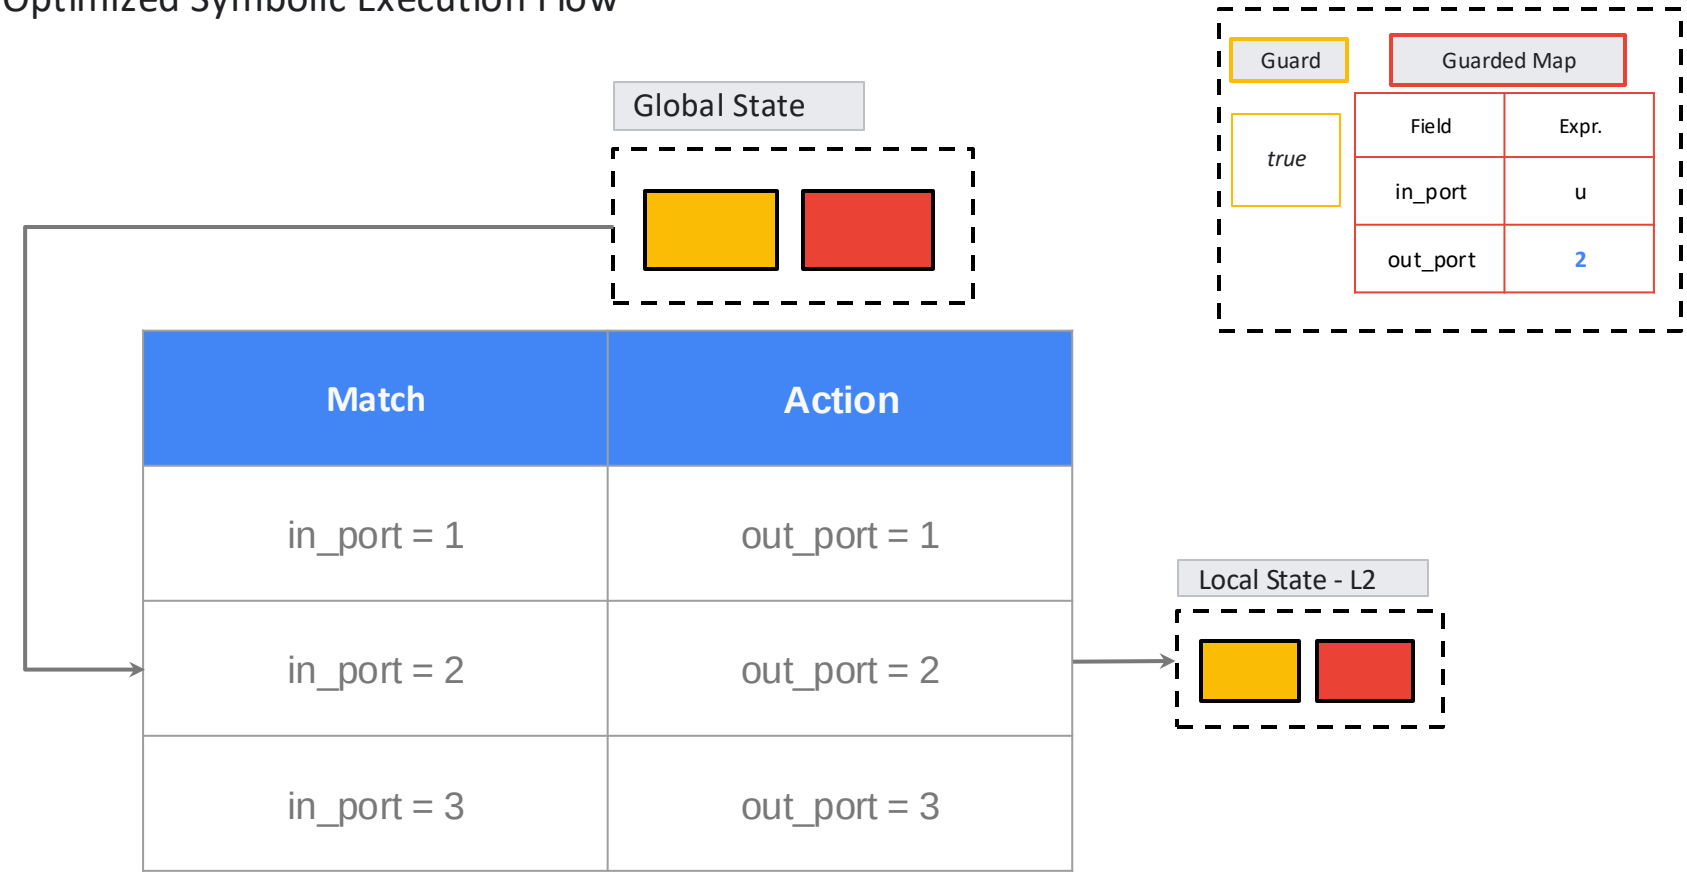

4. Guard Factorization

Optimized Symbolic Execution Flow

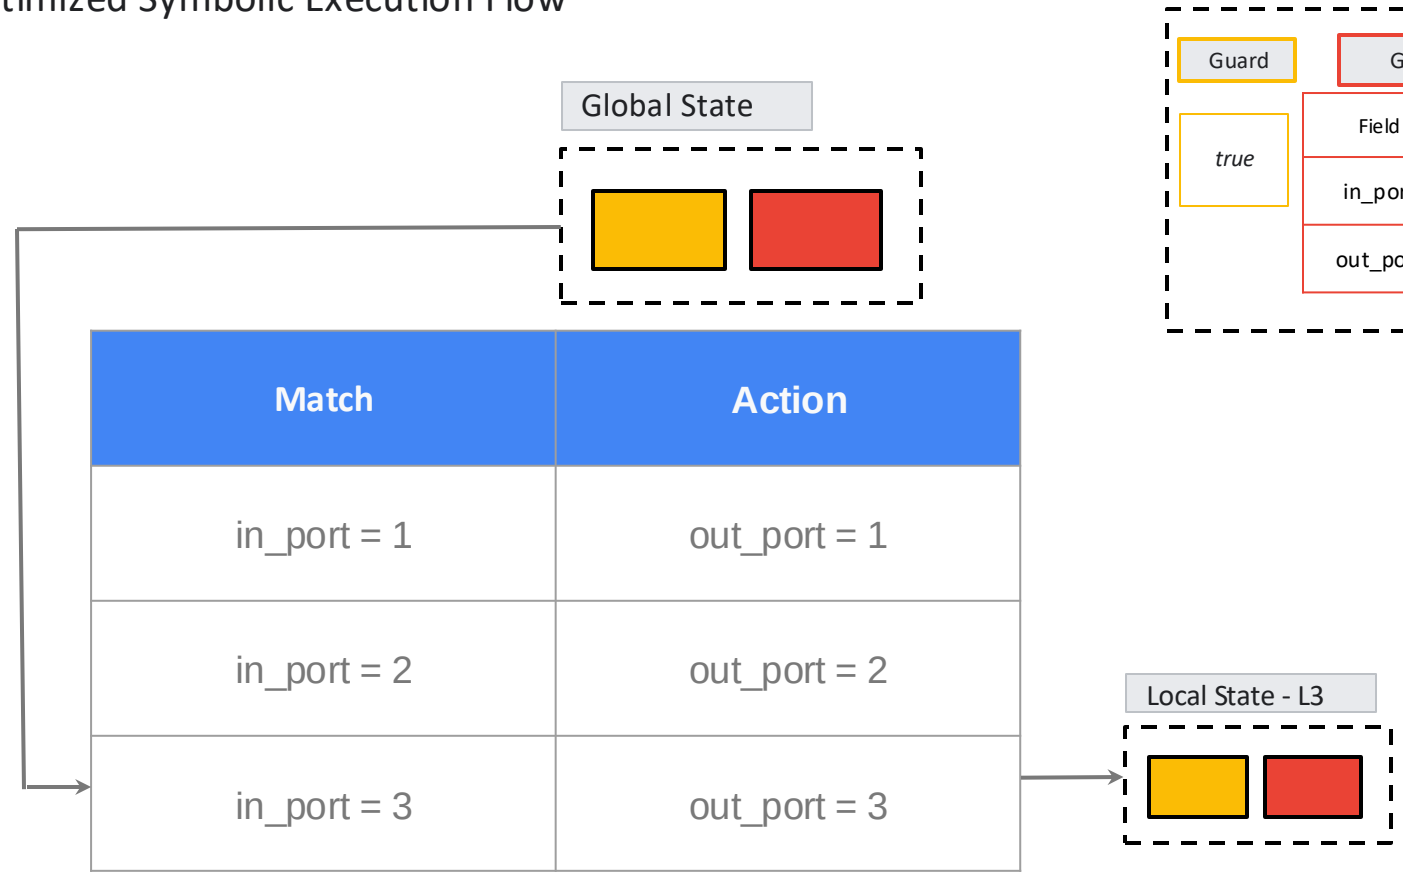

4. Guard Factorization

Optimized Symbolic Execution Flow

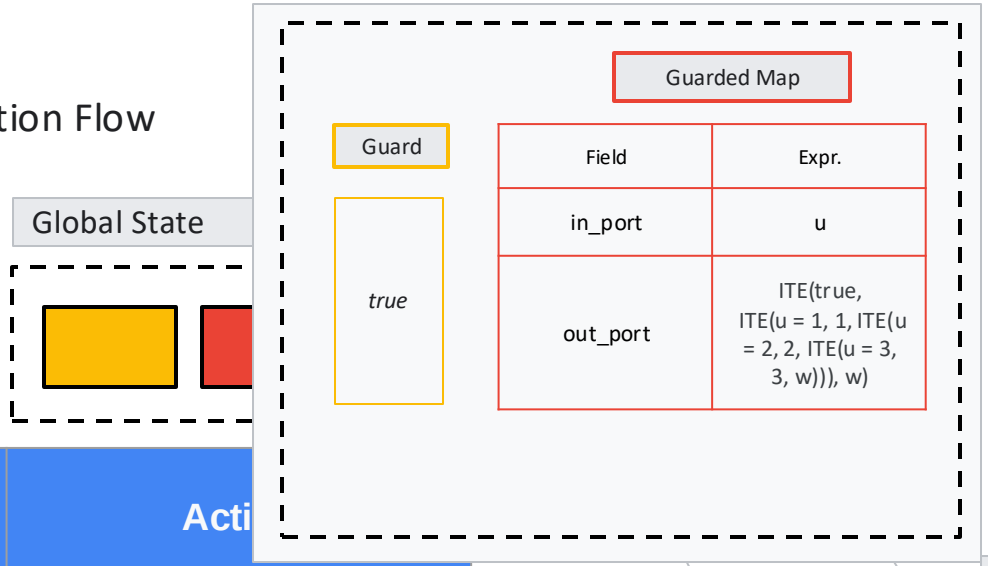

Match	Acti
in_port = 1	out_port = 1
in_port = 2	out_port = 2
in_port = 3	out_port = 3

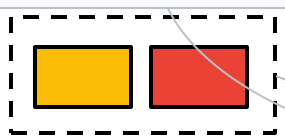

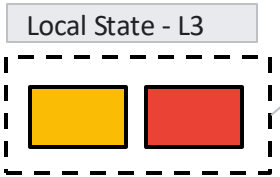

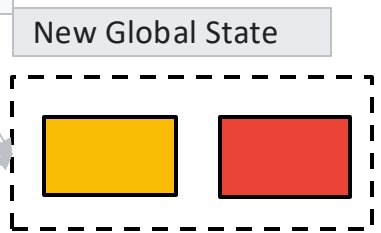

4. Guard Factorization

Optimized Symbolic Execution Flow

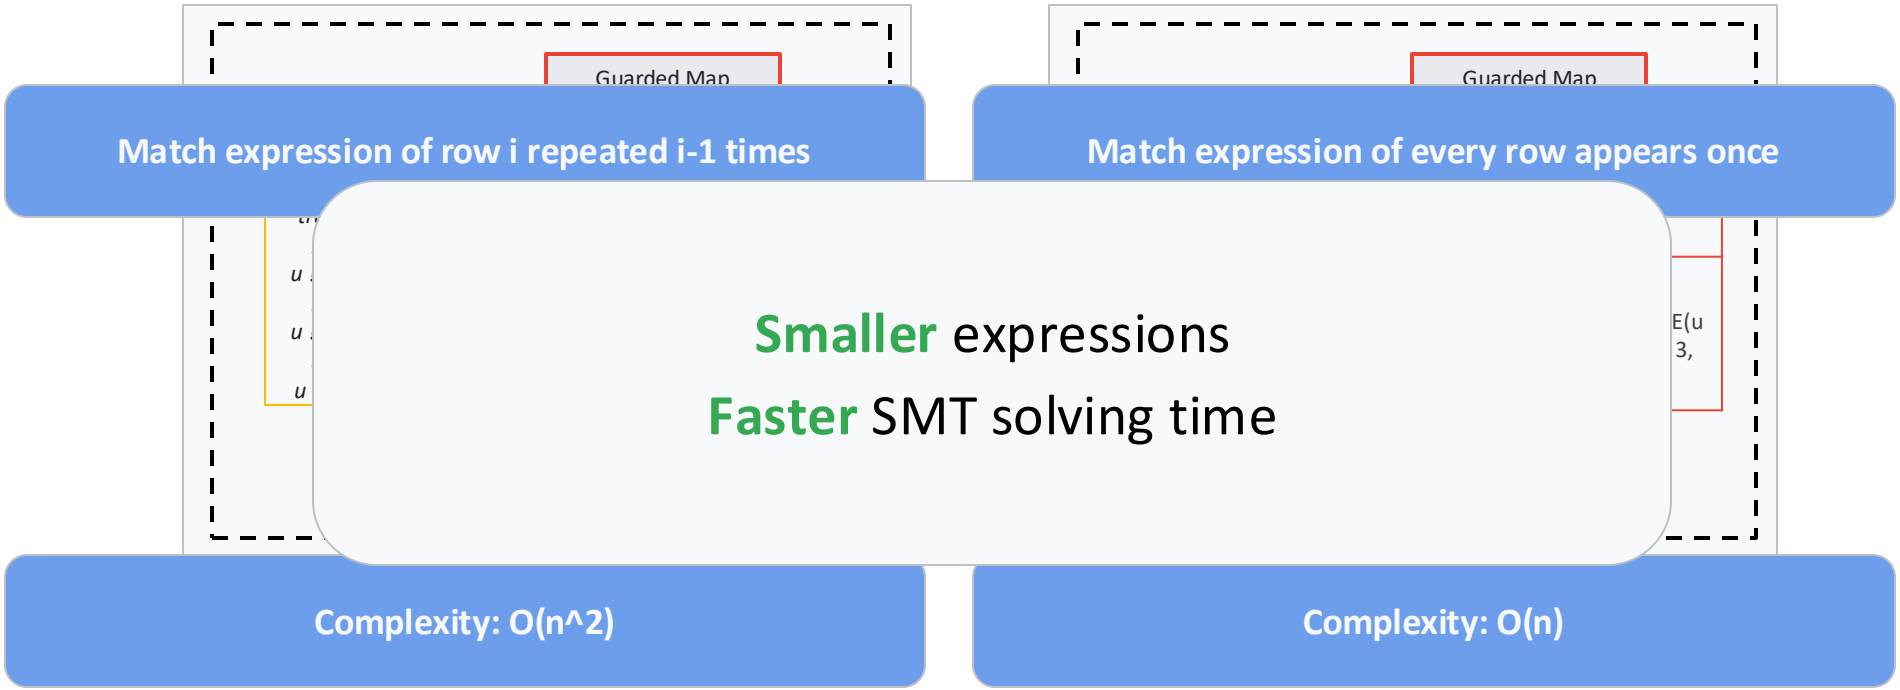

Results

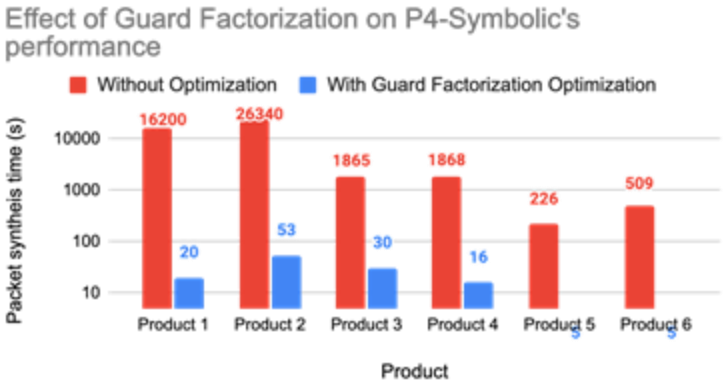

800x speedup!
(2.2hrs -> ~20sec)

Increased developer velocity

Helps testing of the switch better:

- **Test with larger snapshots**
- **Expand coverage goals:** Re-enable more coverage goals (e.g. header coverage)

Current status

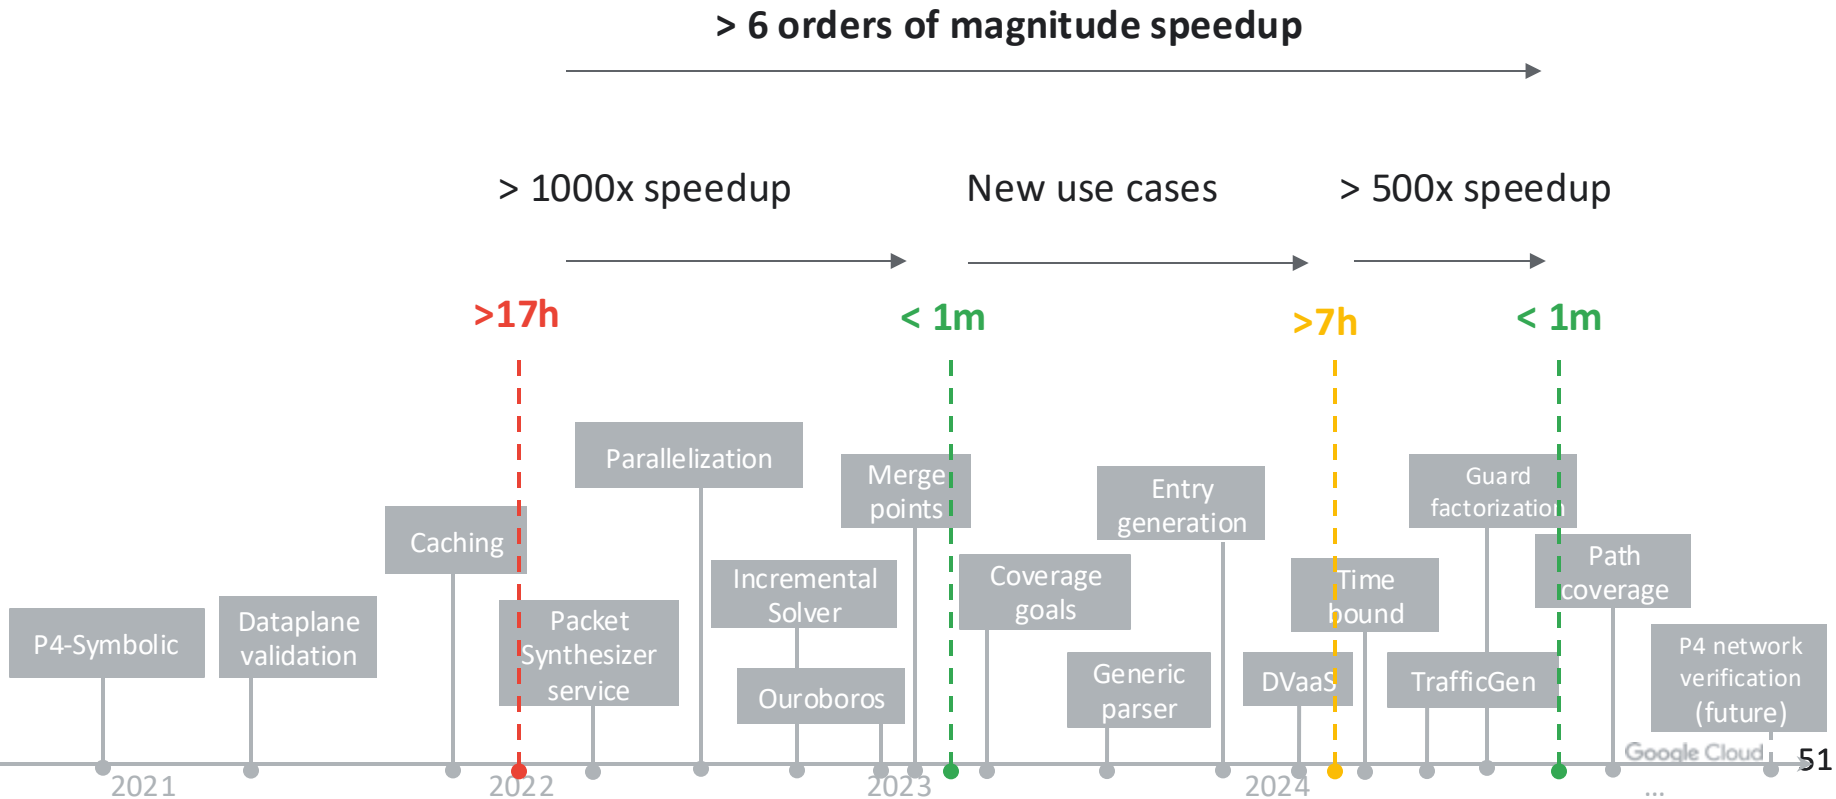

Outline

- Background and Context
- P4-Symbolic
- Performance Improvements
- Coverage Improvements
- Future

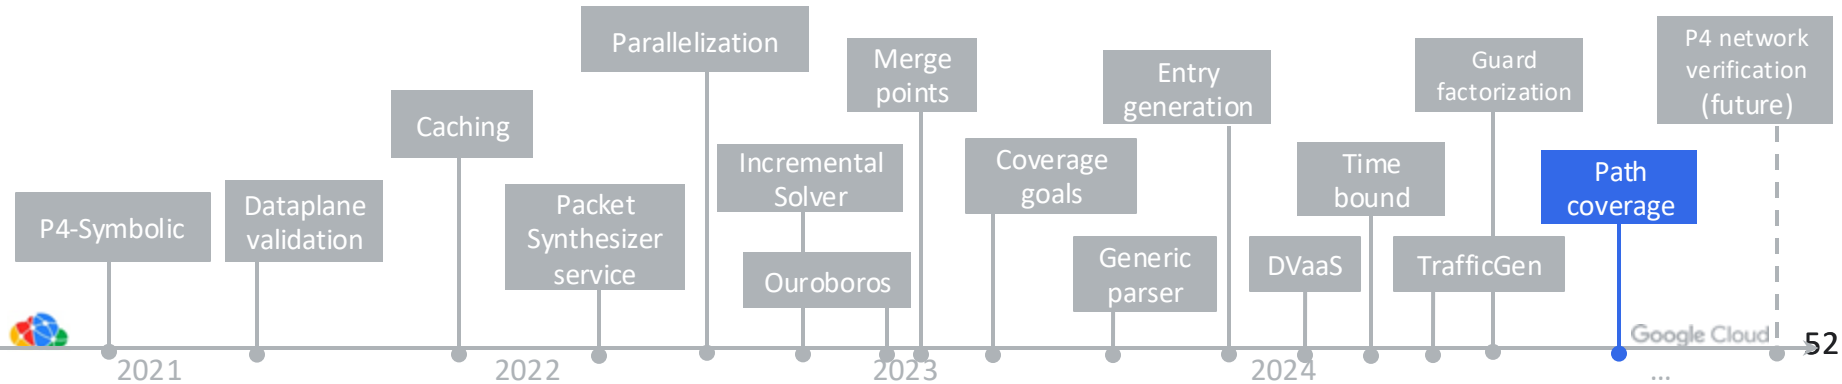

Coverage caveat

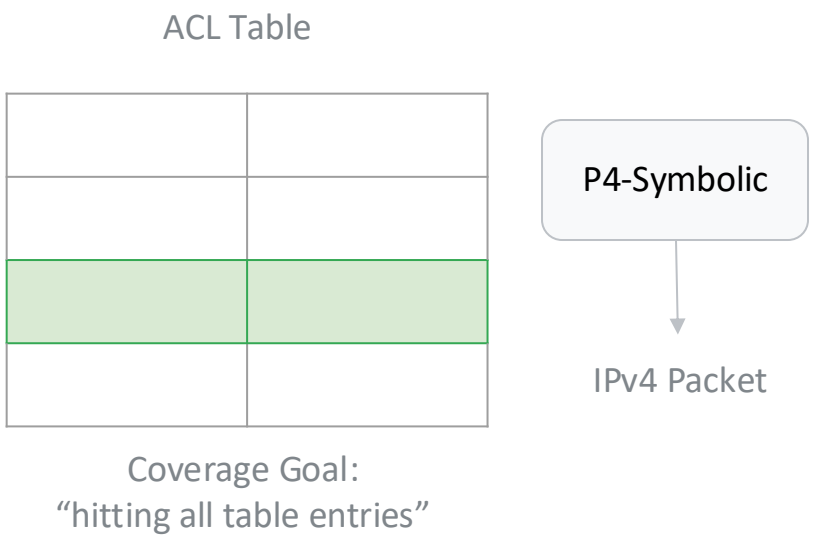

Coverage caveat

ACL Table

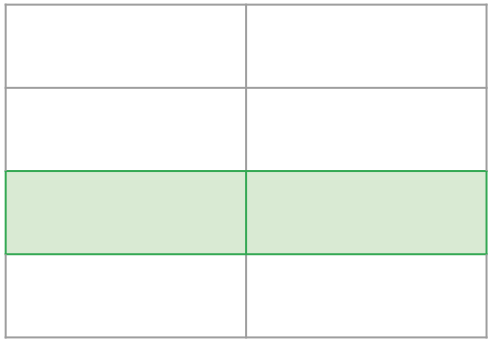

Coverage Goal:
"hitting all table entries"

P4-Symbolic

IPv4 Packet

Production Scenario

ACL Table

IPv6 Packet →  → **Failure**

Could miss the bug!

Nearly missed bug

Good packet: "IPv4 packet hitting ACL table"

Bad packet: "IPv6 packet hitting ACL table"

Coverage

*for **e** in entries:*
*generate a packet hitting **e***

Solution 1: Manually expand coverage goals!

*for **e** in entries:*
*for **h** in headers:*
*generate a packet containing . header **h** and hitting **e***

Add entry coverage, header coverage, ….. , and so on

- Cannot cover every case
- Very complex coverage goals -> more requests -> more time to solve -> slower packet synthesis

Solution 2: Path Coverage! *(Ultimate coverage)*

Path Coverage

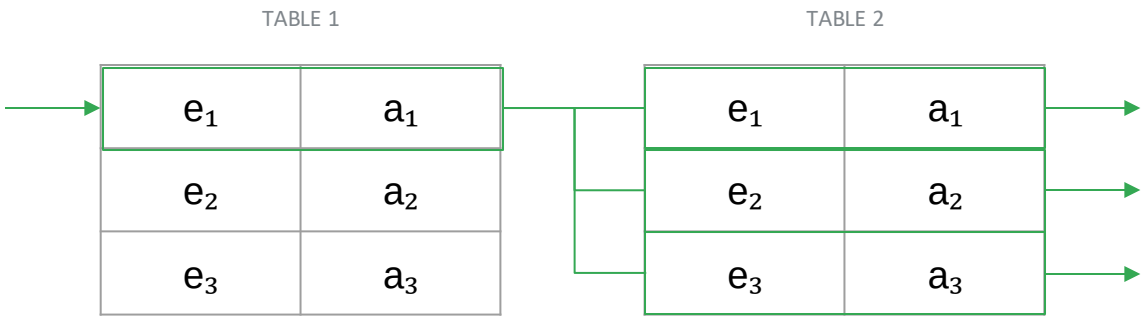

Path Coverage

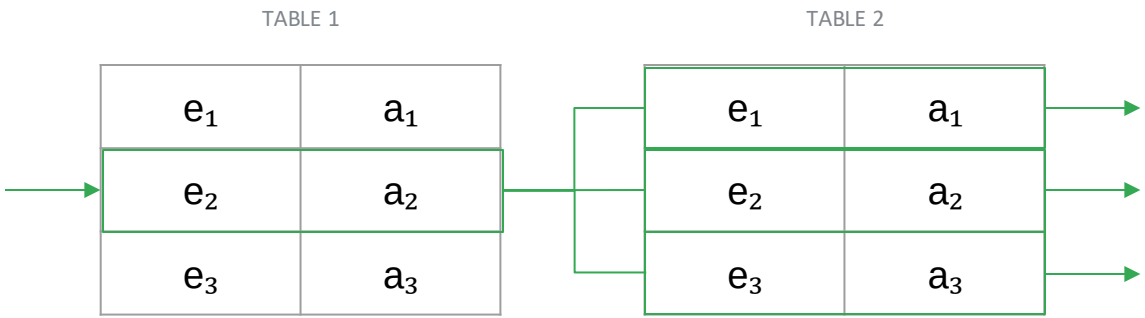

Path Coverage

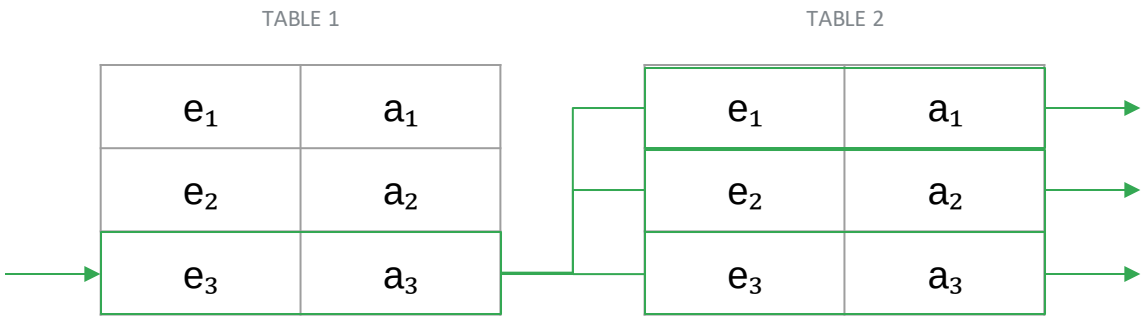

TOTAL PATHS COVERED: 9

Every possible scenario of a packet flow tested!

Path Coverage

Table 1: 1000 entries
Table 2: 1000 entries
Table 3: 1000 entries

Total Paths: **10^9 = 1B**
Covering all paths is **exponential**!

PROBLEM: **PATH EXPLOSION!**

Is there hope? - Yes

- Observation: Not all paths are valid
- Prune paths as you go!

Path Pruning

TABLE 1

Match	Action *egress_port*
e_1	5
e_2	50
e_3	10

TABLE 2

Match *egress_port*	Action
5	a_1
10	a_2
15	a_3

Path Pruning

Valid Path

TABLE 1

Match	Action *egress_port*
e_1	**5**
e_2	50
e_3	10

TABLE 2

Match *egress_port*	Action
5	a_1
10	a_2
15	a_3

**Valid Path -> An actual packet would take this path*

Path Pruning

Invalid Path

TABLE 1

Match	Action *egress_port*
e_1	**5**
e_2	50
e_3	10

TABLE 2

Match *egress_port*	Action
5	a_1
10	a_2

TOTAL PATHS: 9

VALID PATHS: 2

**Invalid Path -> An actual packet would not take this path*

Initial Results

Switch: Product 1 Clos Stage 2

#Paths: > 10^14

#Valid Paths: ~2M (~10^6)

Time taken: 8hrs

Initial Results

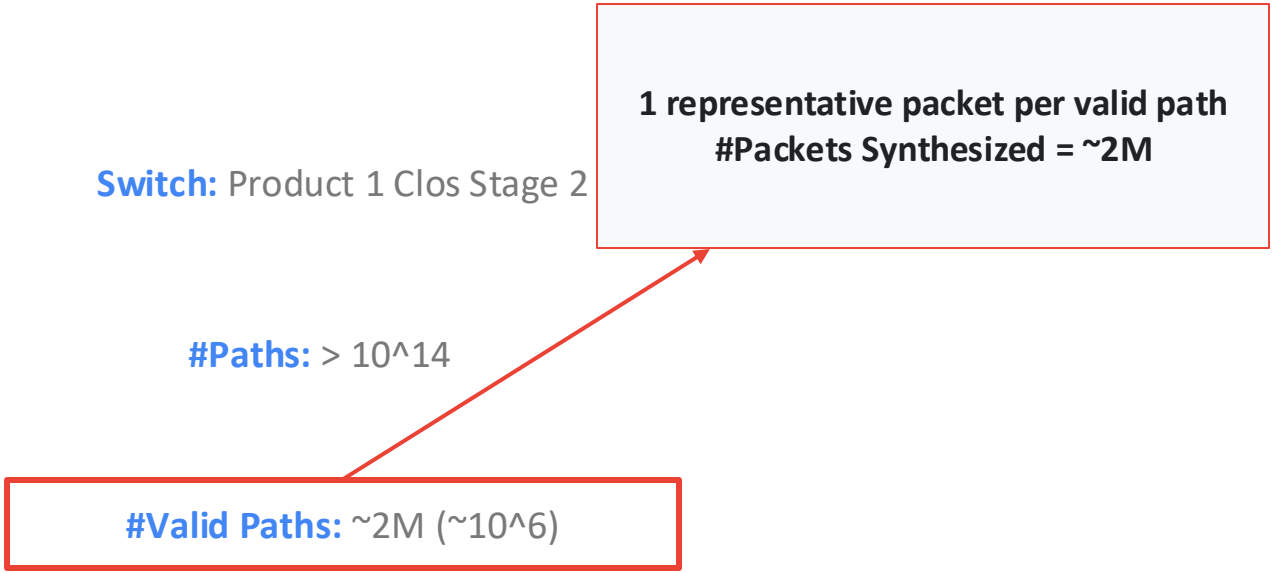

Can we do better?

Observation: Lesser calls to solver => faster execution times.

- Only 12.9% of calls to solver are satisfiable

Can we make fewer calls to solver somehow?

We plan to explore ideas from literature that address this problem

Outline

- Background and Context
- P4-Symbolic
- Performance Improvements
- Coverage Improvements
- Future

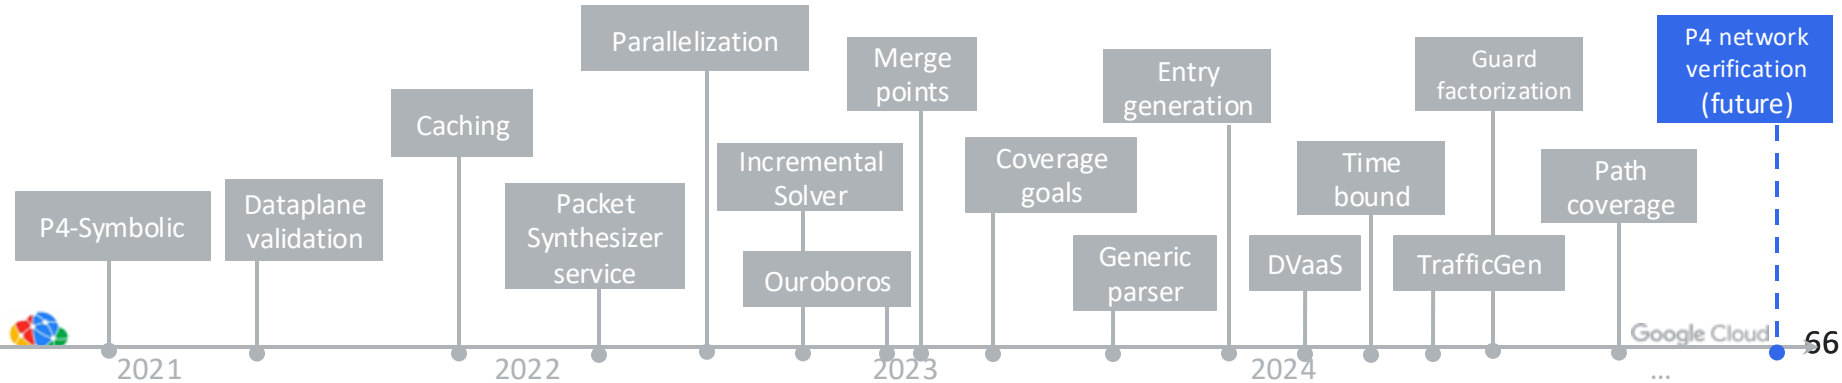

Overall goal: Ensure network works as intended

Subgoal 1: Ensure controller produces correct table entries (according to intents)

Subgoal 2: Ensure switch works as expected (according to table entries)

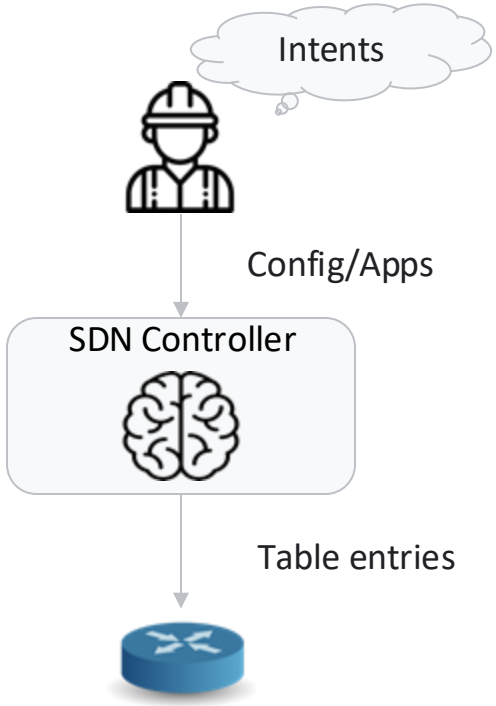

Network Verification

Existing system

- Hardcoded, incomplete model of switch
  - No guarantees on fidelity
  - Hard to evolve

In essence: symbolic execution at network level

Idea

- Extend P4 based symbolic execution to network level
  - Guaranteed high fidelity
    - Due to dataplane validation
  - Effortless evolution

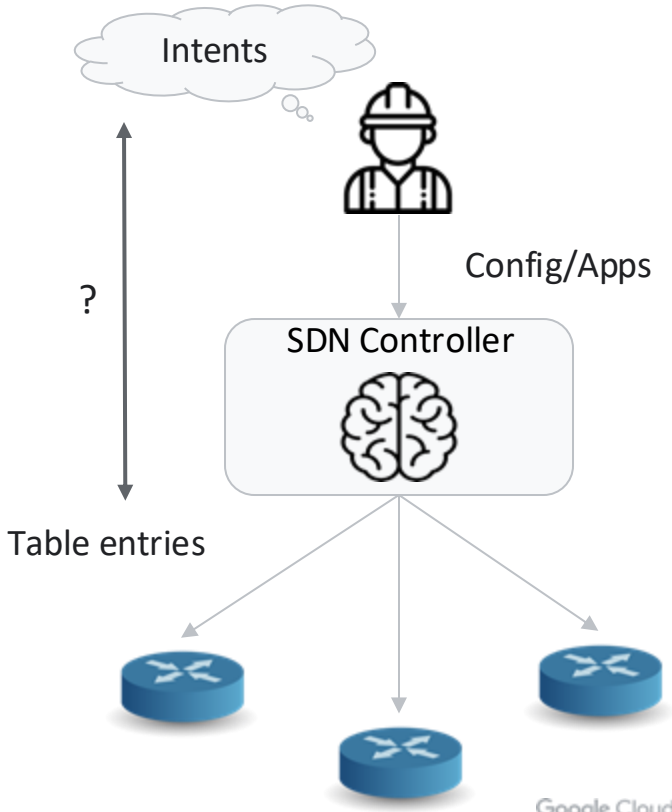

Thank you.

Google Cloud

Milestones and highlights

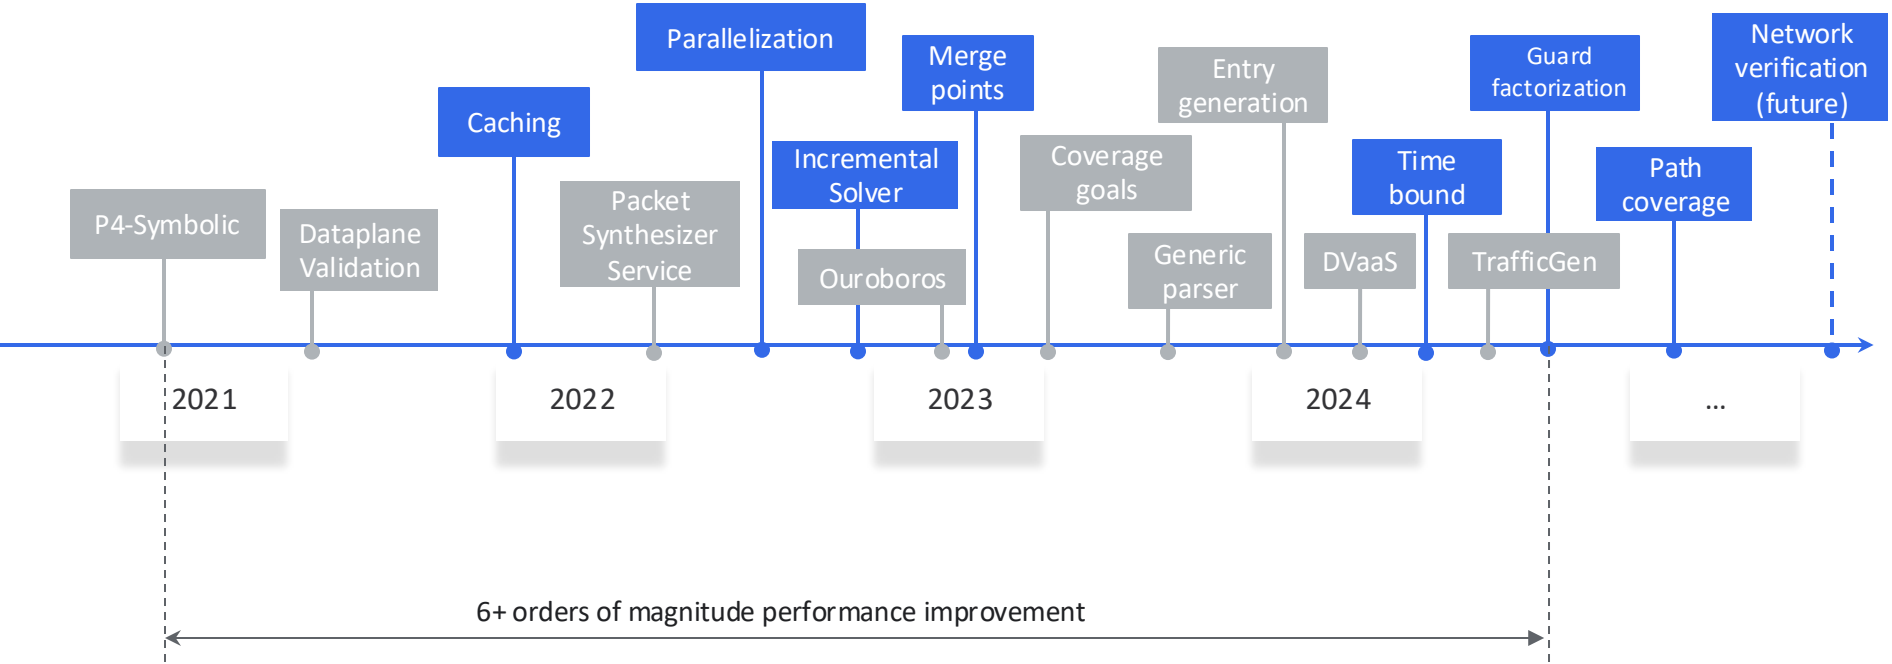

Relevant work

- HSA (NSDI'12), APV (ICNP'13), ddNF (HVC'16), #PEC (ICNP'19), etc.
  - Domain optimized "solvers" for network verification
  - Better performance, but more limitations (e.g. in packet rewrites)
- P4-Symbolic
  - Generic SMT solver (Z3)
  - More flexibility
  - Less performant
    - Good enough (for now)
    - Can employ ideas from above if needed