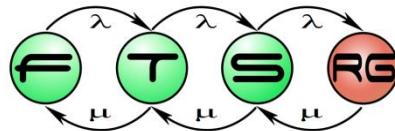


# Incremental Graph Queries with openCypher

Gábor Szárnyas

GraphDevroom @ FOSDEM 2017

Budapest University of Technology and Economics  
McGill University, Montréal



McGill

# Incremental Graph Queries with openCypher

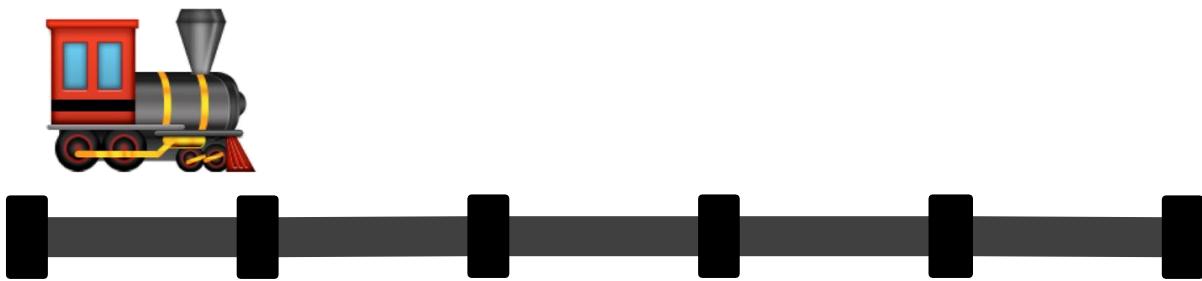
# Incremental Graph Queries

## with openCypher

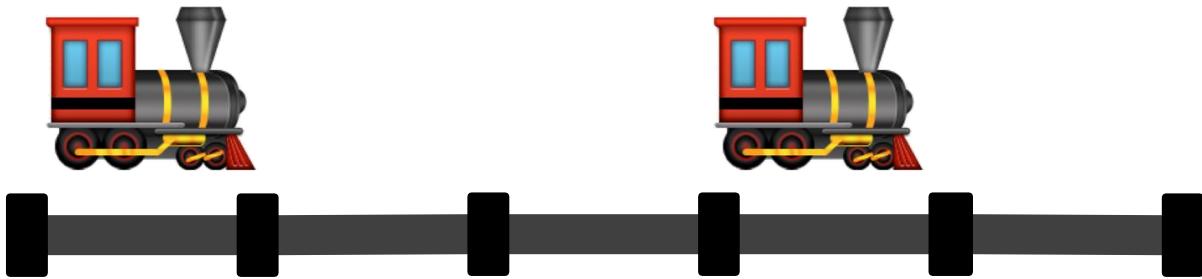
# Live railway model



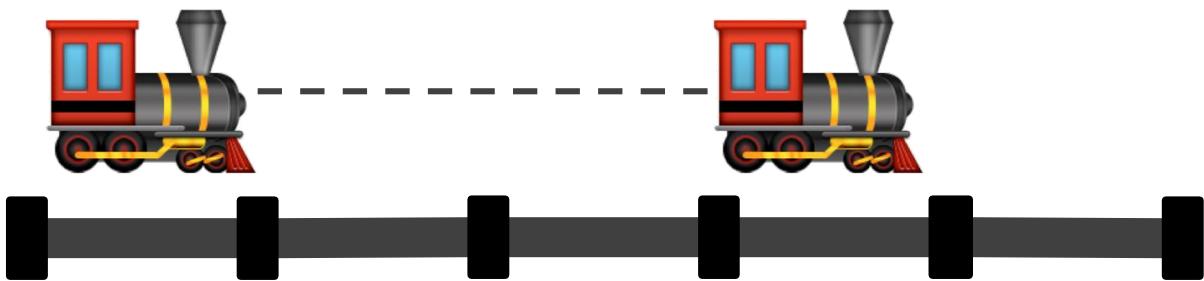
# Live railway model



# Live railway model

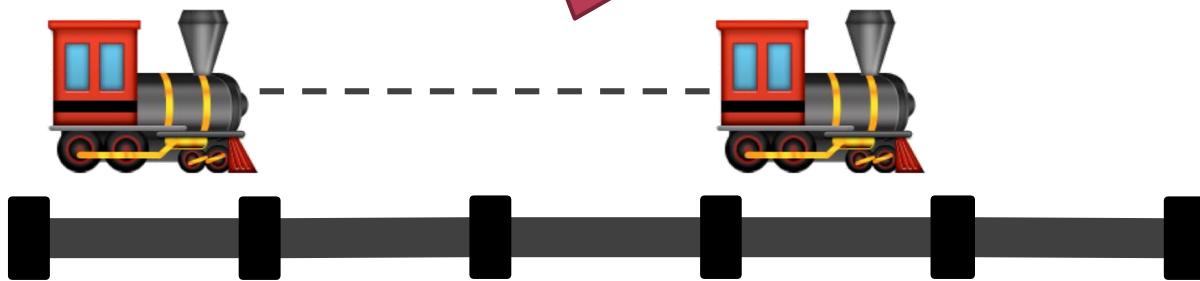


# Live railway model

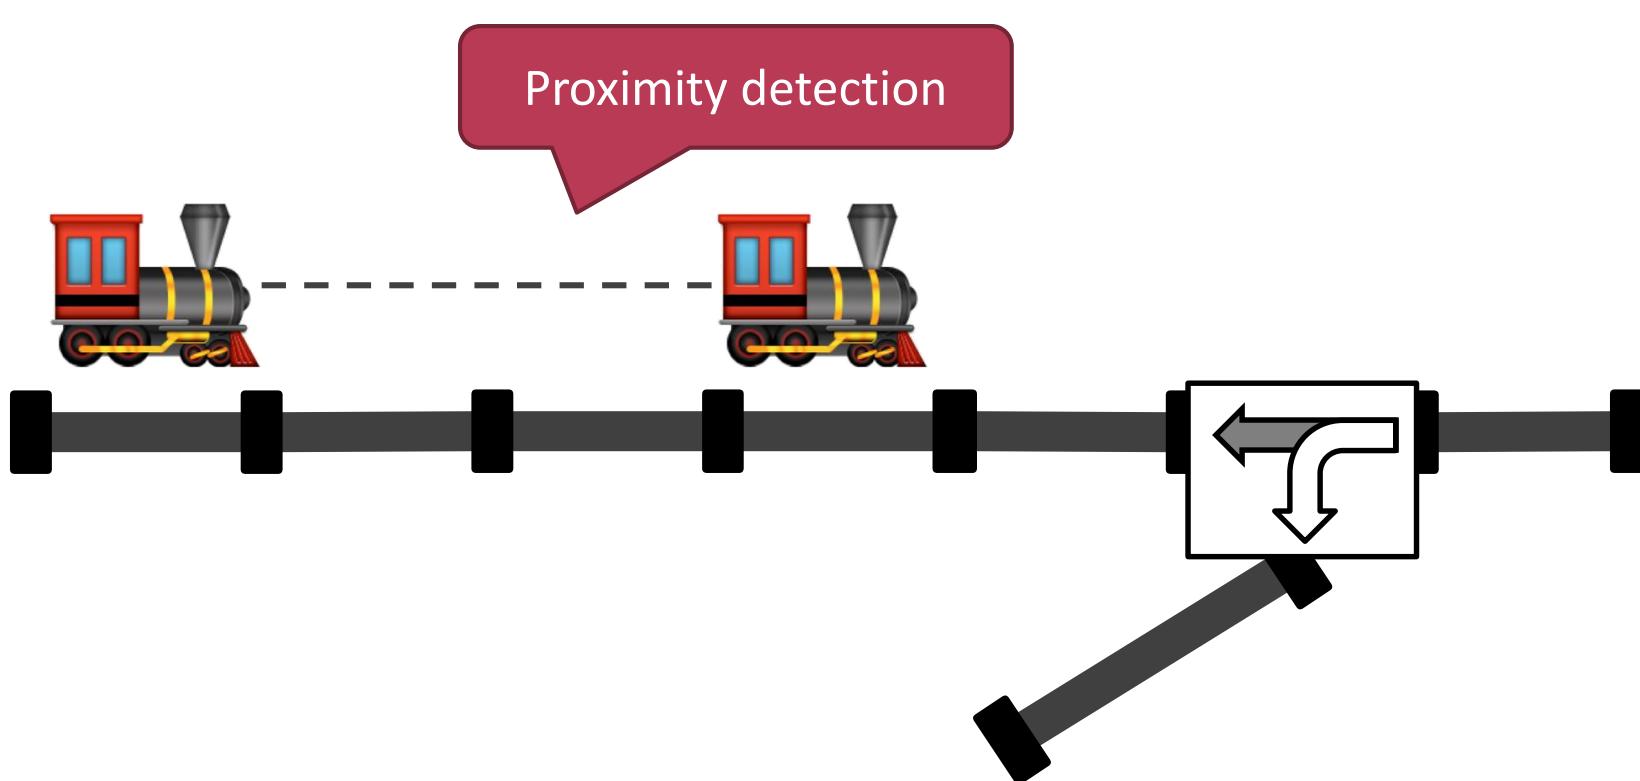


# Live railway model

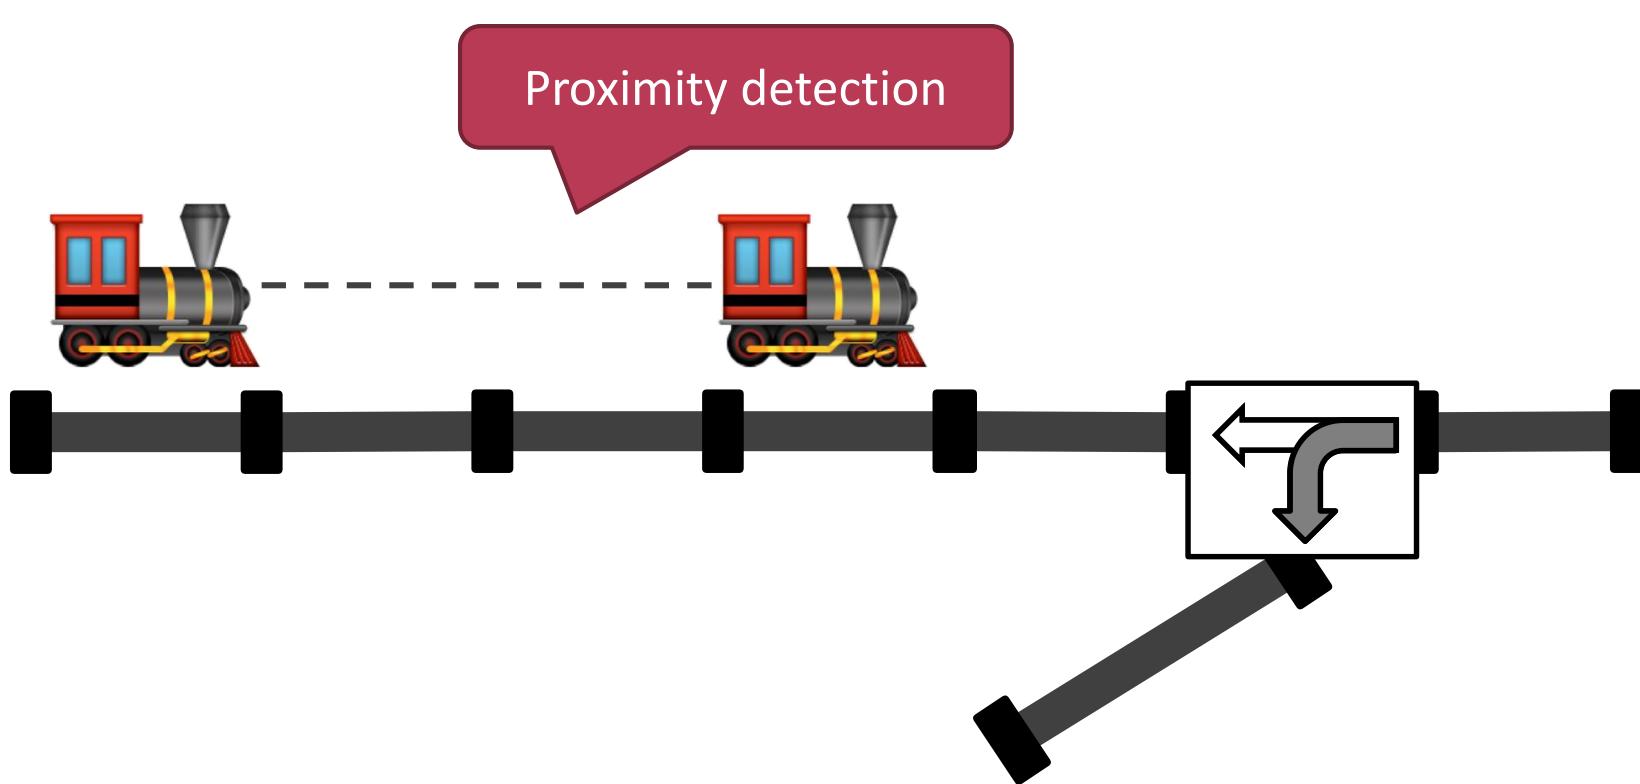
Proximity detection



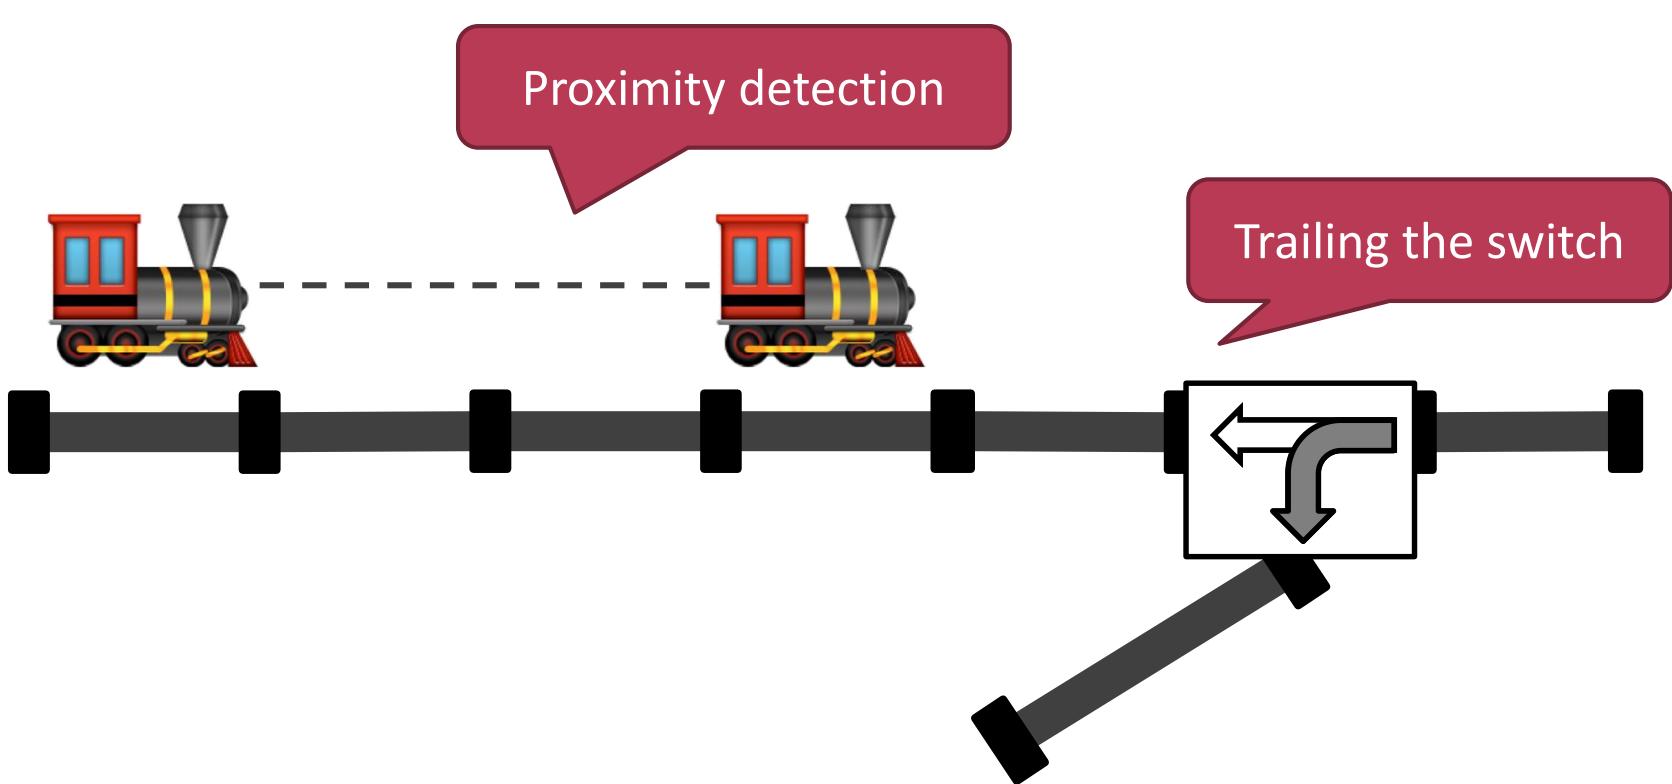
# Live railway model



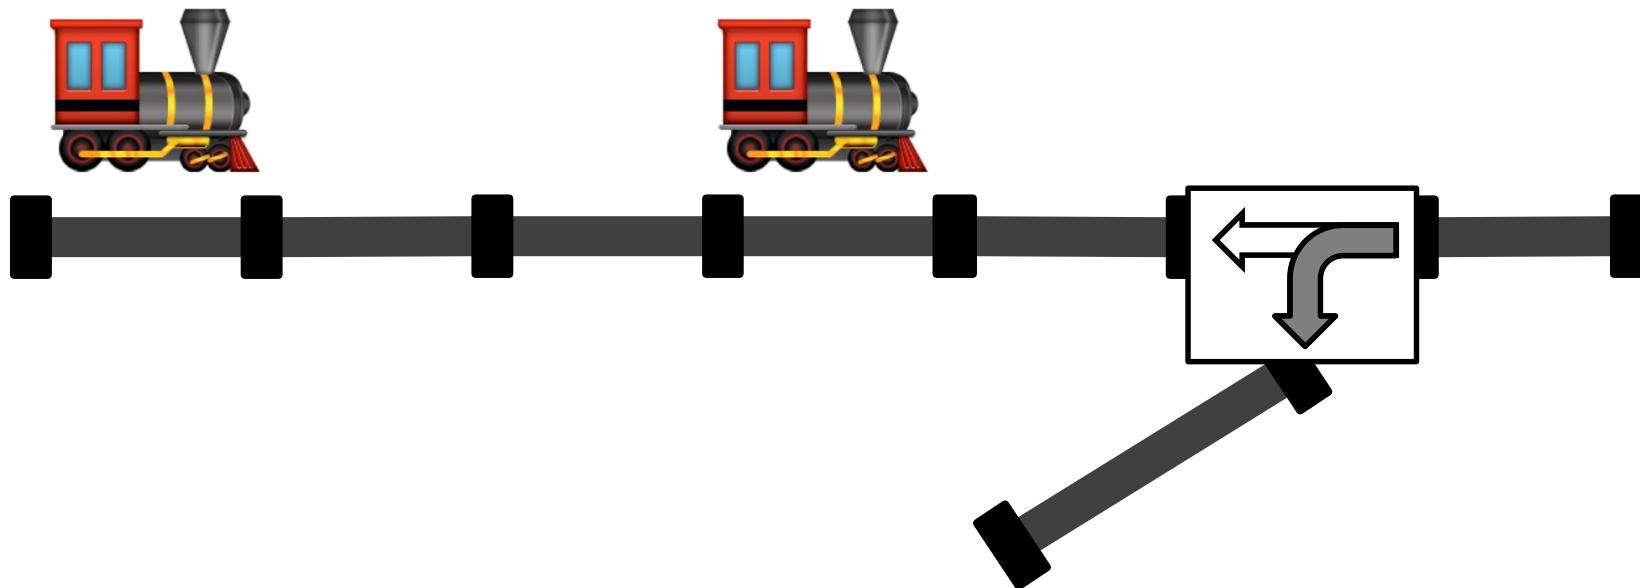
# Live railway model



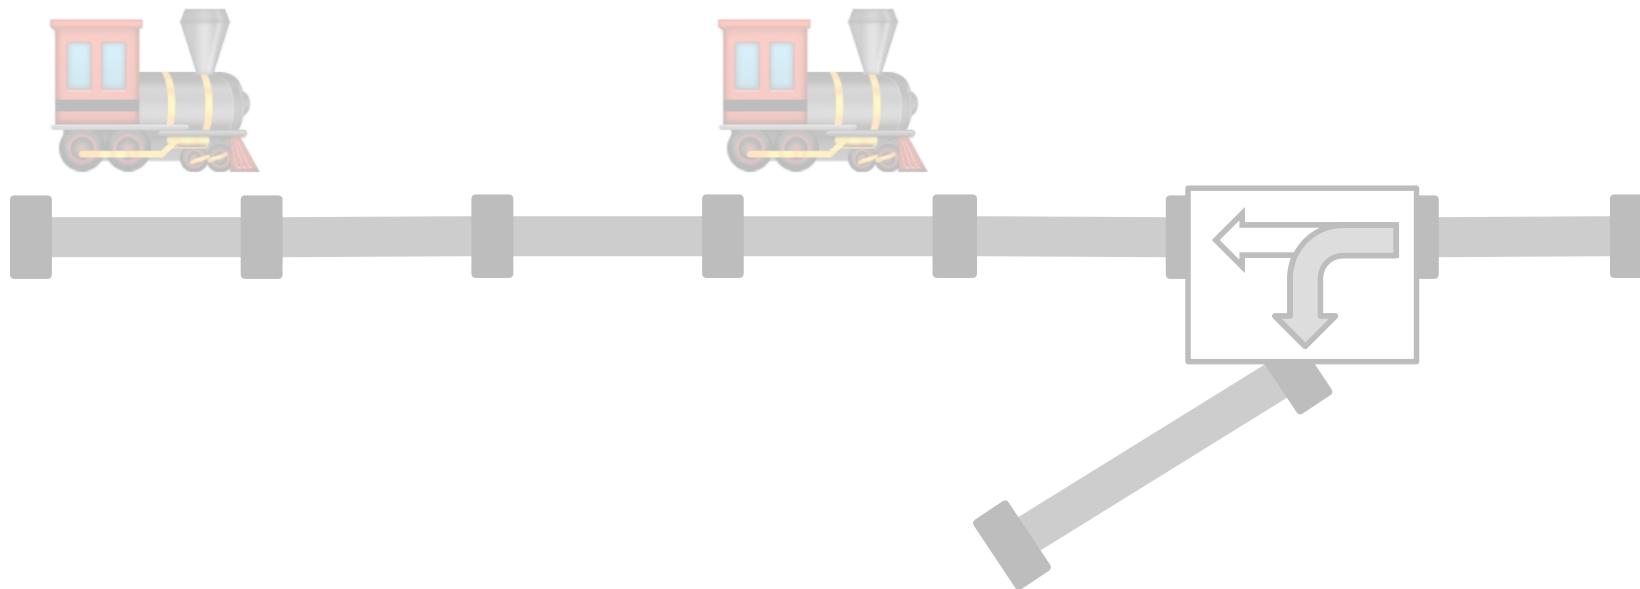
# Live railway model



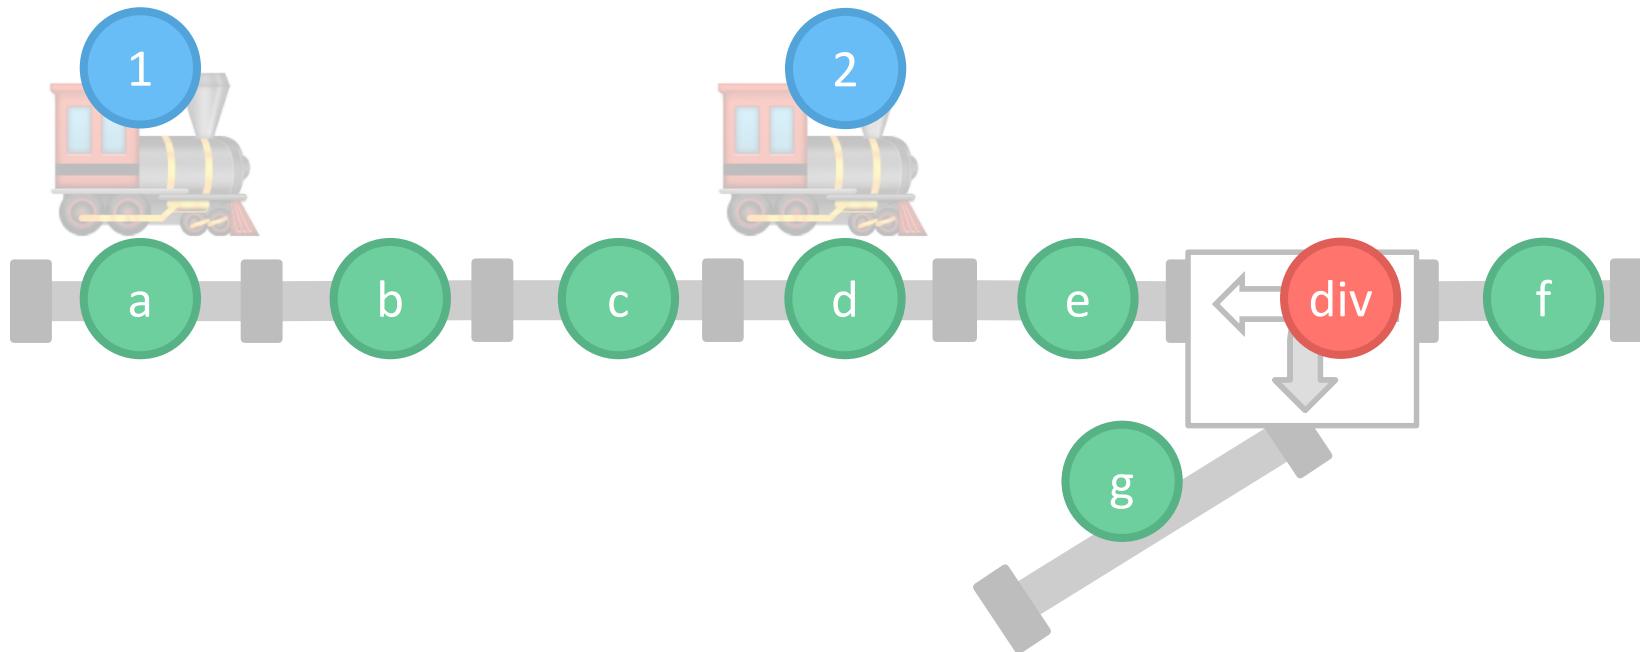
# Live railway model



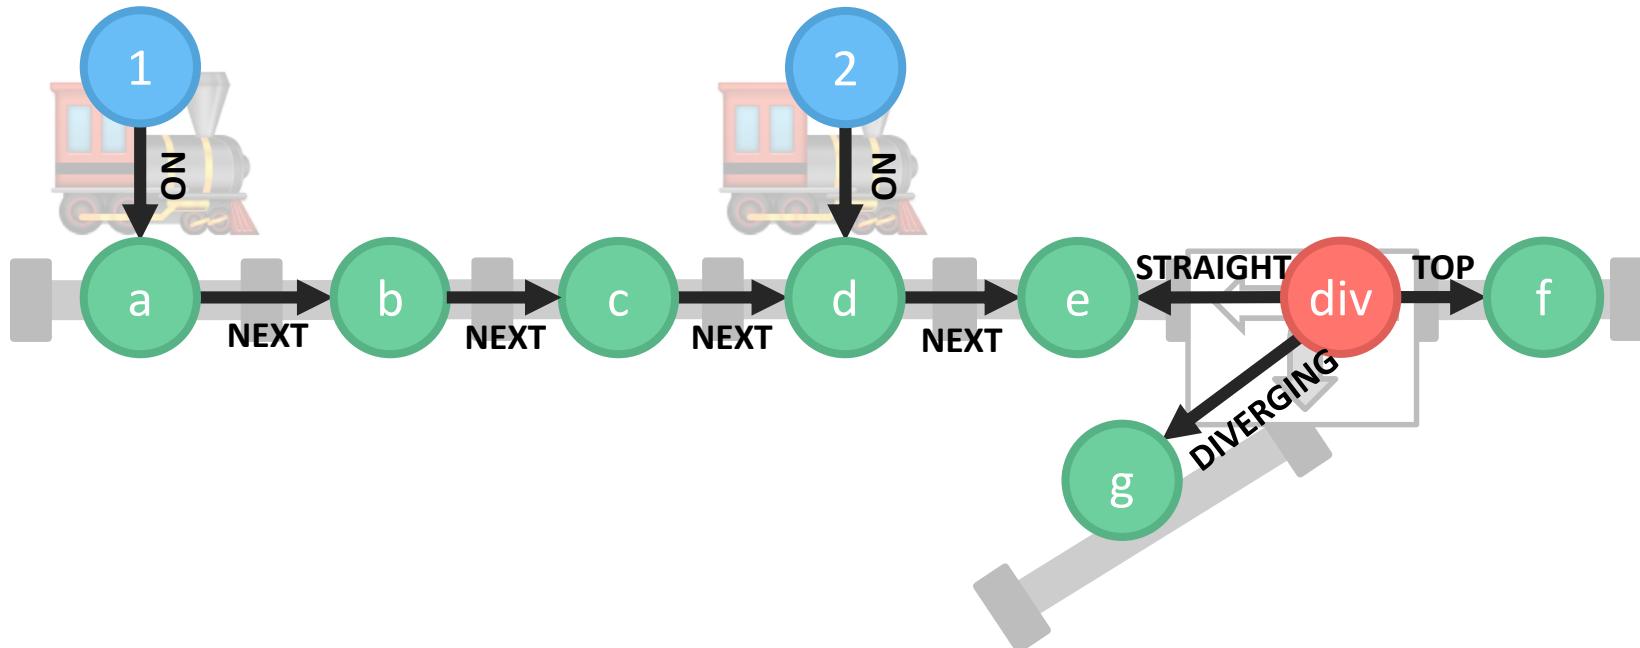
# Live railway model



# Live railway model

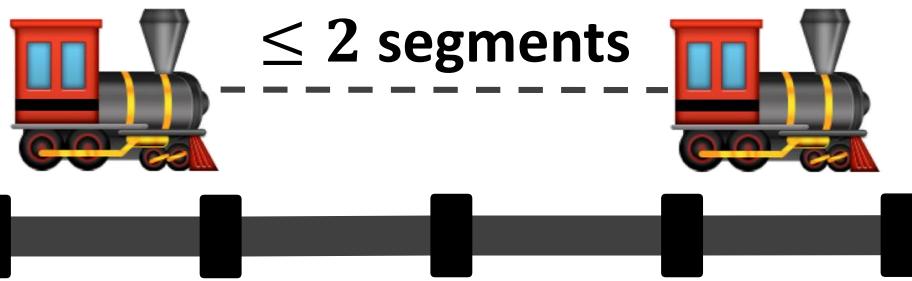


# Live railway model



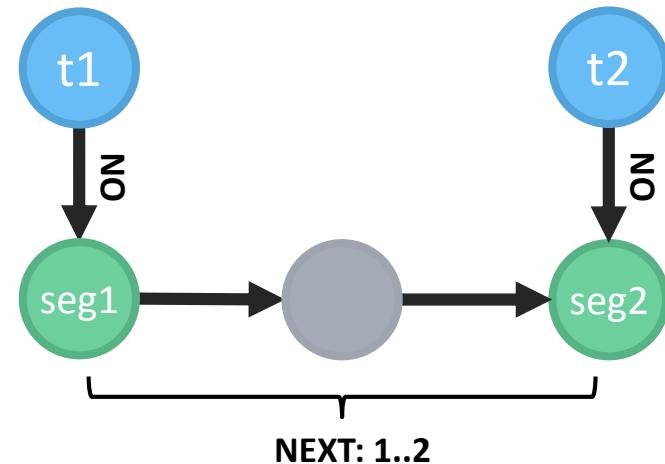
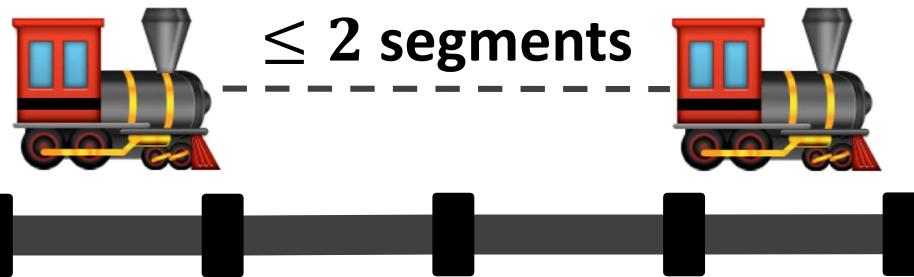
# Proximity detection

Proximity detection



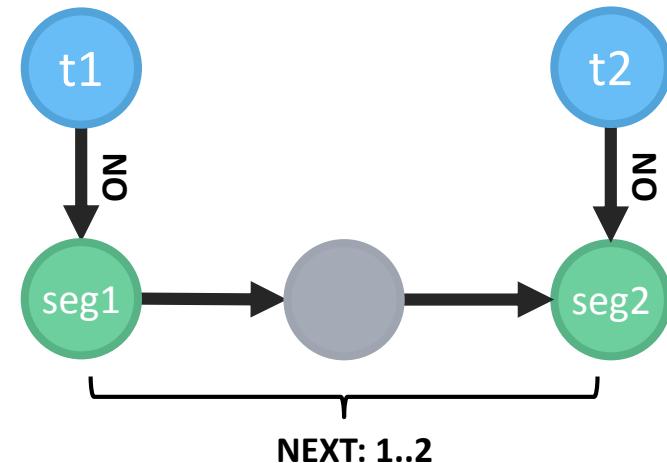
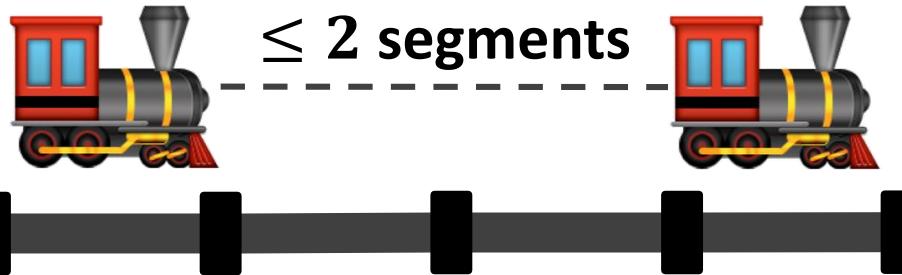
# Proximity detection

## Proximity detection



# Proximity detection

## Proximity detection



MATCH

$(t1:Train) - [:ON] -> (seg1:Segment)$

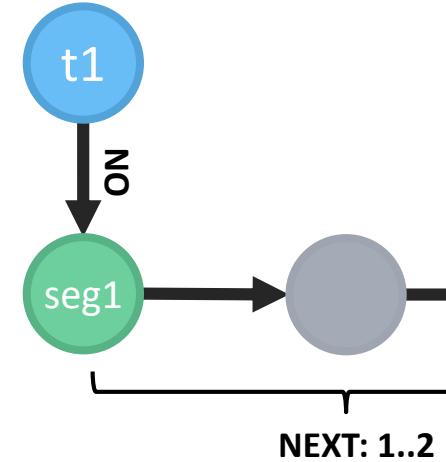
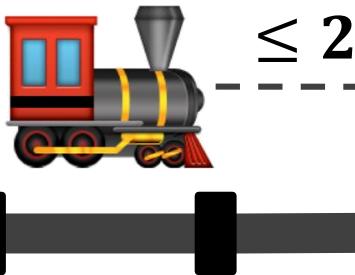
$- [:NEXT*1..2] -> (seg2:Segment)$

$<- [:ON] - (t2:Train)$

RETURN  $t_1, t_2, seg1, seg2$

# Proximity detection

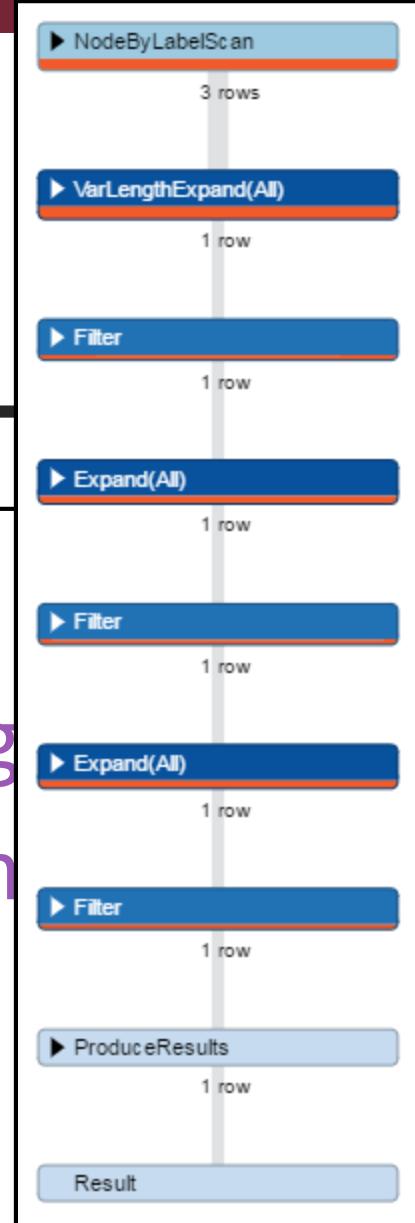
## Proximity detection



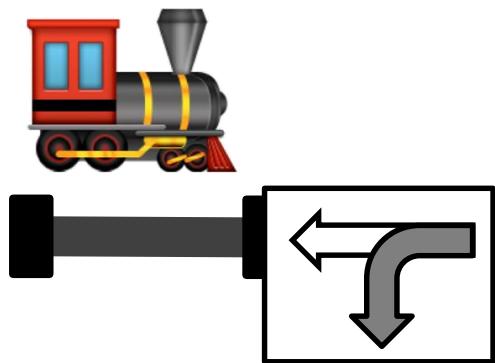
MATCH

```
(t1:Train)-[:ON]->(seg1:Segment)
-[:NEXT*1..2]->(seg2:Segment)
<-[:ON]-(t2:Train)
```

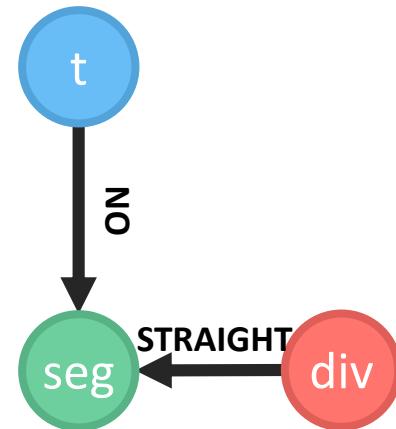
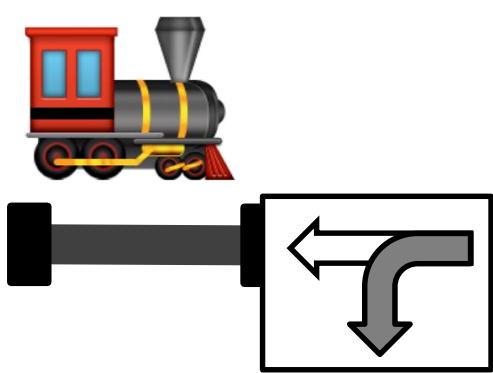
RETURN t1, t2, seg1, seg2



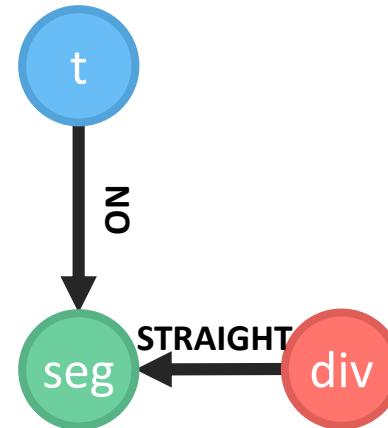
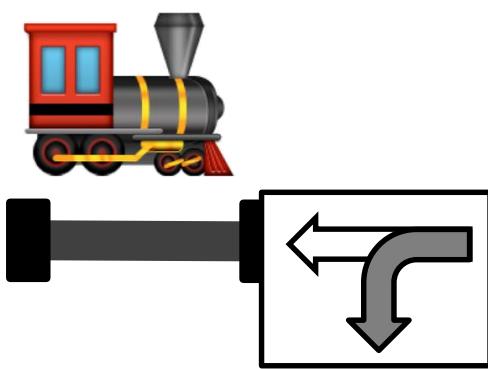
# Trailing the switch



# Trailing the switch

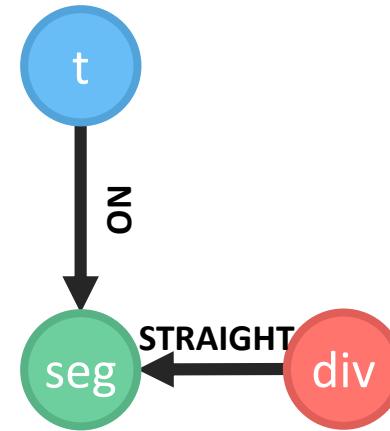
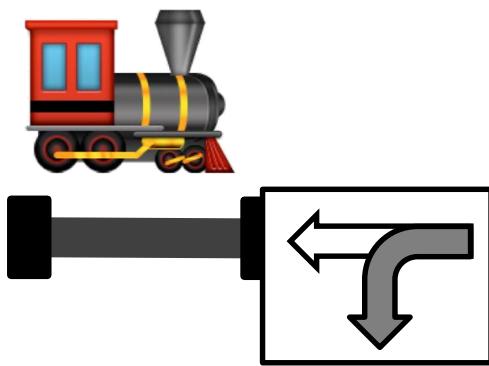


# Trailing the switch

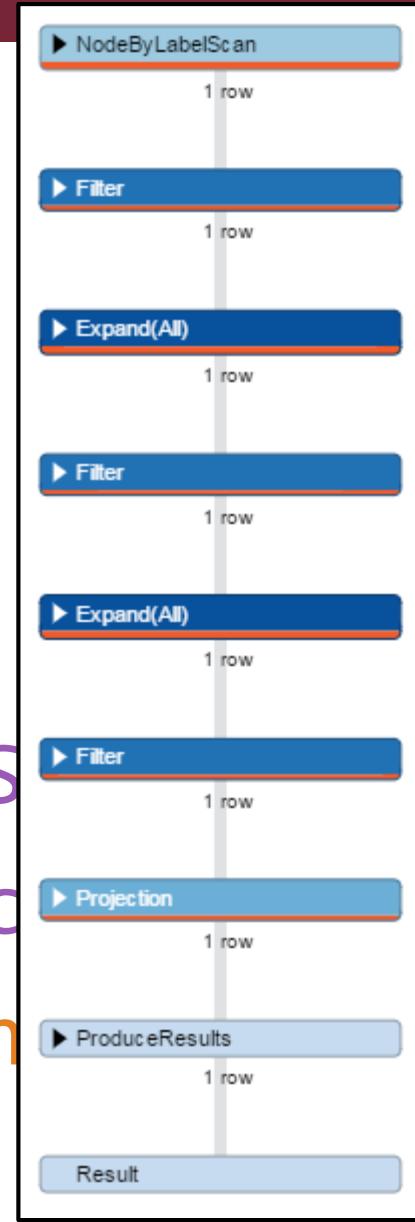


```
MATCH (t:Train)-[:ON]->(seg:Segment)
      <- [:STRAIGHT]- (sw:Switch)
WHERE sw.position = 'diverging'
RETURN t.number, sw
```

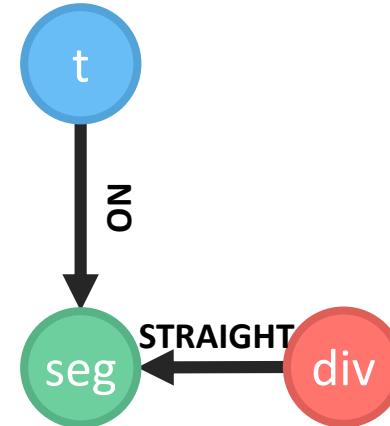
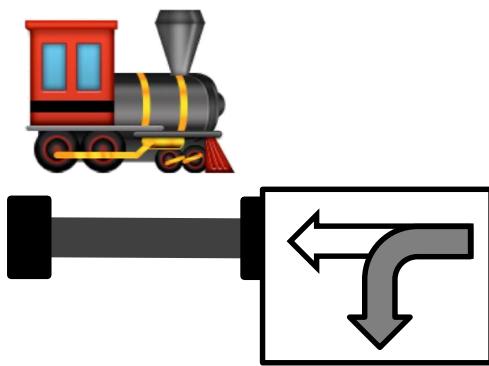
# Trailing the switch



```
MATCH (t:Train)-[:ON]->(seg:Straight)
      <- [:STRAIGHT]- (sw:Switch)
WHERE sw.position = 'diverging'
RETURN t.number, sw
```

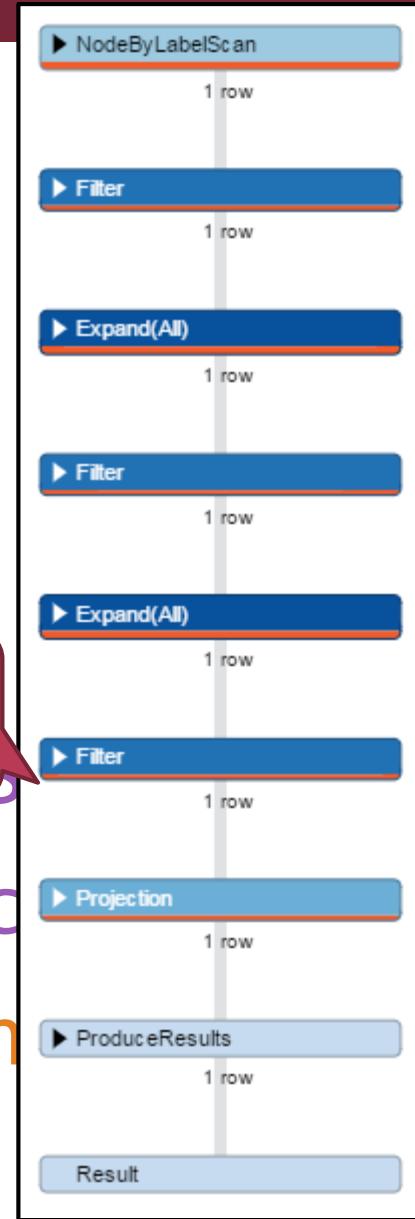


# Trailing the switch



Evaluate  
continuously

```
MATCH (t:Train)-[:ON]-(seg)  
    <- [:STRAIGHT]- (sw:Switch)  
WHERE sw.position = 'diverging'  
RETURN t.number, sw
```



# Incremental queries

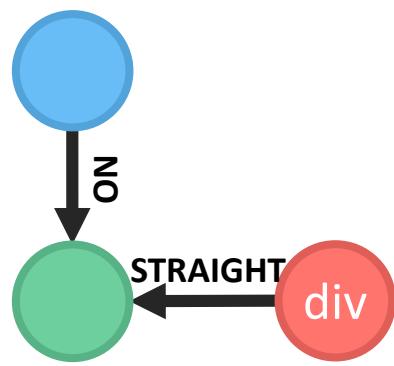
# Incremental queries

- Register a set of **standing queries**
- **Continuously evaluate** queries on changes

# Incremental queries

- Register a set of **standing queries**
- **Continuously evaluate** queries on changes
- The **Rete algorithm** (1974)
  - Originally for rule-based expert systems
  - Indexes the graph and caches interim query results

## Trailing the switch



$\pi_{t.number, sw}$

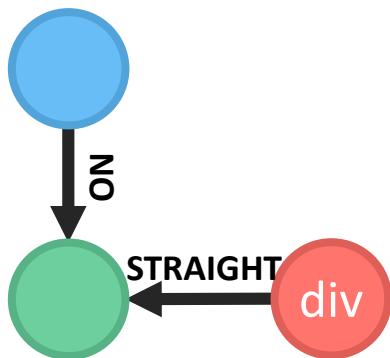
$\sigma_{sw.position = 'diverging'}$



ON

STRAIGHT

## Trailing the switch



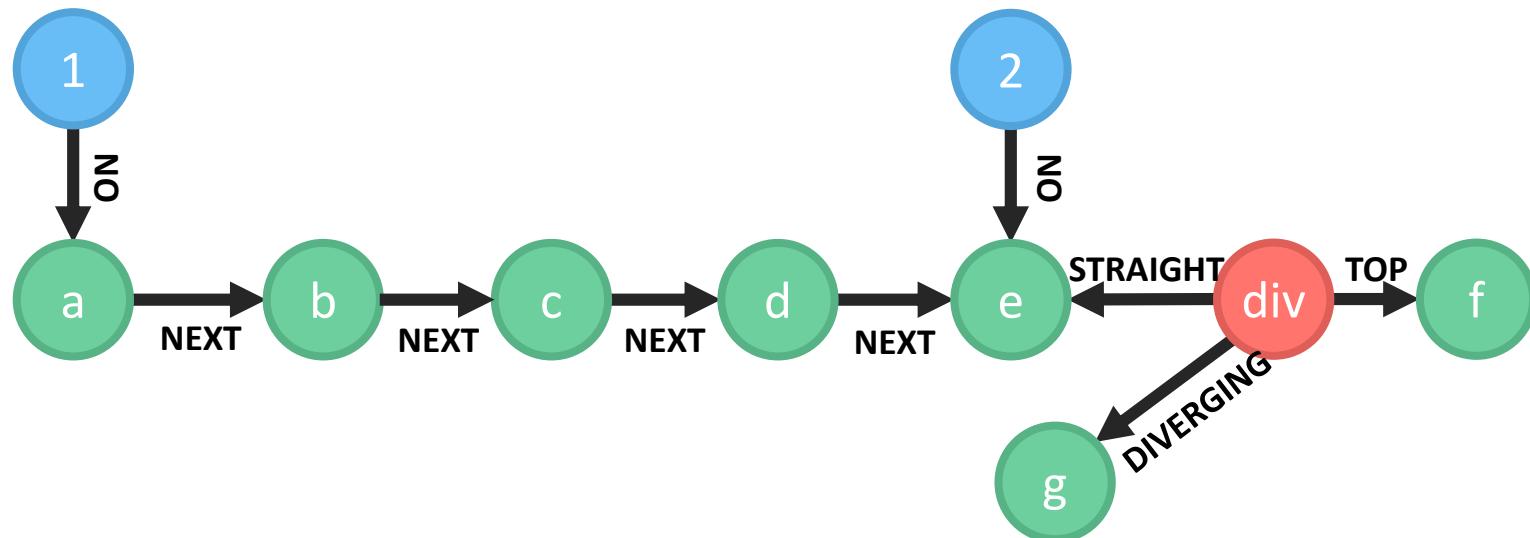
$\pi_{t.number, sw}$

$\sigma_{sw.position} = 'diverging'$

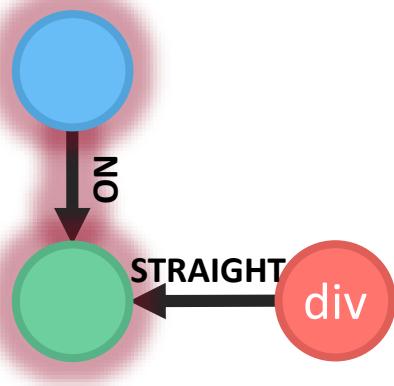


ON

STRAIGHT



## Trailing the switch



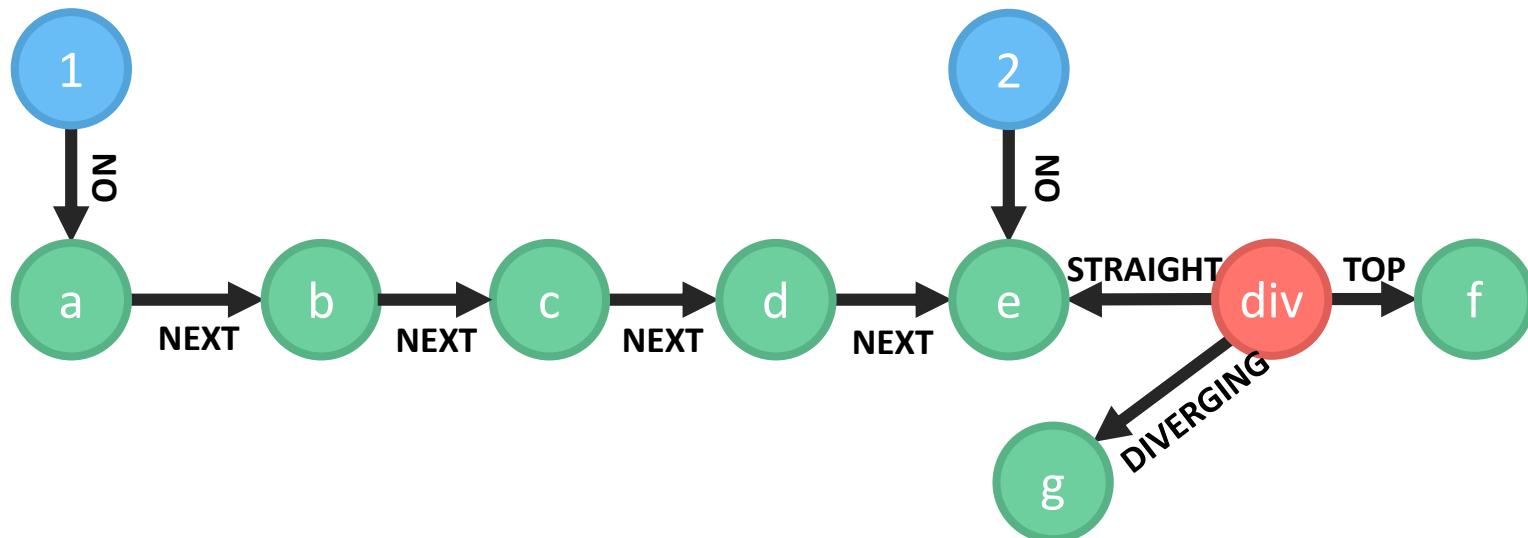
$\pi_{t.number, sw}$

$\sigma_{sw.position} = 'diverging'$

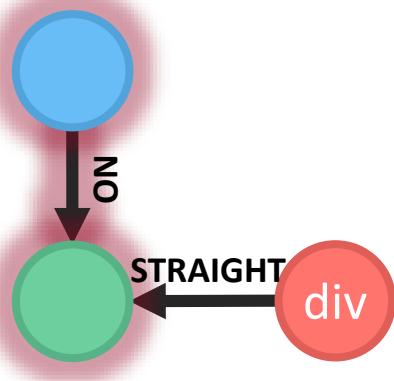
☒

ON

STRAIGHT



## Trailing the switch



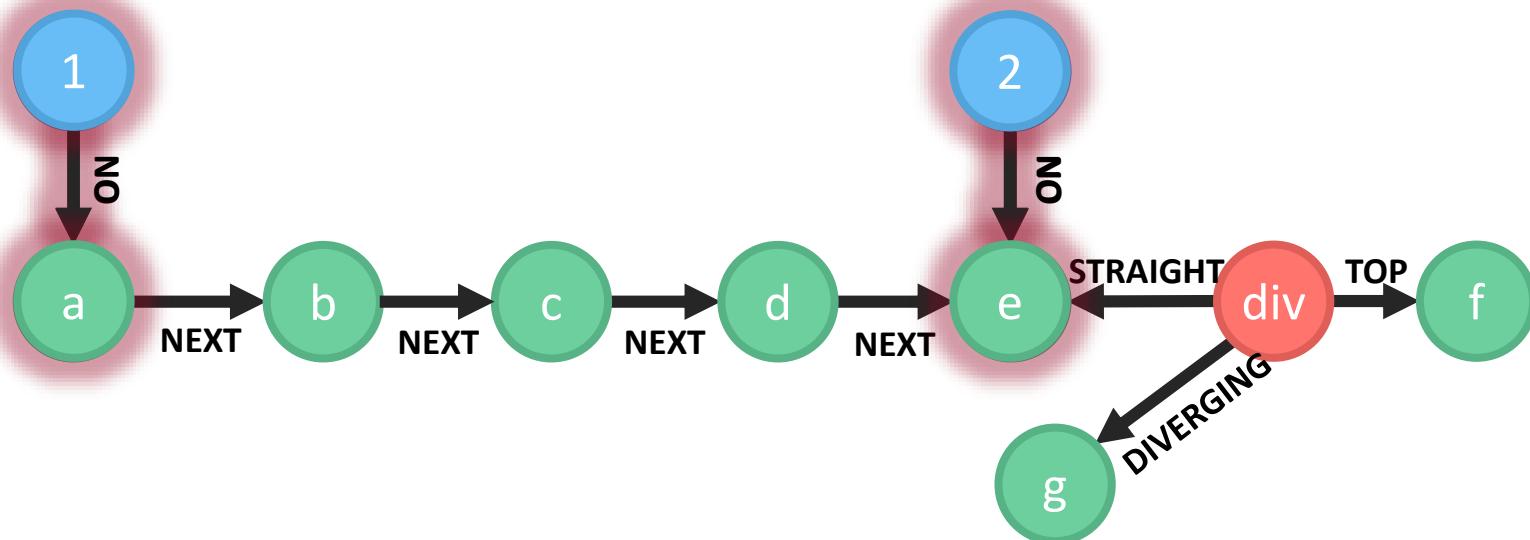
$\pi_{t.number, sw}$

$\sigma_{sw.position} = 'diverging'$

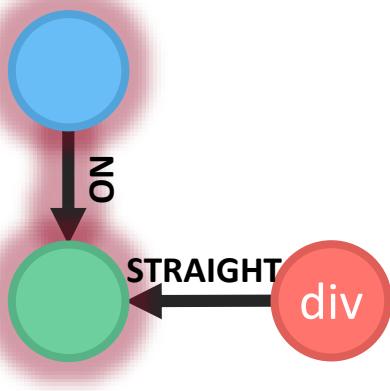
⊗

ON

STRAIGHT



## Trailing the switch



$\pi_{t.number, sw}$

$\sigma_{sw.position} = 'diverging'$

$\bowtie$

ON

1

ON

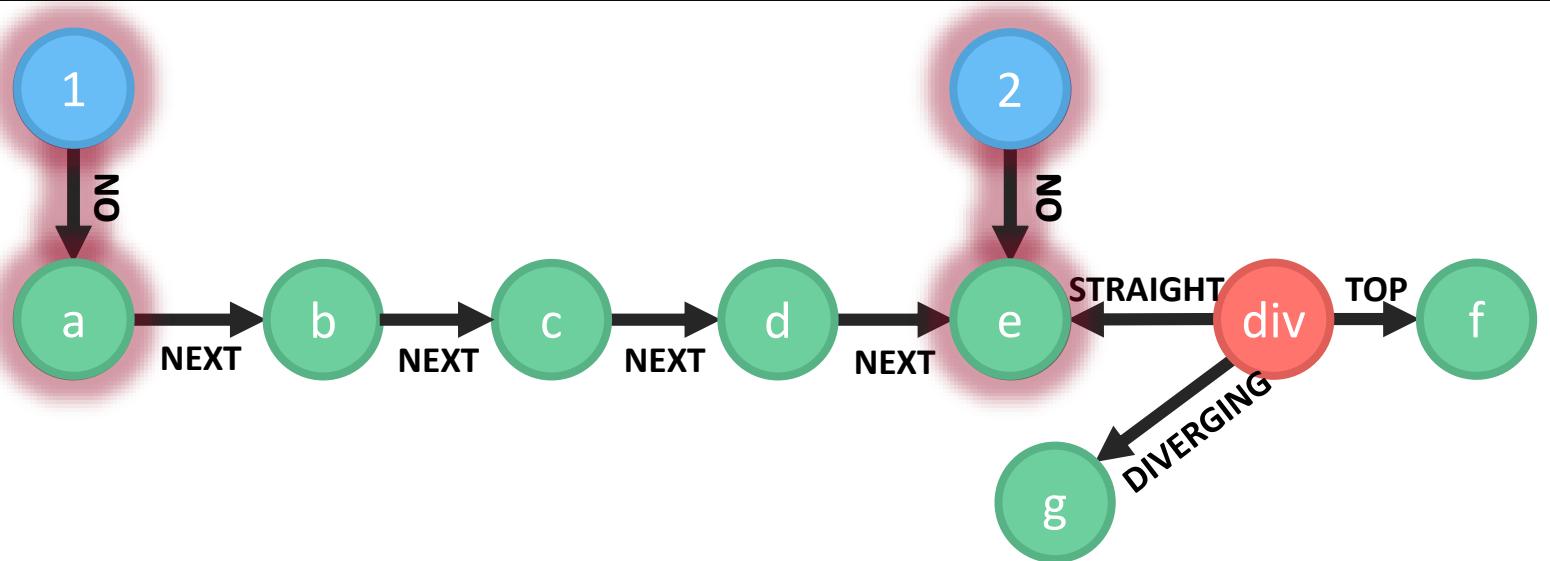
a

2

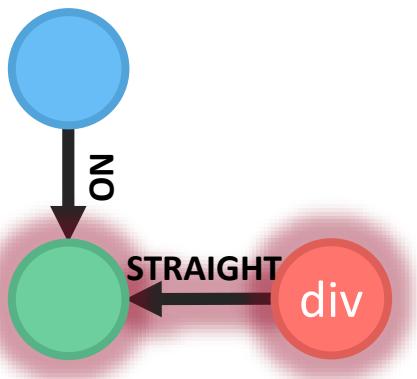
ON

e

STRAIGHT



## Trailing the switch



$\pi_{t.number, sw}$

$\sigma_{sw.position} = 'diverging'$

⊗

ON

1

ON

a

2

ON

e

STRAIGHT

1

ON

a

b

c

d

2

ON

e

g

DIVERGING

TOP

STRAIGHT

div

f

NEXT

NEXT

NEXT

NEXT

NEXT

g

DIVERGING

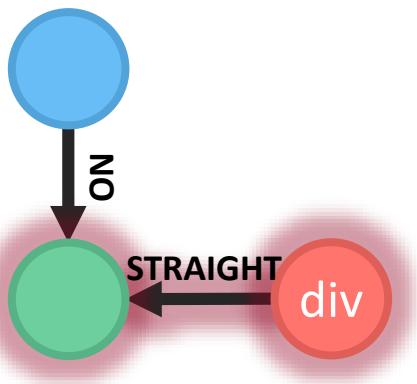
TOP

STRAIGHT

div

f

## Trailing the switch



$\pi_{t.number, sw}$

$\sigma_{sw.position} = 'diverging'$

$\bowtie$

ON

1

ON

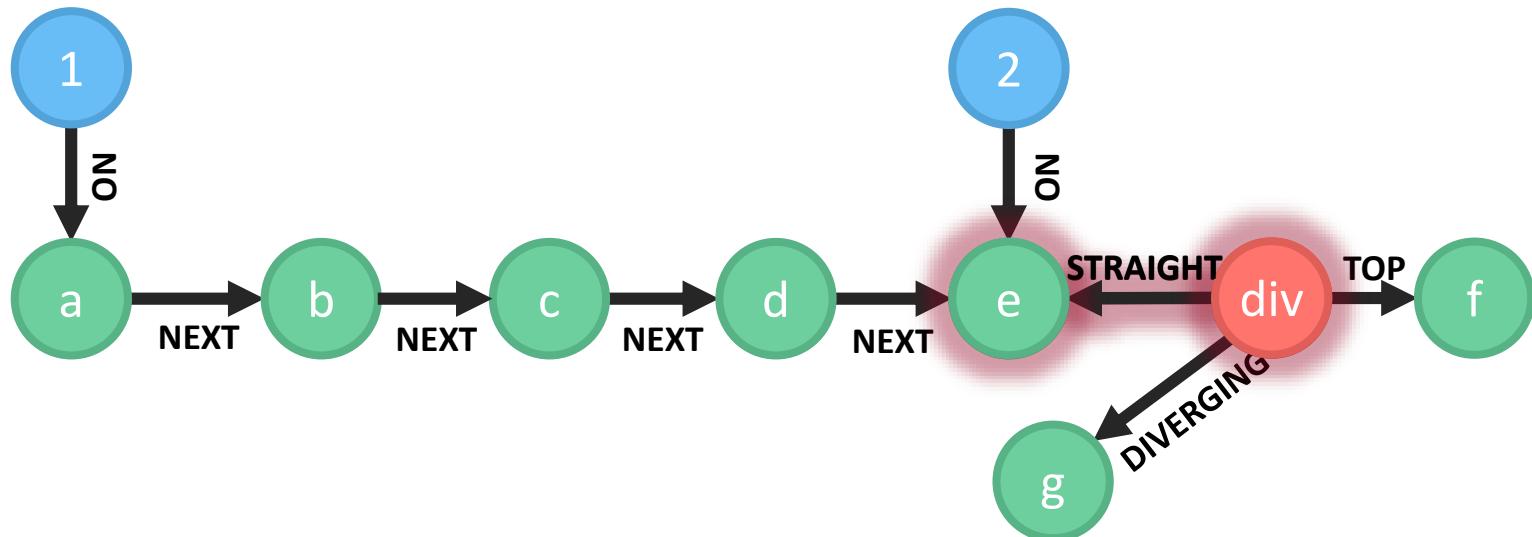
a

2

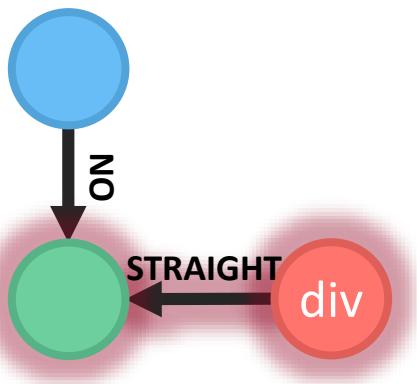
ON

e

STRAIGHT



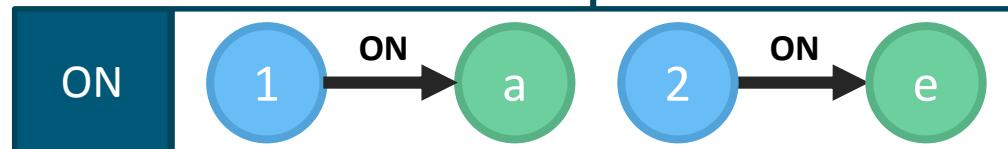
## Trailing the switch



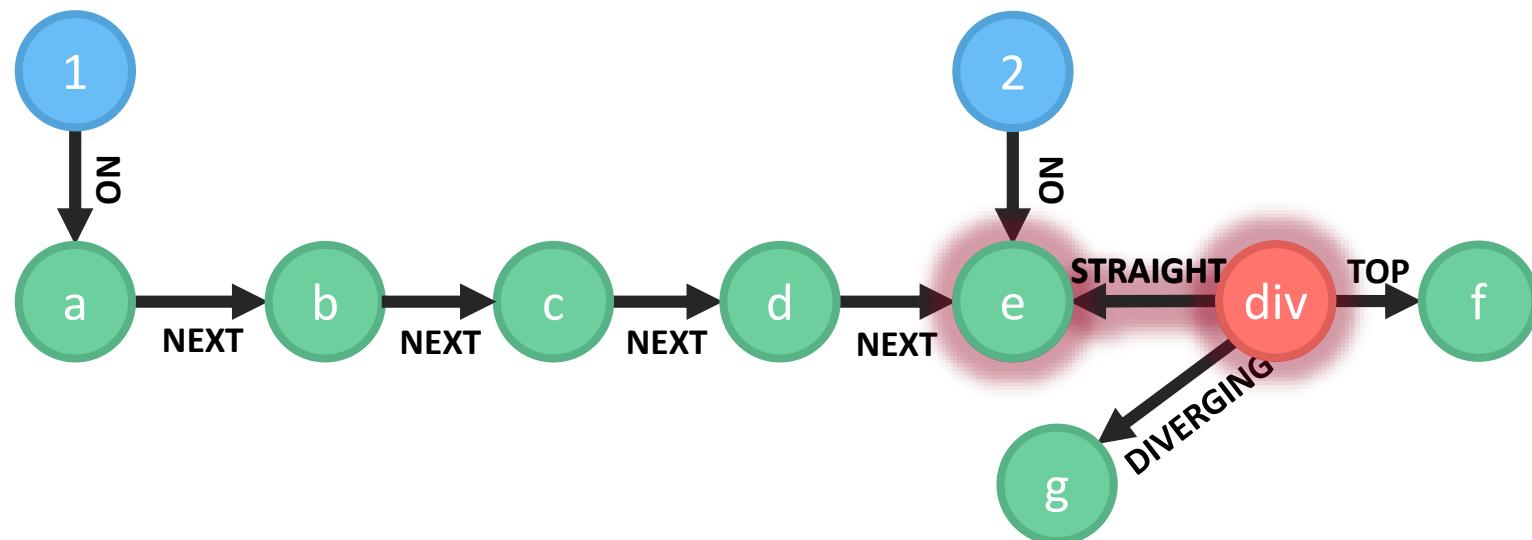
$\pi_{t.number, sw}$

$\sigma_{sw.position} = 'diverging'$

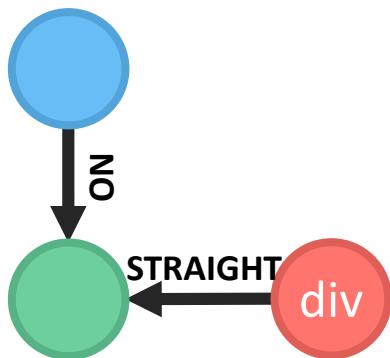
⊗



STRAIGHT

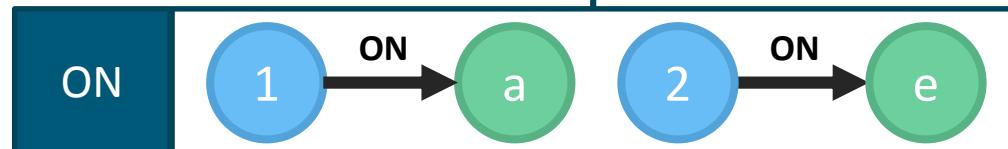


## Trailing the switch

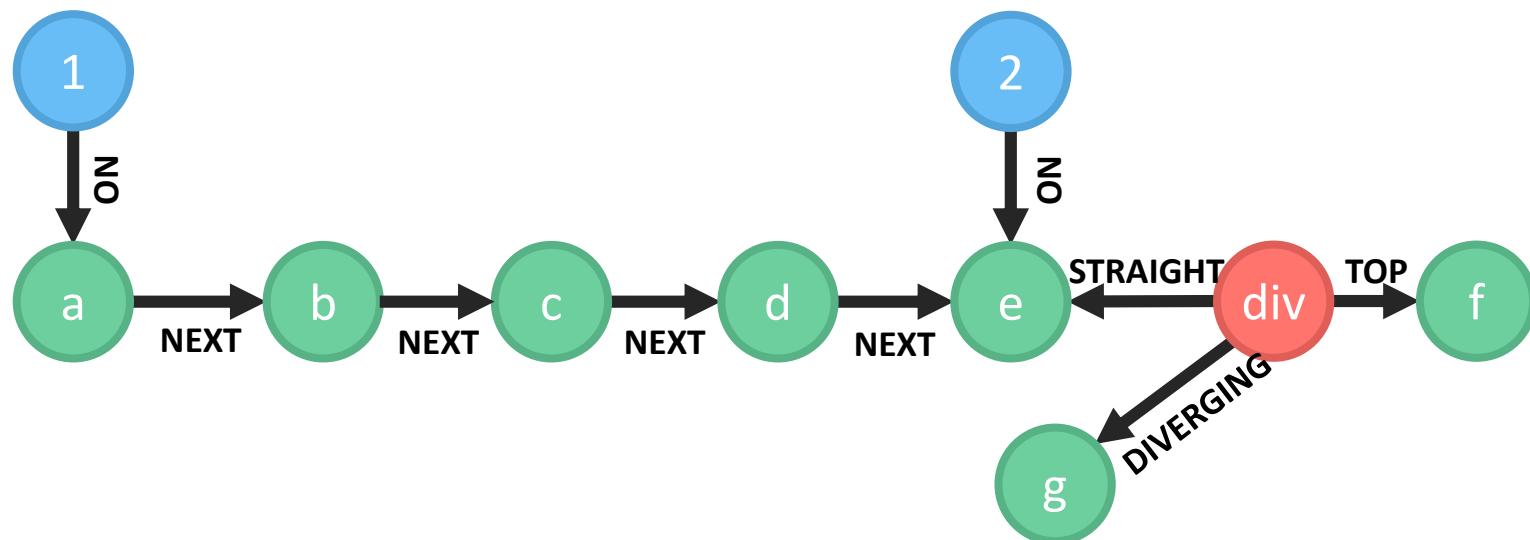


$\pi_{t.number, sw}$

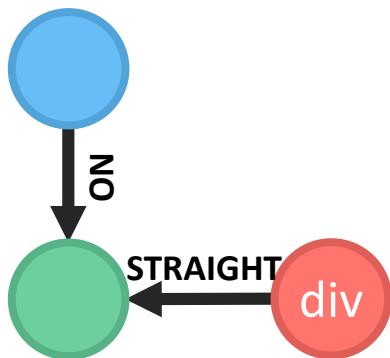
$\sigma_{sw.position} = 'diverging'$



STRAIGHT



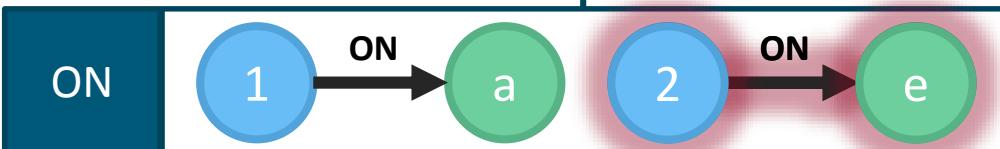
## Trailing the switch



$\pi_{t.number, sw}$

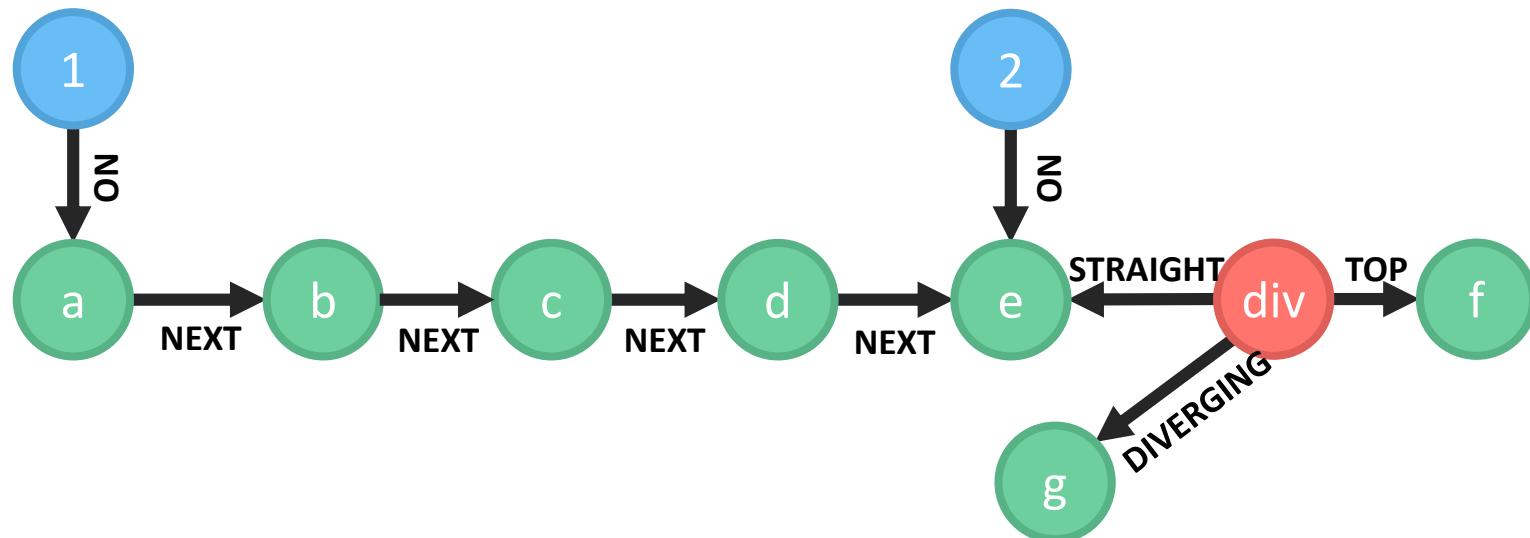
$\sigma_{sw.position} = 'diverging'$

⊗

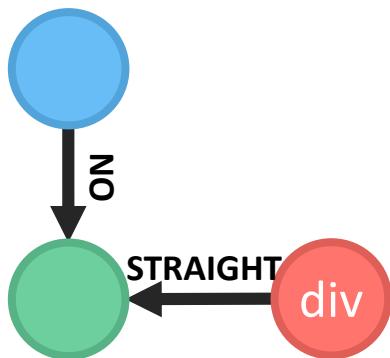


STRAIGHT

div

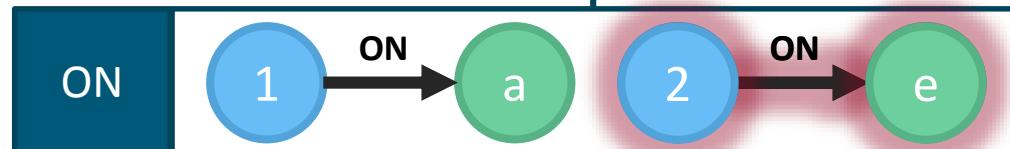


## Trailing the switch

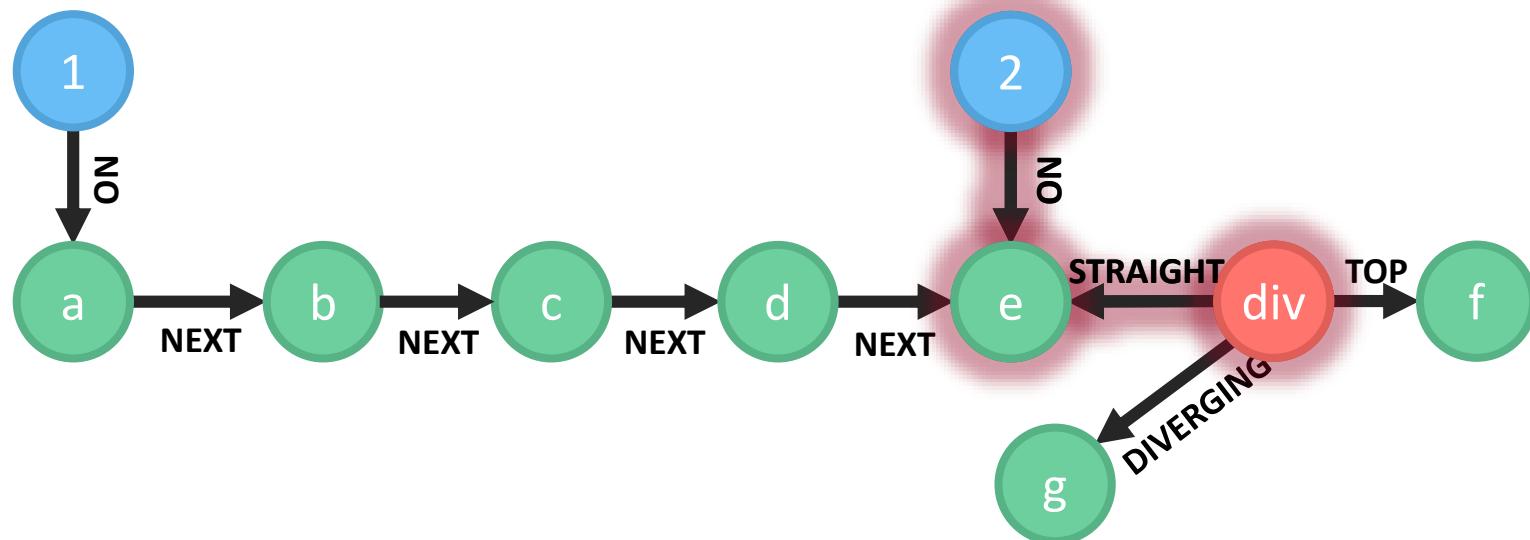


$\pi_{t.number, sw}$

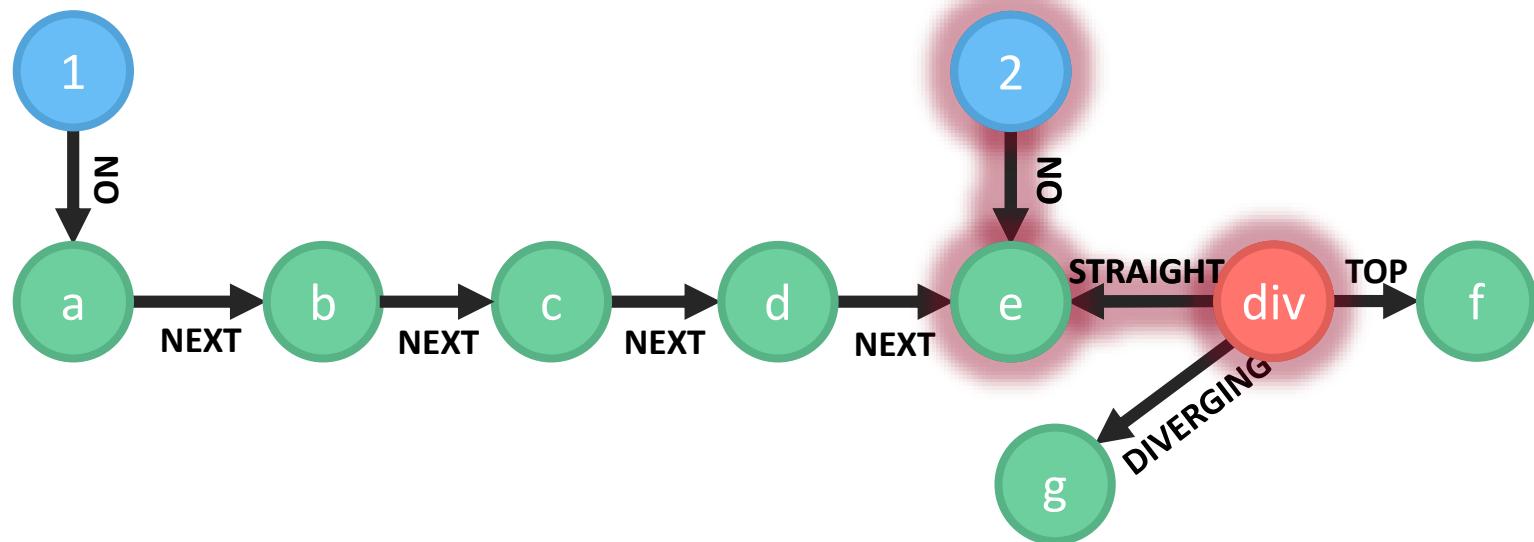
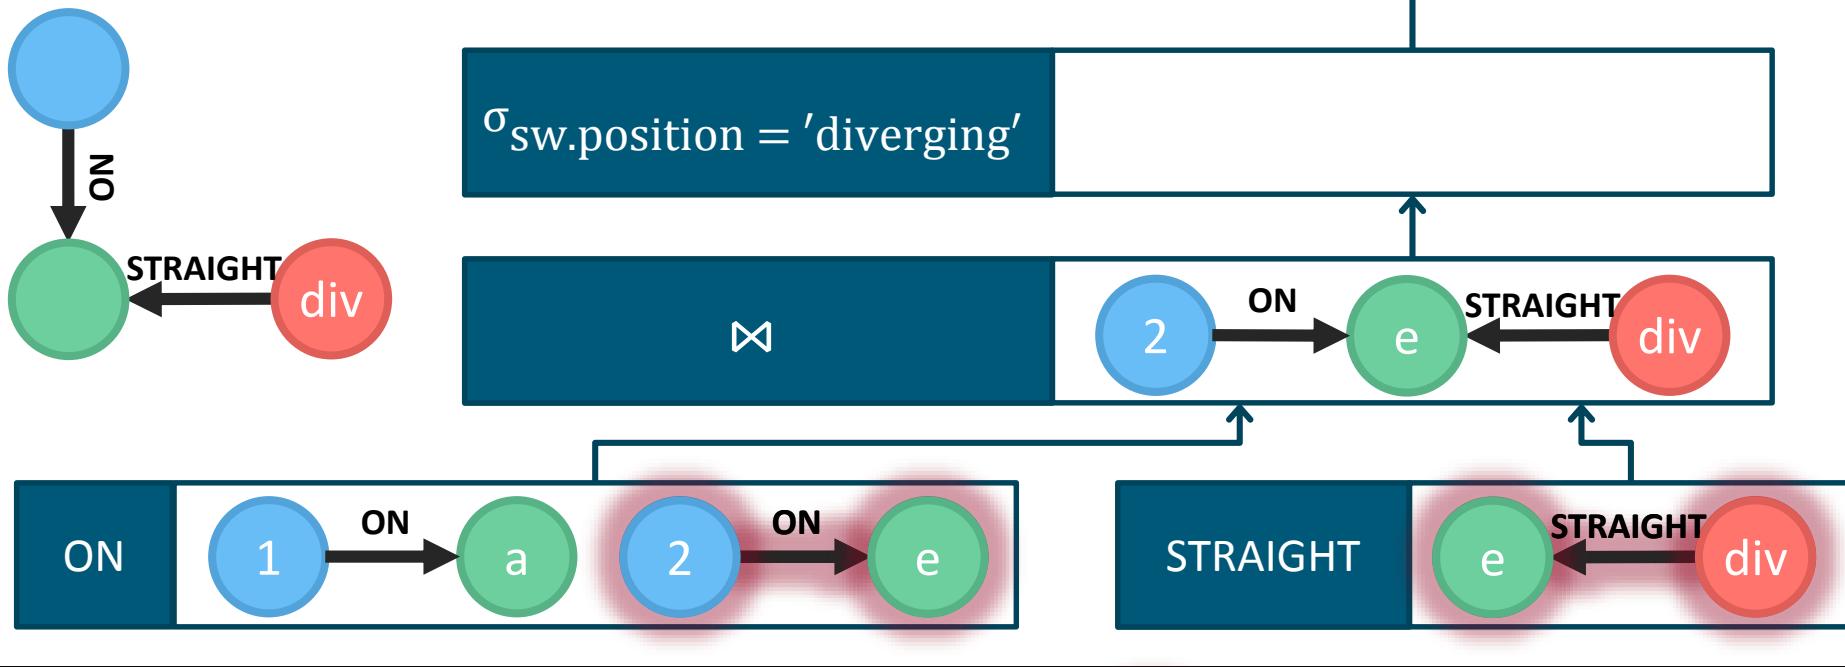
$\sigma_{sw.position} = 'diverging'$



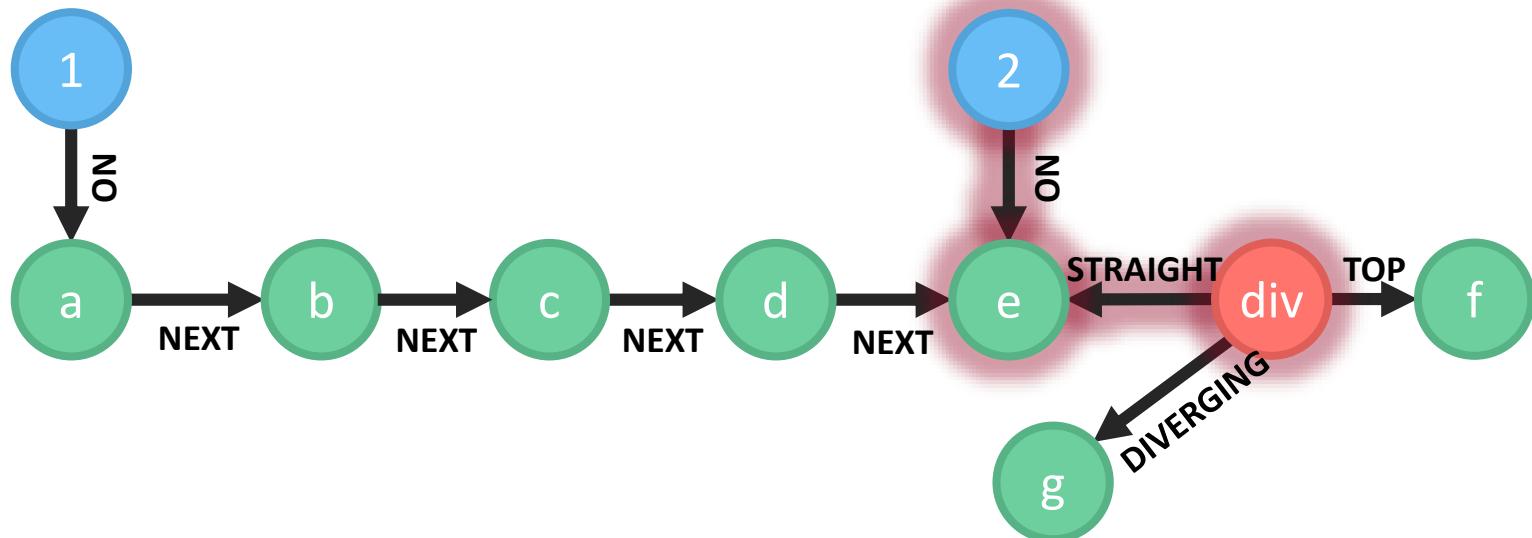
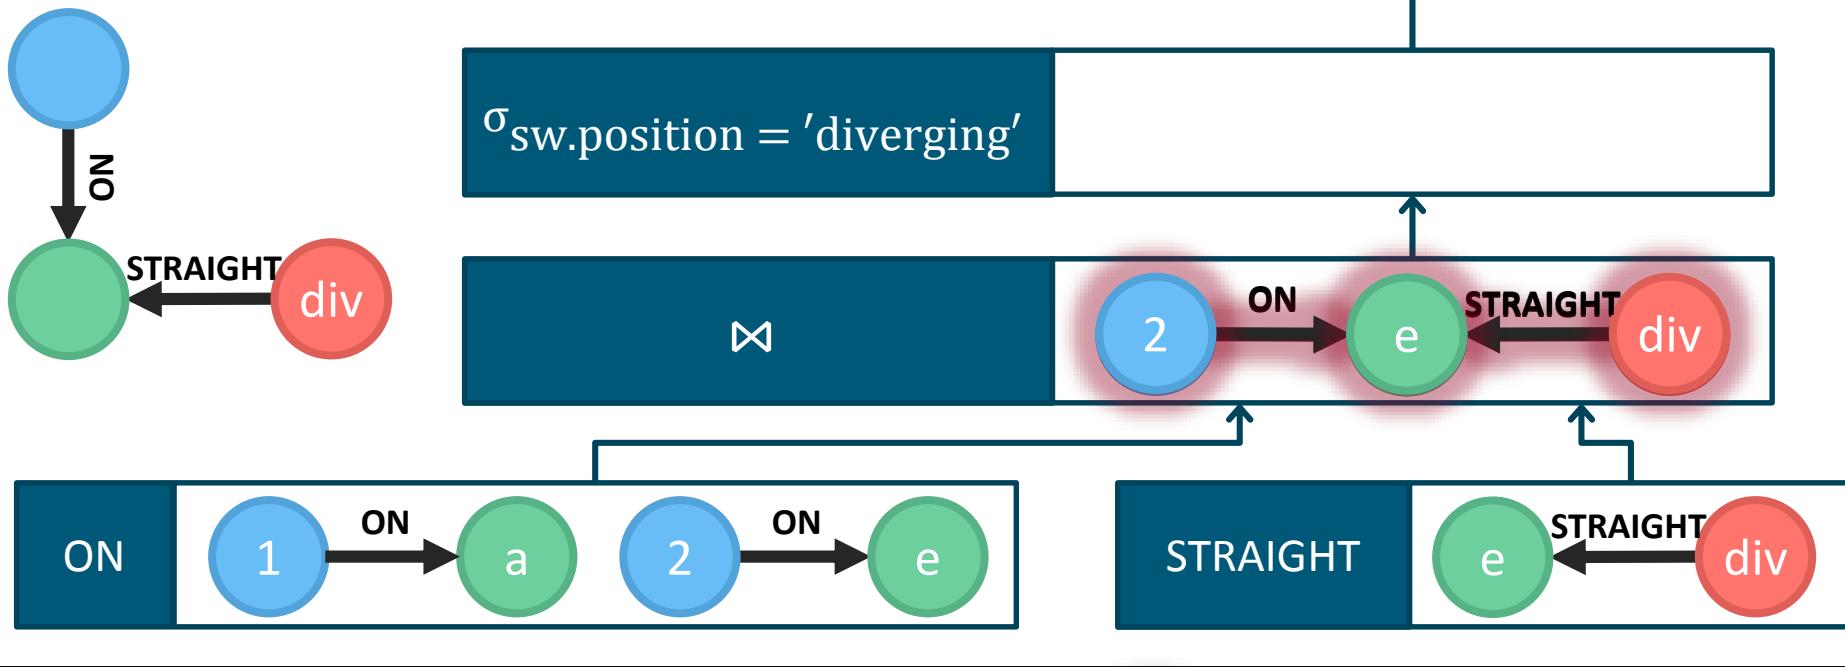
STRAIGHT



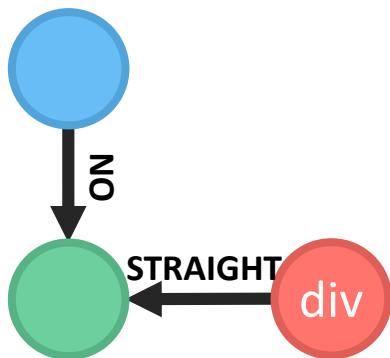
## Trailing the switch



## Trailing the switch



## Trailing the switch

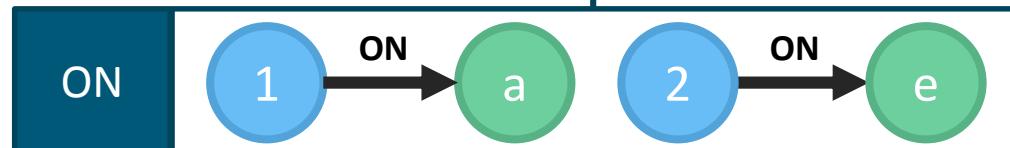


$\pi_{t.number, sw}$

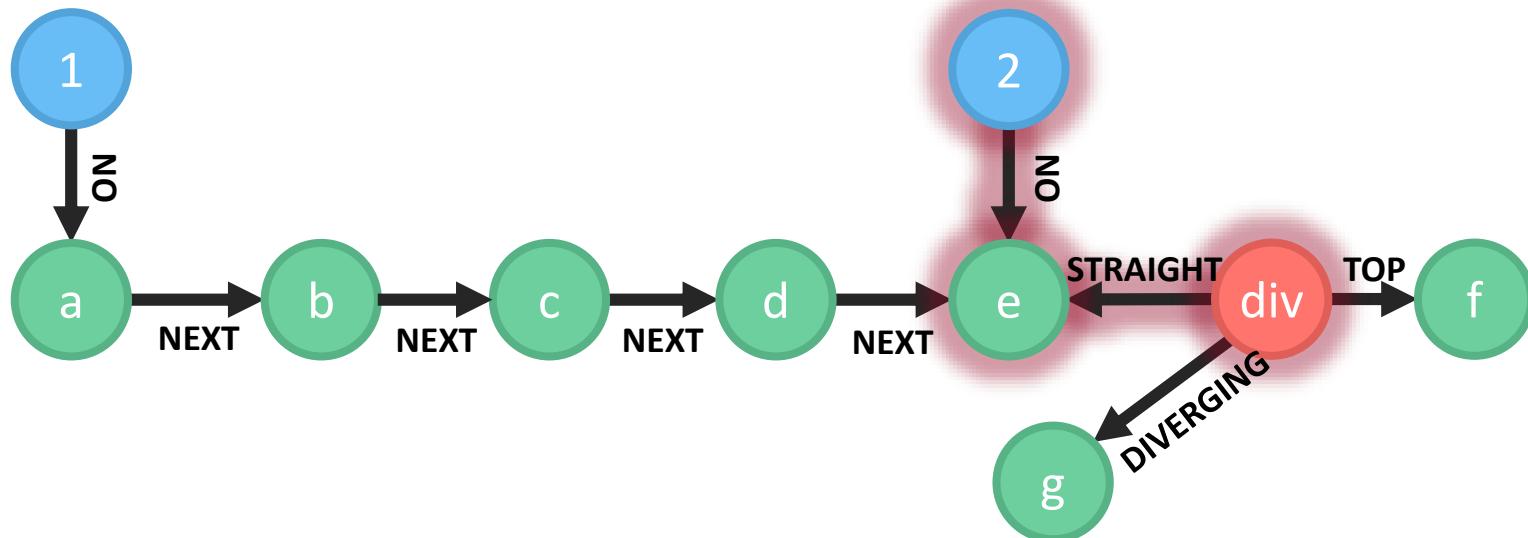
$\sigma_{sw.position} = 'diverging'$



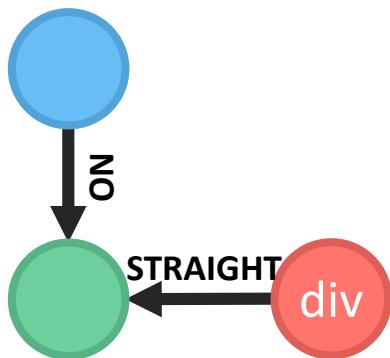
$\bowtie$



STRAIGHT



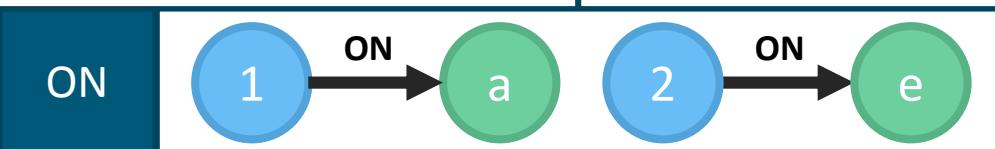
## Trailing the switch



$\pi_{t.number, sw}$

$\sigma_{sw.position} = 'diverging'$

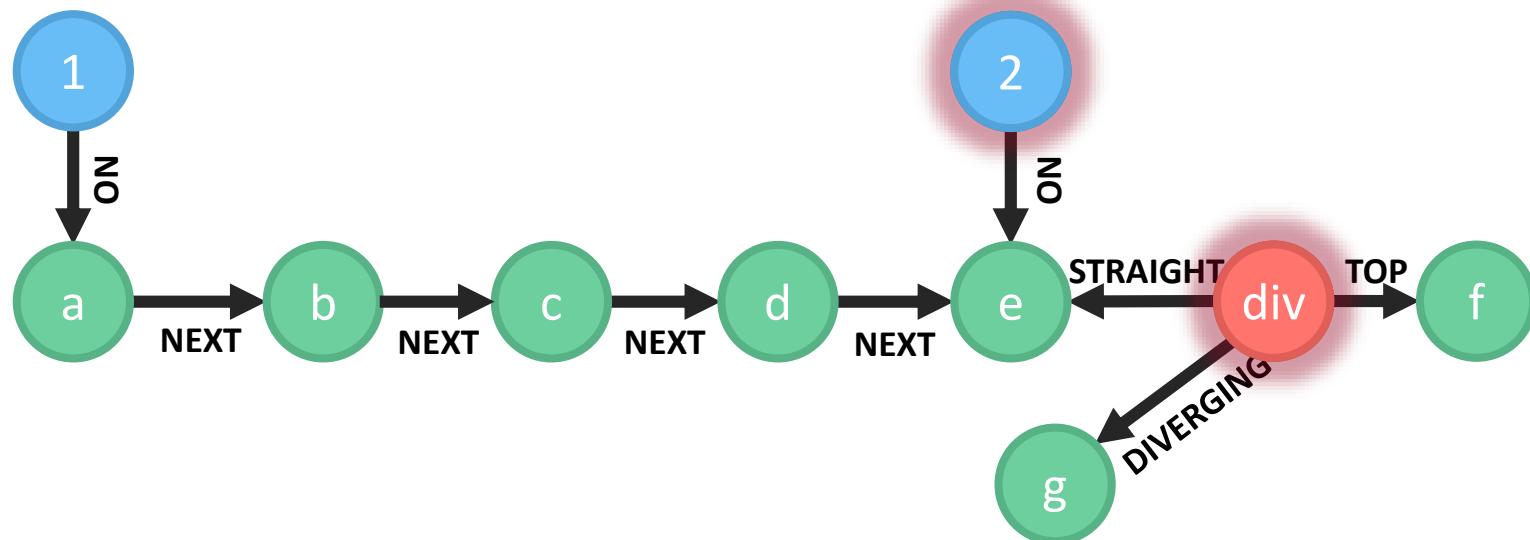
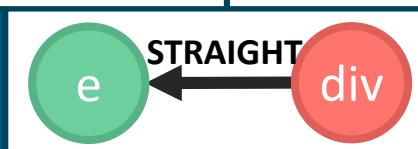
$\bowtie$



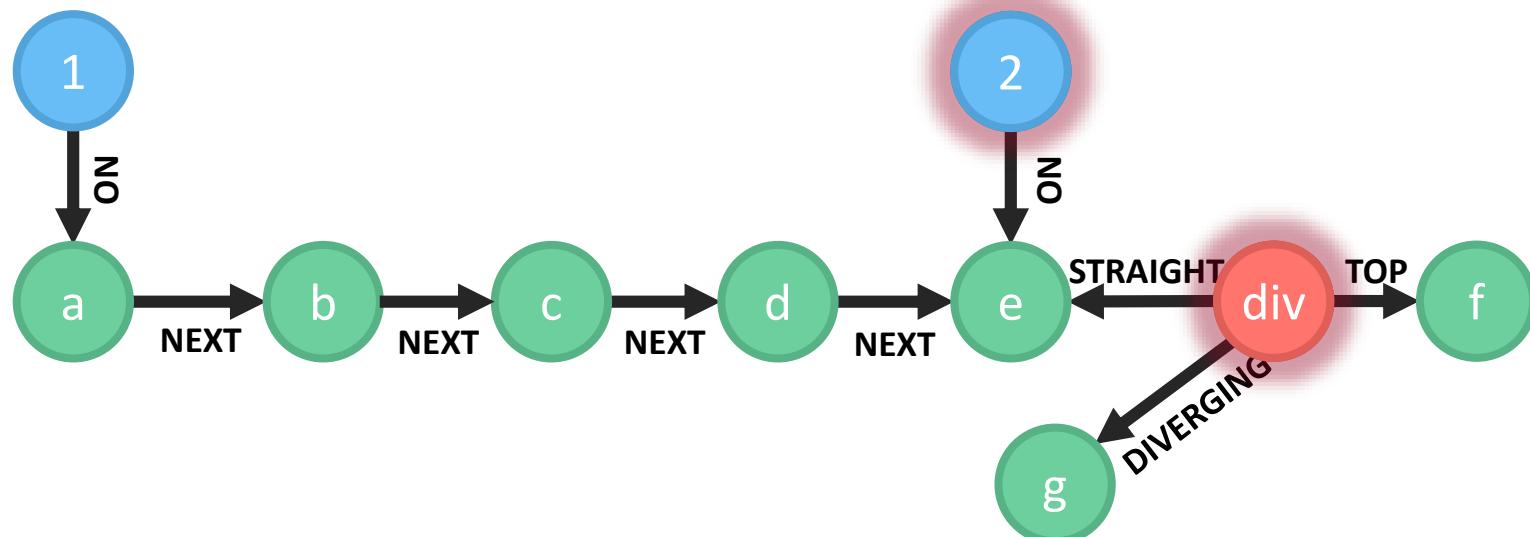
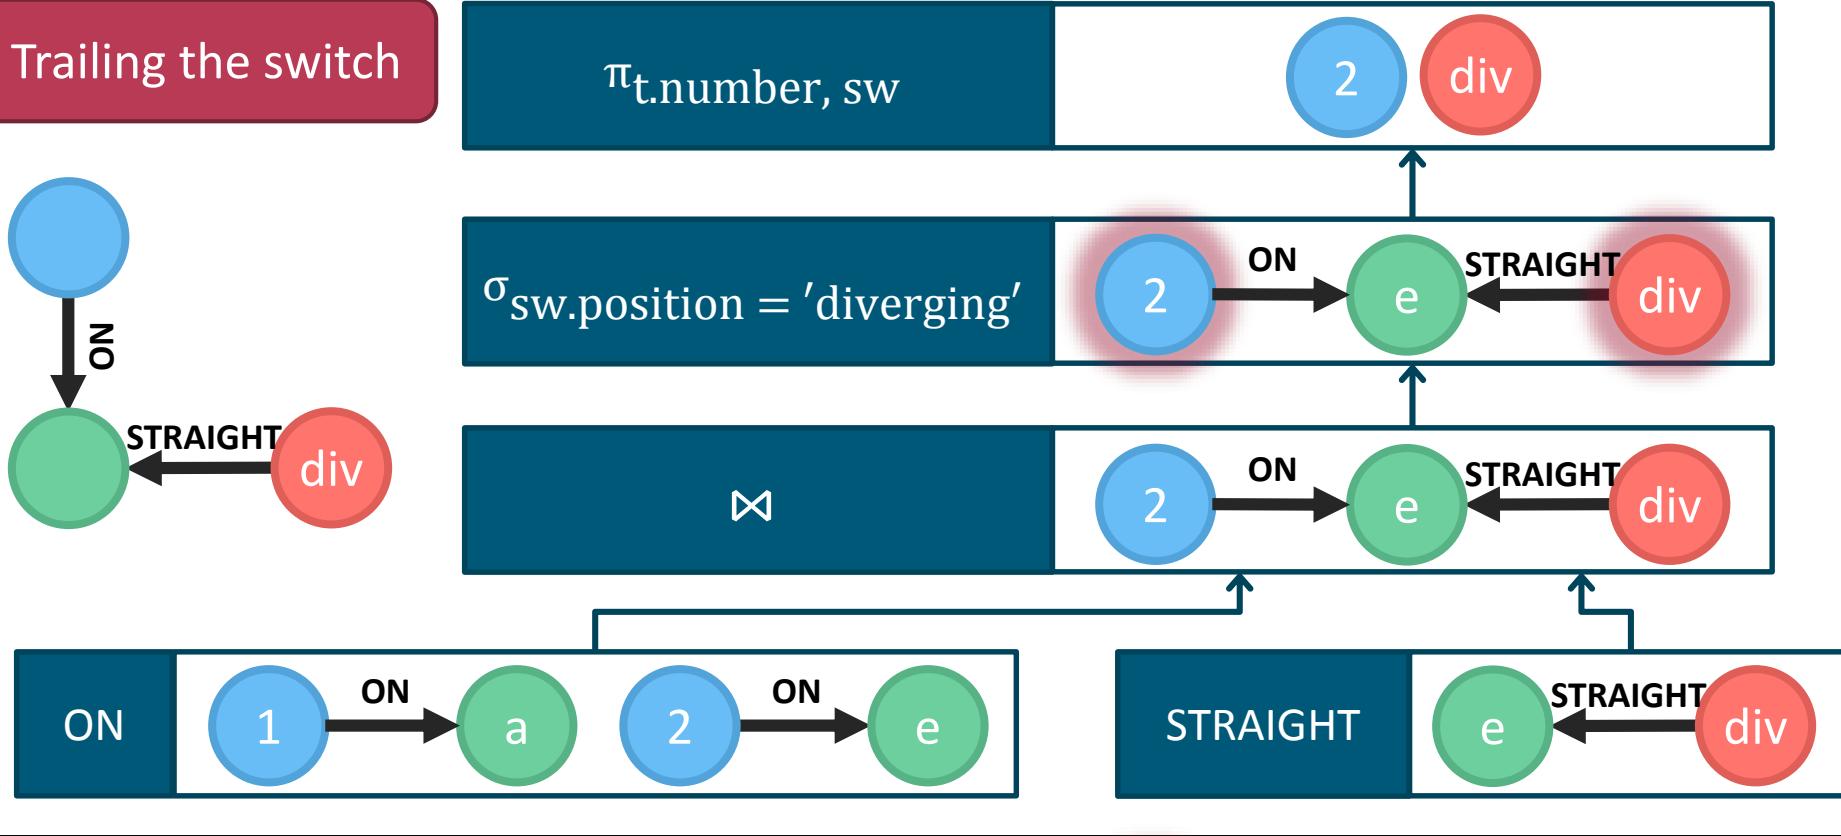
ON → e → STRAIGHT ← div

ON → e → STRAIGHT ← div

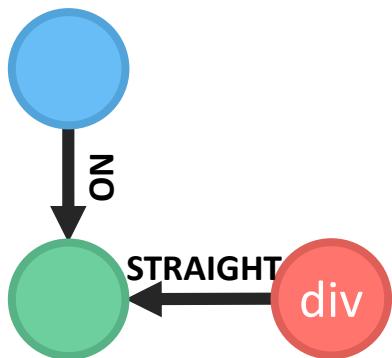
STRAIGHT



## Trailing the switch



## Trailing the switch



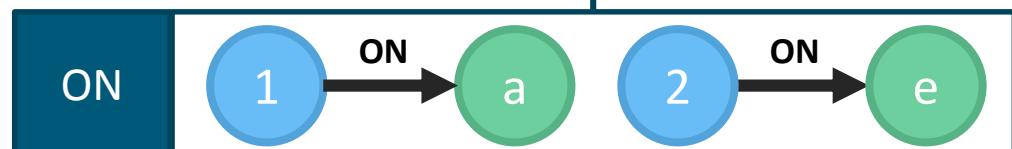
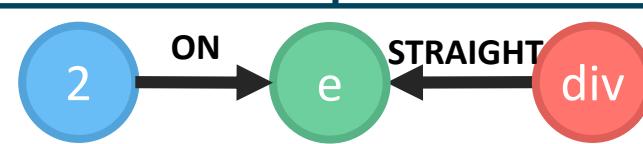
$\pi_{t.number, sw}$



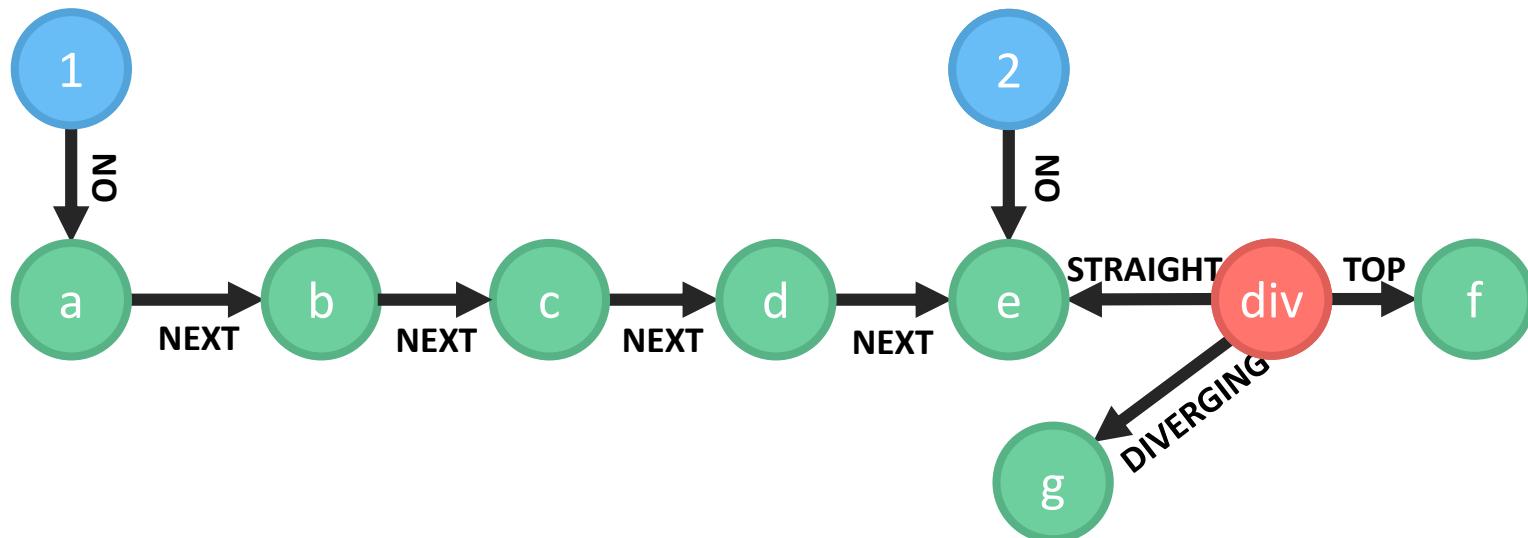
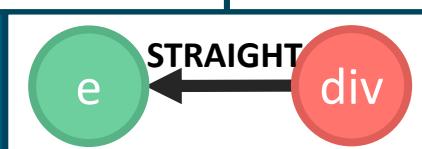
$\sigma_{sw.position} = 'diverging'$



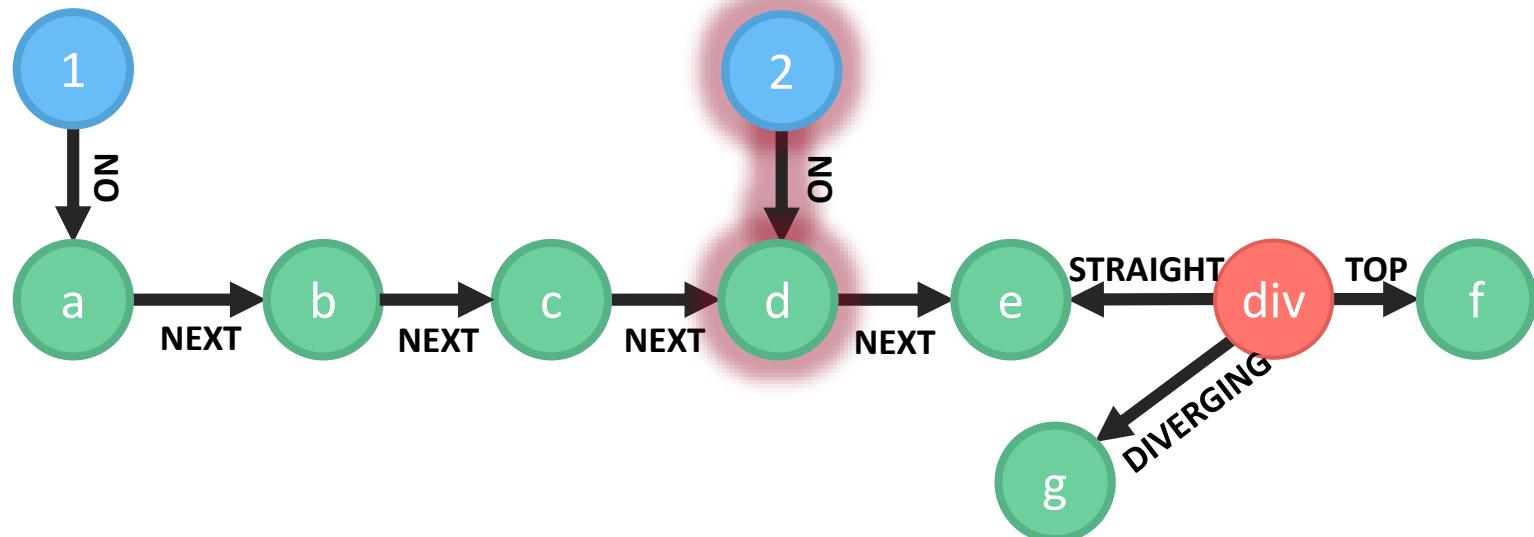
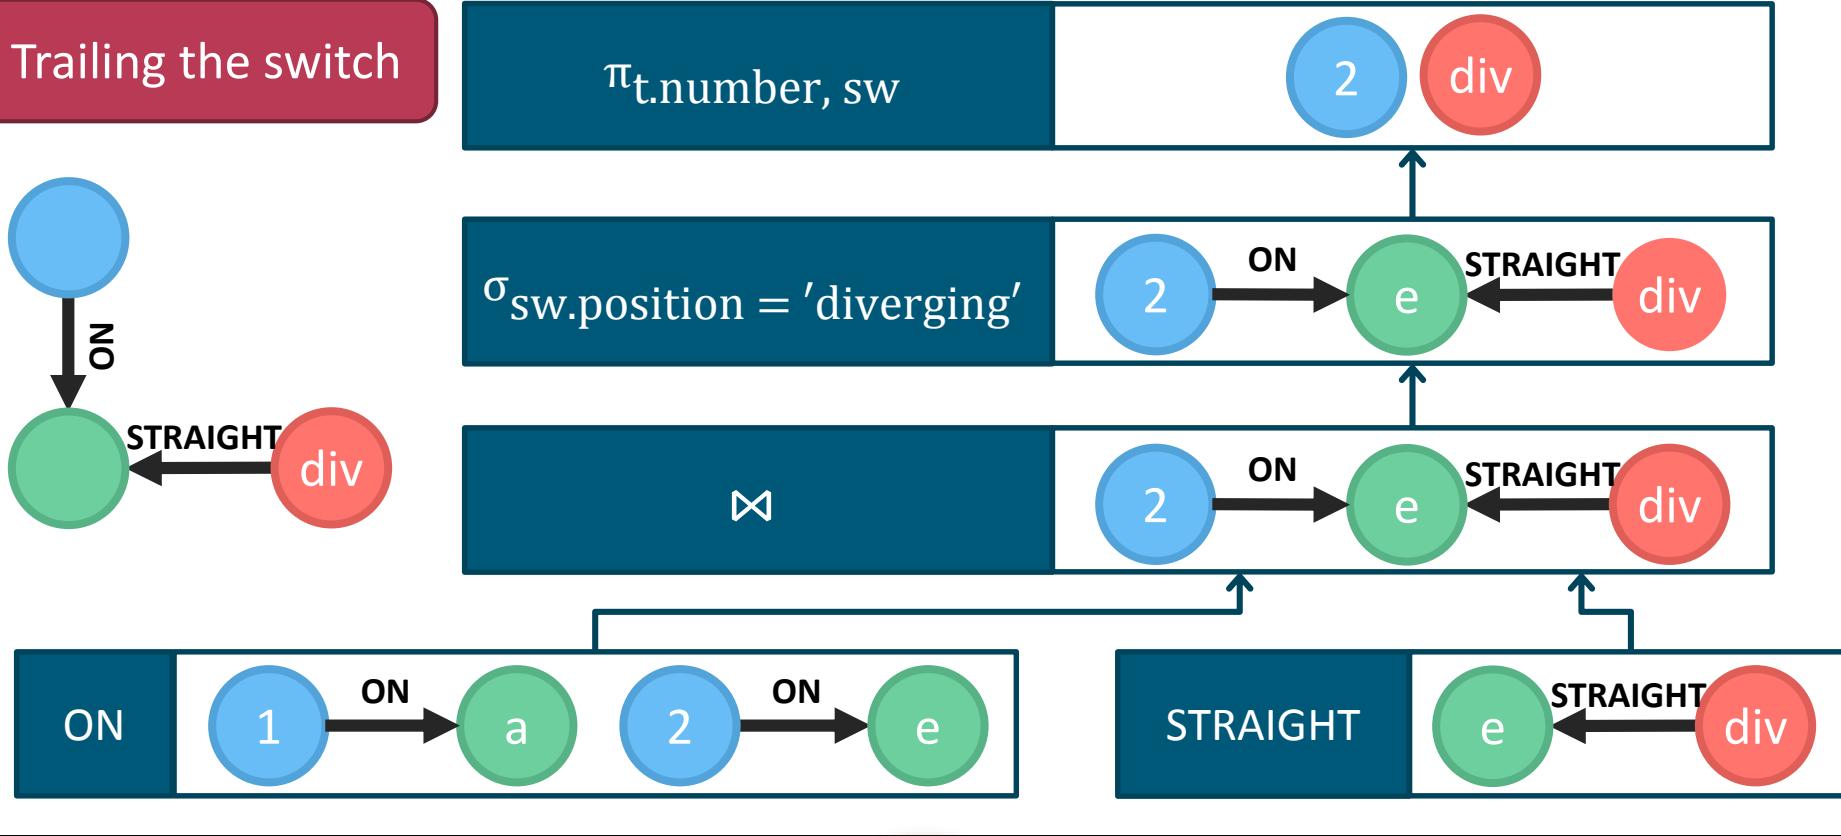
$\bowtie$



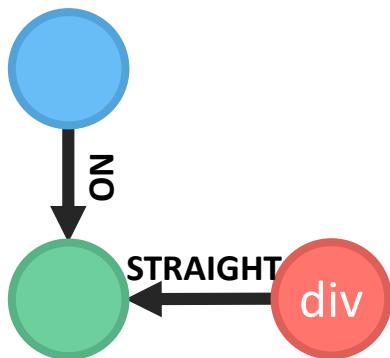
STRAIGHT



## Trailing the switch



## Trailing the switch



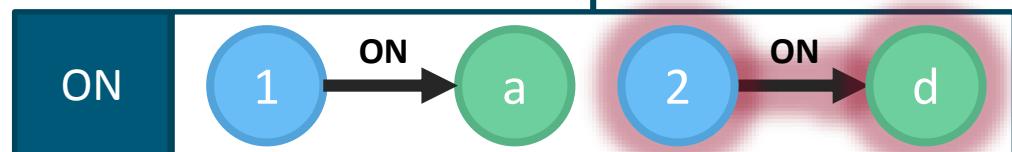
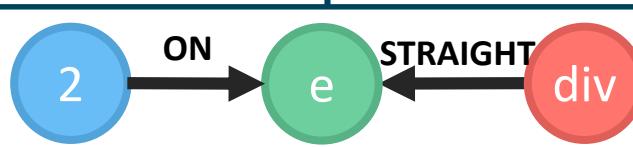
$\pi_{t.\text{number}, \text{sw}}$



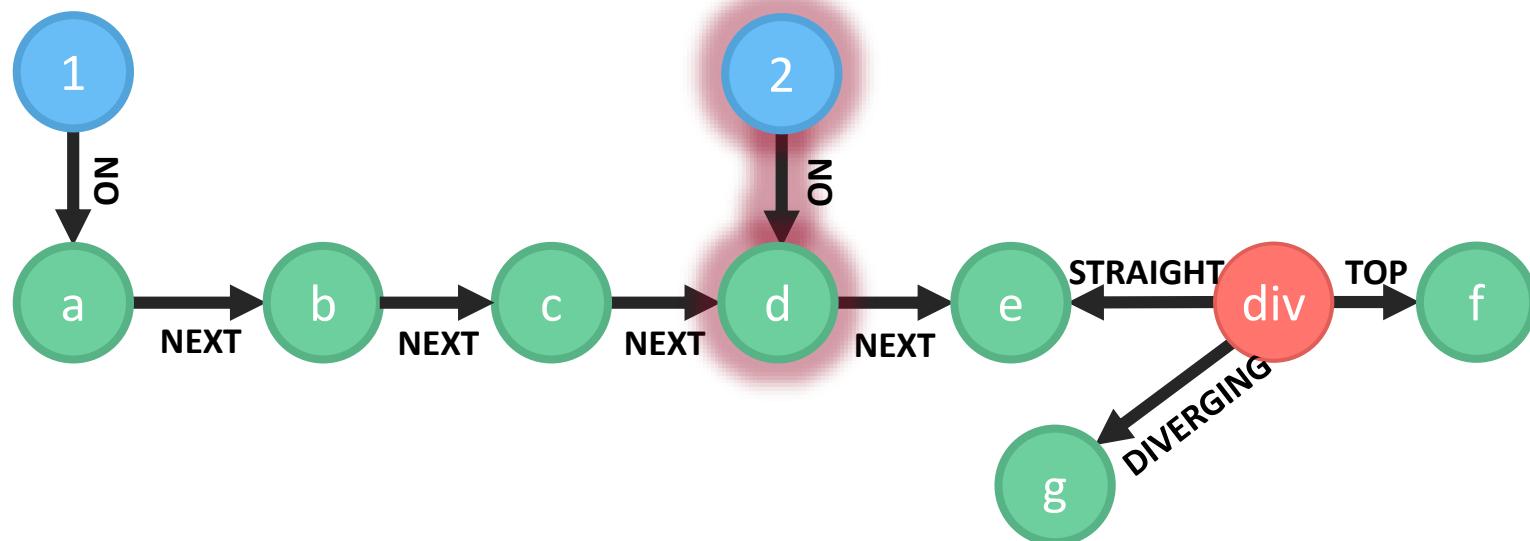
$\sigma_{\text{sw.position}} = \text{'diverging'}$



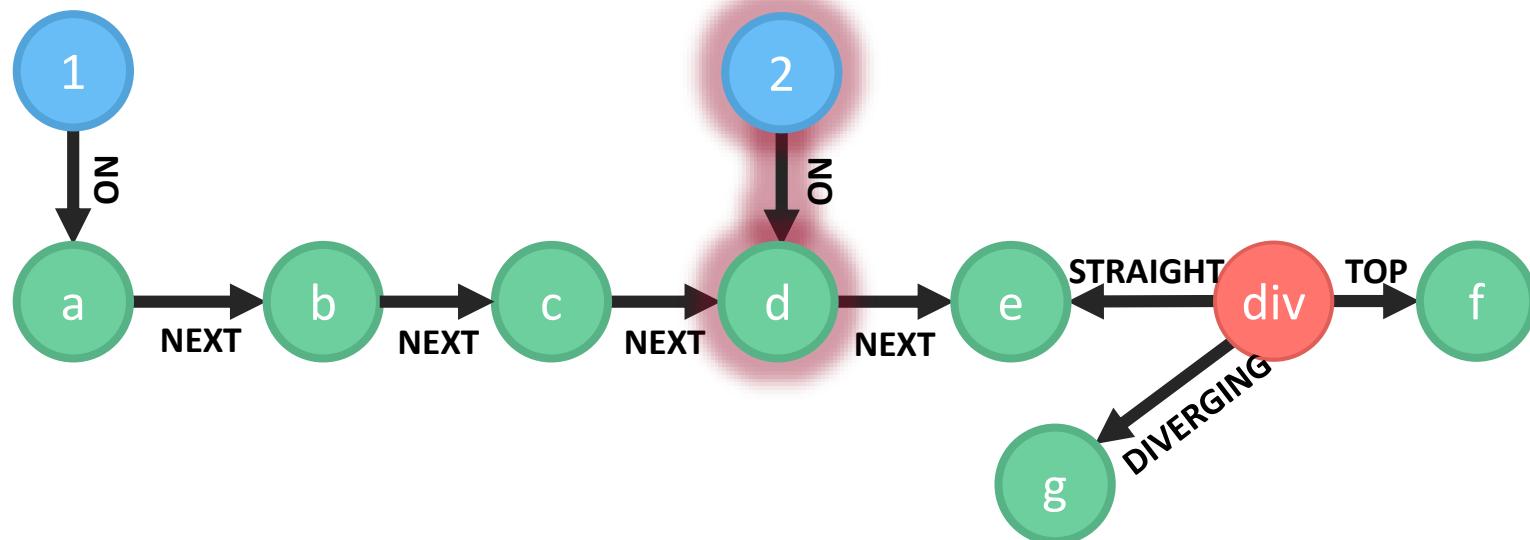
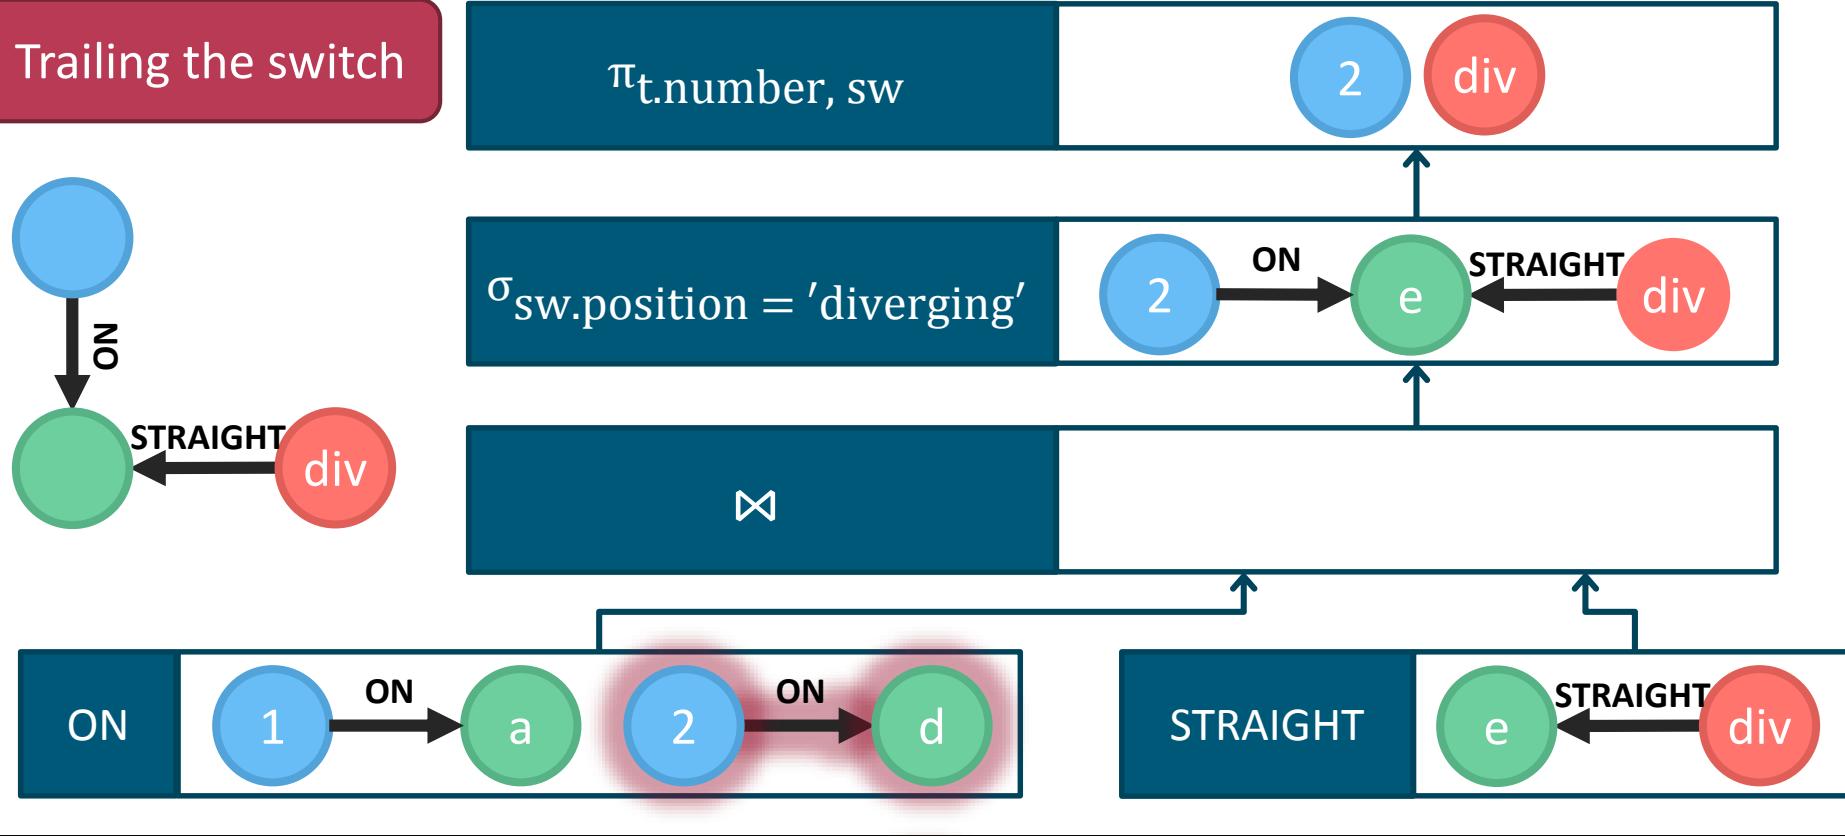
$\bowtie$



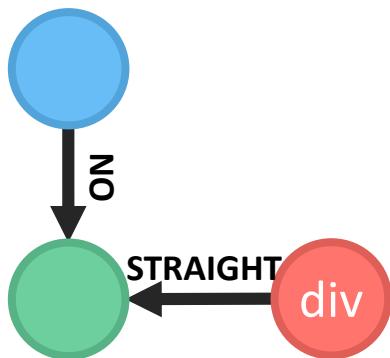
STRAIGHT



## Trailing the switch



## Trailing the switch

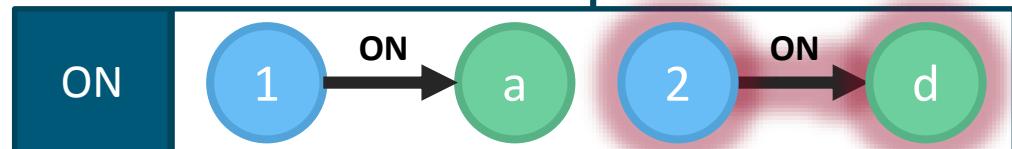


$\pi_{t.\text{number}, \text{sw}}$

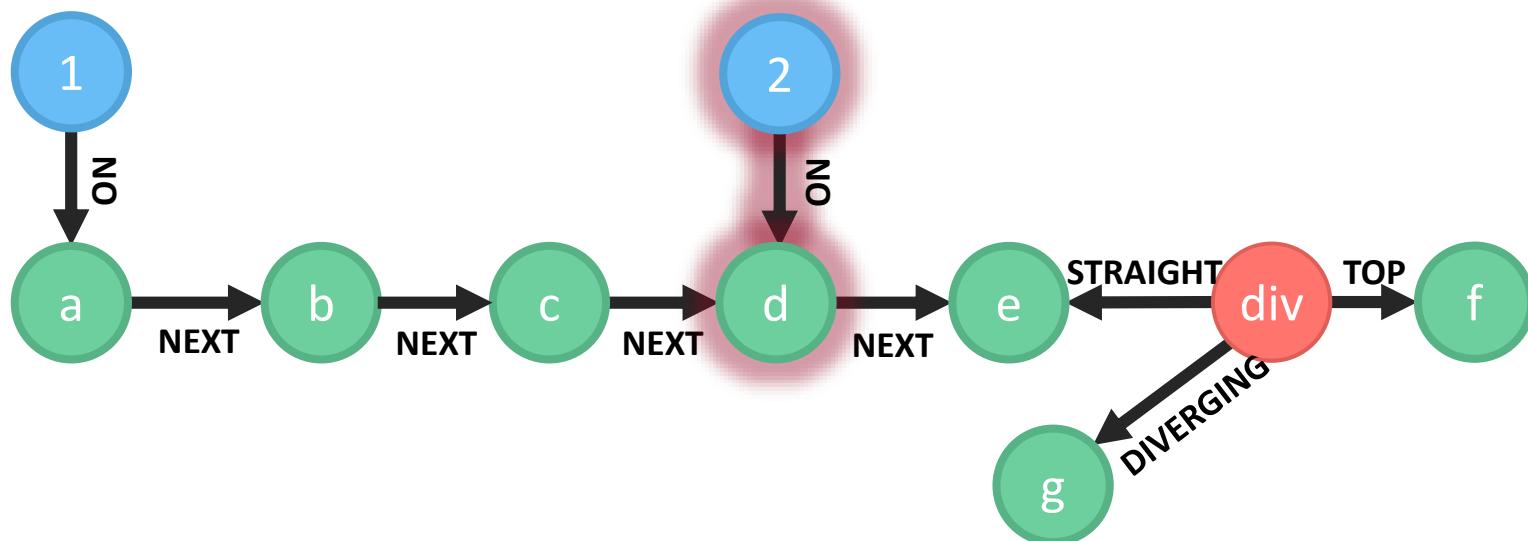
2

div

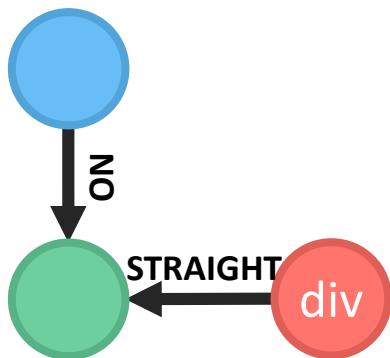
$\sigma_{\text{sw.position}} = \text{'diverging'}$



STRAIGHT

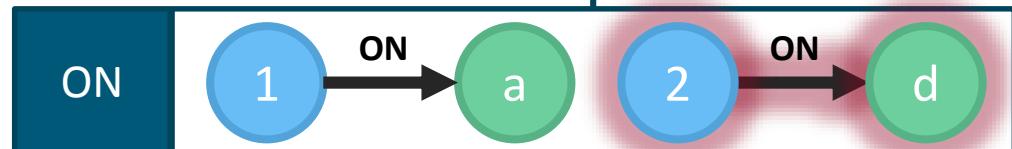


## Trailing the switch

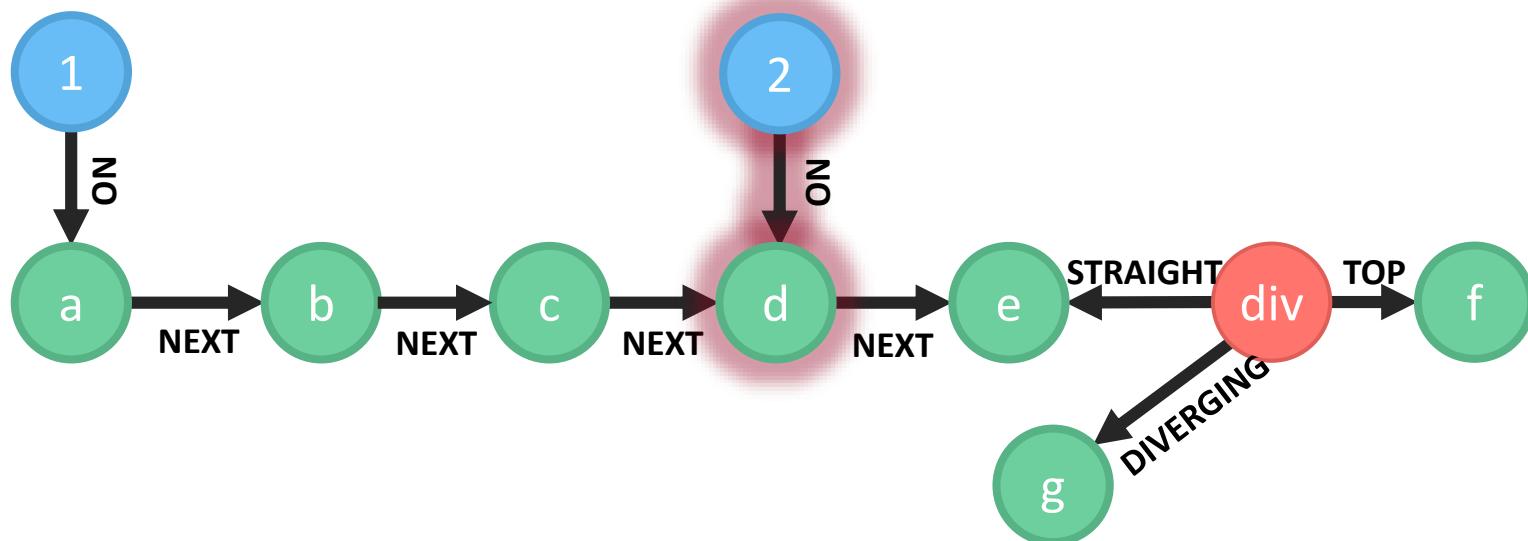


$\pi_{t.\text{number}, \text{sw}}$

$\sigma_{\text{sw.position}} = \text{'diverging'}$



STRAIGHT



# Batch vs. incremental queries

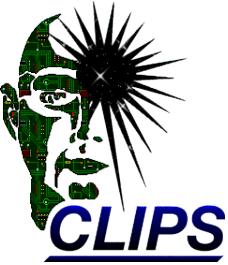
- **Batch queries**  
(pull / request-driven):
  1. Client selects a query
  2. Results are calculated
- Query results  
obtained on demand

# Batch vs. incremental queries

- **Batch queries**  
(pull / request-driven):
  1. Client selects a query
  2. Results are calculated
- Query results obtained on demand
- **Incremental queries**  
(push / event-driven):
  1. Client registers queries
  2. Graph is changed
  3. Results are maintained
  4. Goto 2
- Query results are always available

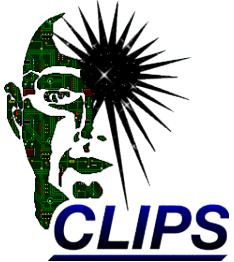
# Incremental query engines

- CLIPS                      C structures      NASA
- Drools                      POJO                  Red Hat
- VIATRA                      EMF                  BME / IncQuery Labs.



# Incremental query engines

■ <a href="#"><u>CLIPS</u></a>	C structures	NASA
■ <a href="#"><u>Drools</u></a>	POJO	Red Hat
■ <a href="#"><u>VIATRA</u></a>	EMF	BME / IncQuery Labs.



■ <a href="#"><u>INSTANS</u></a>	RDF	Aalto University
■ <a href="#"><u>i3QL</u></a>	POJO	TU Darmstadt
■ <a href="#"><u>IncQuery-D</u></a>	RDF	BME

# Incremental query engines

■ <a href="#"><u>CLIPS</u></a>	C structures	NASA
■ <a href="#"><u>Drools</u></a>	POJO	Red Hat
■ <a href="#"><u>VIATRA</u></a>	EMF	BME / IncQuery Labs.



■ <a href="#"><u>INSTANS</u></a>	RDF	Aalto University
■ <a href="#"><u>i3QL</u></a>	POJO	TU Darmstadt
■ <a href="#"><u>IncQuery-D</u></a>	RDF	BME
■ <b>No implementations for property graphs yet</b>		

# Incremental Graph Queries with openCypher

# Incremental Graph Queries

## with openCypher

# openCypher

„The openCypher project aims to deliver a full and open specification of the industry’s most widely adopted graph database query language: Cypher.”

- Grammar specification
- Technology Compatibility Kit (TCK)
- Reference implementation

# openCypher constructs

## ■ Standard constructs

- pattern matching
- filtering
- lists, maps
- data manipulation
- variable length paths

# openCypher constructs

- Standard constructs
  - pattern matching
  - filtering
  - lists, maps
  - data manipulation
  - variable length paths
  
- Legacy constructs
  - indexing, constraints
  - regular expressions
  - some list functions, including reduce
  - most predicate functions
  - shortest path functions
  - **CASE** expressions
  - **id()**

# openCypher constructs

## ■ Standard constructs

- pattern matching
- filtering
- lists, maps
- data manipulation
- variable length paths

Difficult to handle  
incrementally

## ■ Legacy constructs

- indexing, constraints
- regular expressions
- some list functions,  
including reduce
- most predicate functions
- shortest path functions
- **CASE** expressions
- **id()**

# Mapping openCypher to relational algebra

## Combining and filtering pattern matches

<code>MATCH «p»</code>	$\exists_{\text{edges of } p} (p)$
<code>MATCH «p1», «p2»</code>	$\exists_{\text{edges of } p1 \text{ and } p2} (p_1 \bowtie p_2)$
<code>MATCH «p1» MATCH «p2»</code>	$\exists_{\text{edges of } p1} (p_1) \bowtie \exists_{\text{edges of } p2} (p_2)$
<code>MATCH «p1» OPTIONAL MATCH «p2»</code>	$\exists_{\text{edges of } p1} (p_1) \bowtie \exists_{\text{edges of } p2} (p_2)$
<code>MATCH «p» WHERE «condition»</code>	$\sigma_{\text{condition}}(r)$ , where condition may specify patterns and arithmetic constraints on existing variables

Result and sub-result operations. Rules for `RETURN` also apply to `WITH`.

<code>RETURN «variables»</code>	$\pi_{\text{variables}}(r)$
<code>RETURN «v1» AS «alias1» ...</code>	$\pi_{v1 \rightarrow \text{alias1}, \dots}(r)$
<code>RETURN DISTINCT «variables»</code>	$\delta(\pi_{\text{variables}}(r))$
<code>RETURN «variables», «aggregates»</code>	$\gamma_{\text{variables, aggregates}}(r)$

## List operations

<code>ORDER BY «v1» [ASC DESC] ...</code>	$\tau_{\uparrow/\downarrow v1, \dots}(r)$
<code>LIMIT «l»</code>	$\lambda_l(r)$

# Mapping openCypher to relational algebra

Combining and filtering pattern matches

<code>MATCH «p»</code>	$\not\exists_{\text{edges of } p} (p)$
<code>MATCH «p1», «p2»</code>	$\not\exists_{\text{edges of } p1 \text{ and } p2} (p_1 \bowtie p_2)$
<code>MATCH «p1» MATCH «p2»</code>	$\not\exists_{\text{edges of } p1} (p_1) \bowtie \not\exists_{\text{edges of } p2} (p_2)$
<code>MATCH «p1» OPTIONAL MATCH «p2»</code>	$\not\exists_{\text{edges of } p1} (p_1) \bowtie \not\exists_{\text{edges of } p2} (p_2)$
<code>MATCH «p» WHERE «condition»</code>	$\sigma_{\text{condition}}(r)$ , where condition may specify patterns and arithmetic constraints on existing variables

Result and sub-result operations. Rules for `RETURN` also apply to `WITH`.

<code>RETURN «variables»</code>	$\pi_{\text{variables}}(r)$
<code>RETURN «v1» AS «alias1» ...</code>	$\pi_{v1 \rightarrow \text{alias1}, \dots}(r)$
<code>RETURN DISTINCT «variables»</code>	$\delta(\pi_{\text{variables}}(r))$
<code>RETURN «variables», «aggregates»</code>	$\gamma_{\text{variables, aggregates}}(r)$

List operations

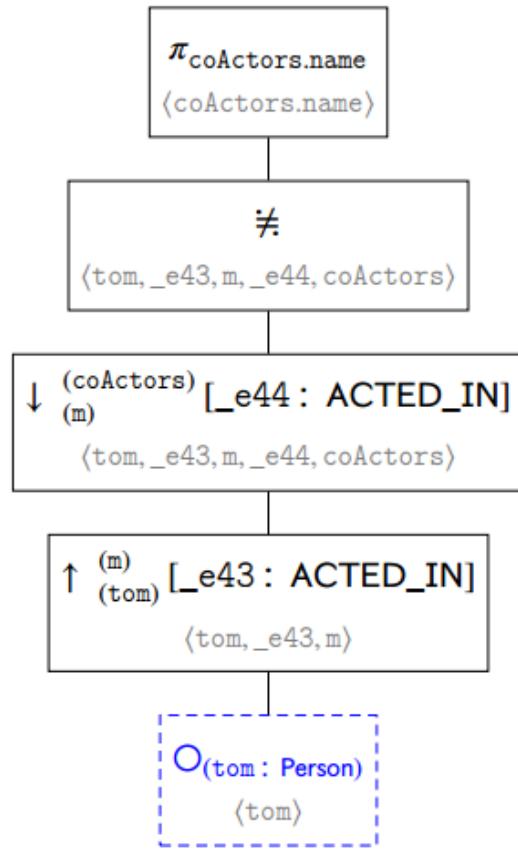
<code>ORDER BY «v1» [ASC DESC] ...</code>	$\tau_{\uparrow/\downarrow v1}$	<b>Technical report</b>
<code>LIMIT «l»</code>	$\lambda_l(r)$	

## Query specification (query-18)

```
1 MATCH (tom:Person {name:"Tom Hanks"})-[:ACTED_IN]->(m)<-[:ACTED_IN]-(coActors)
2 RETURN coActors.name
```

## Query specification (query-18)

```
1 MATCH (tom:Person {name:"Tom Hanks"})-[:ACTED_IN]->(m)<-[:ACTED_IN]-(coActors)
2 RETURN coActors.name
```



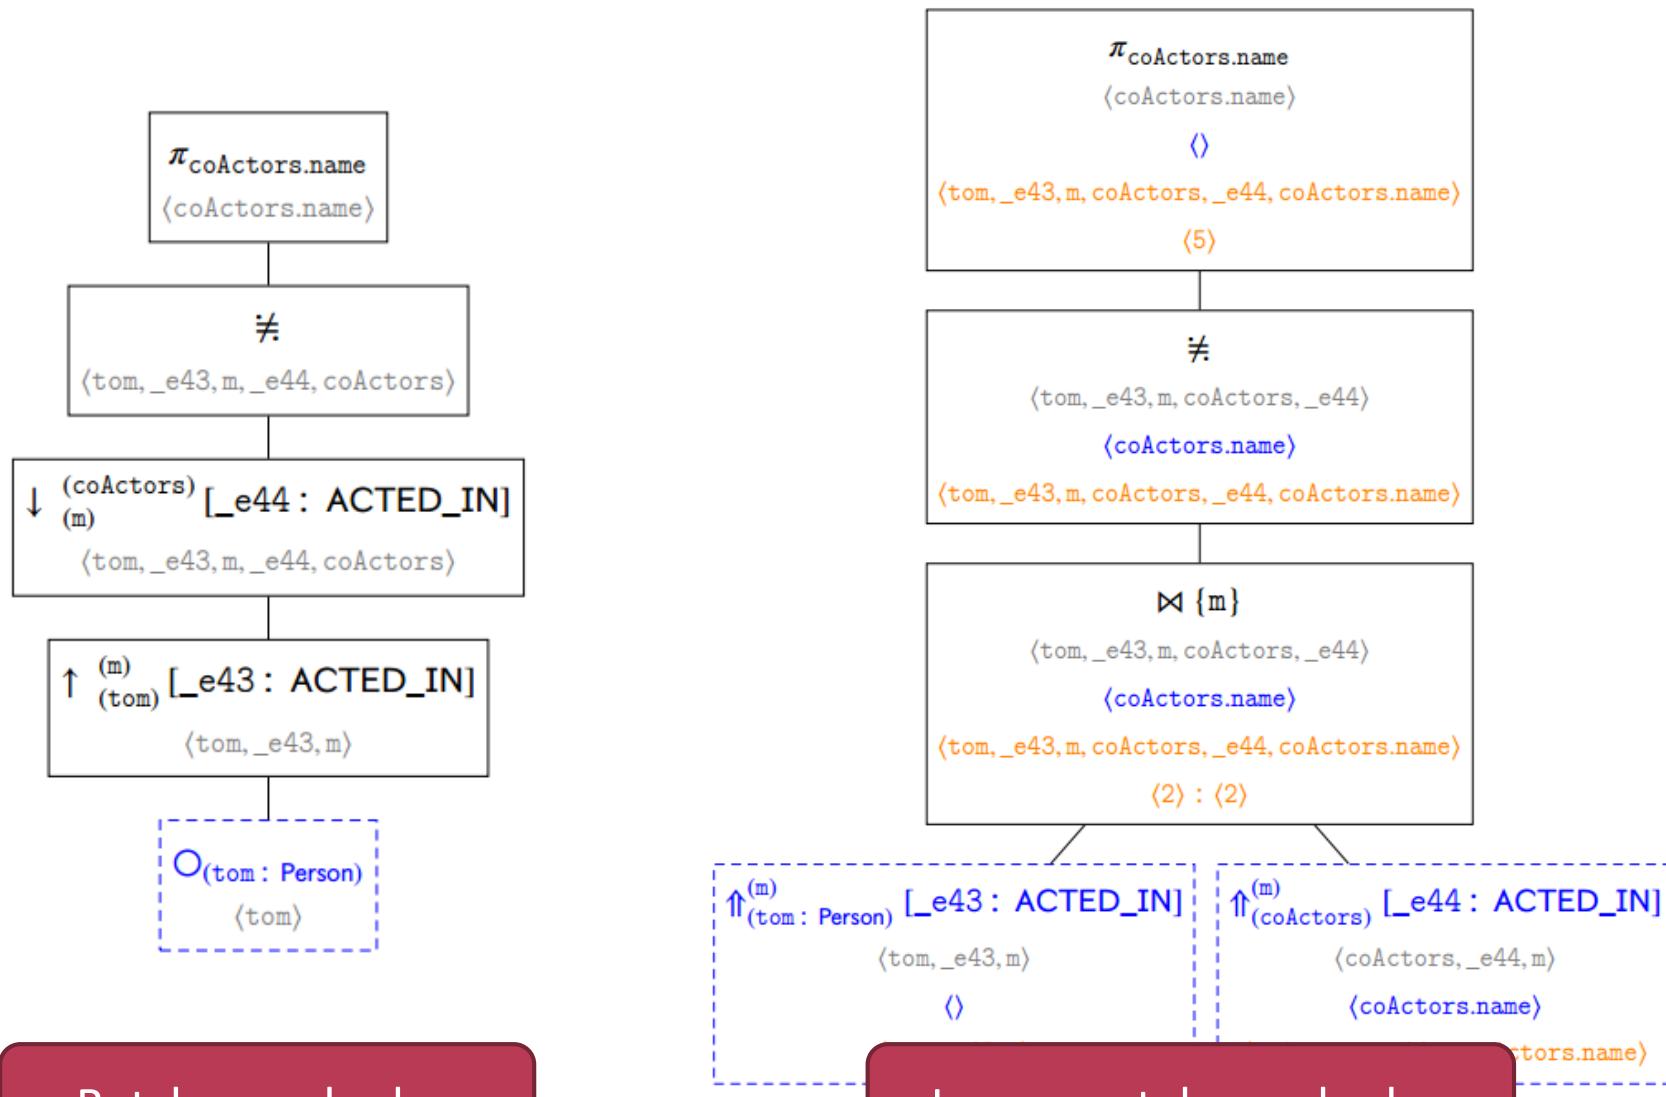
Batch search plan

## Query specification (query-18)

```

1 MATCH (tom:Person {name:"Tom Hanks"})-[:ACTED_IN]->(m)<-[:ACTED_IN]-(coActors)
2 RETURN coActors.name

```



Batch search plan

Incremental search plan

# Incremental openCypher challenges

- Lists
  - `['a', 1, 2, true]`
  - `['a', [1, [2]], true]`
  - `UNWIND`
  - `collect()`
- Bag semantics, `ORDER BY`, `SKIP` and `LIMIT`
- Legacy constructs
  - `shortestPath()`, `allShortestPaths()`
  - `reduce()`

# Incremental openCypher challenges

- Lists
  - `['a', 1, 2, true]`
  - `['a', [1, [2]], true]`
  - `UNWIND`
  - `collect()`
- Bag semantics, `ORDER BY`, `SKIP` and `LIMIT`
- Legacy constructs
  - `shortestPath()`, `allShortestPaths()`
  - `reduce()`

These are not handled  
by traditional Rete  
implementations

# igraph

- An incremental, in-memory graph query engine

# ingraph

- An incremental, in-memory graph query engine

client

ingraph

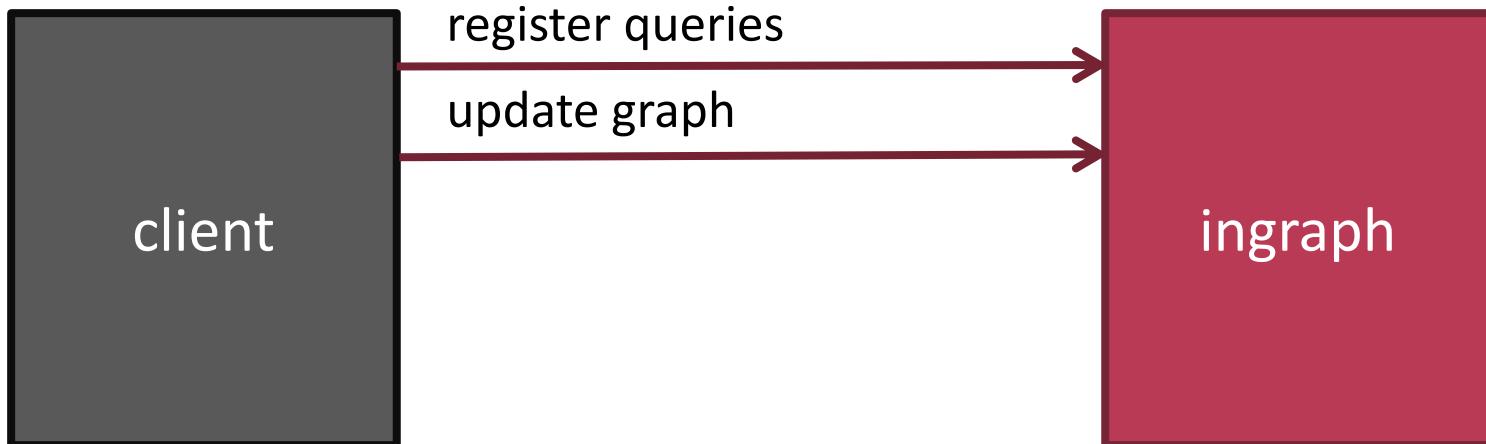
# ingraph

- An incremental, in-memory graph query engine



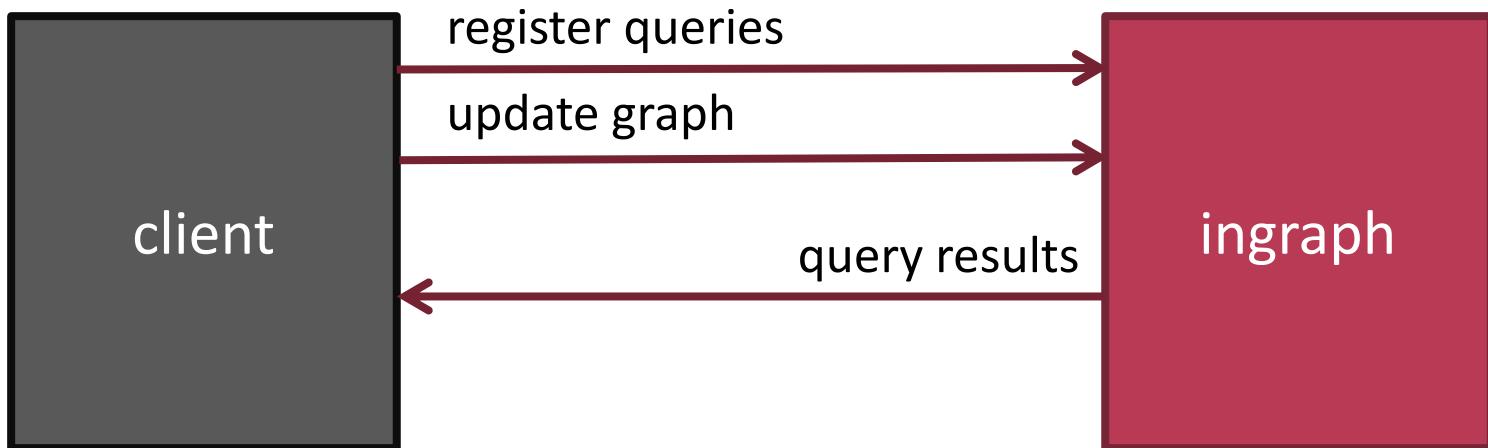
# ingraph

- An incremental, in-memory graph query engine



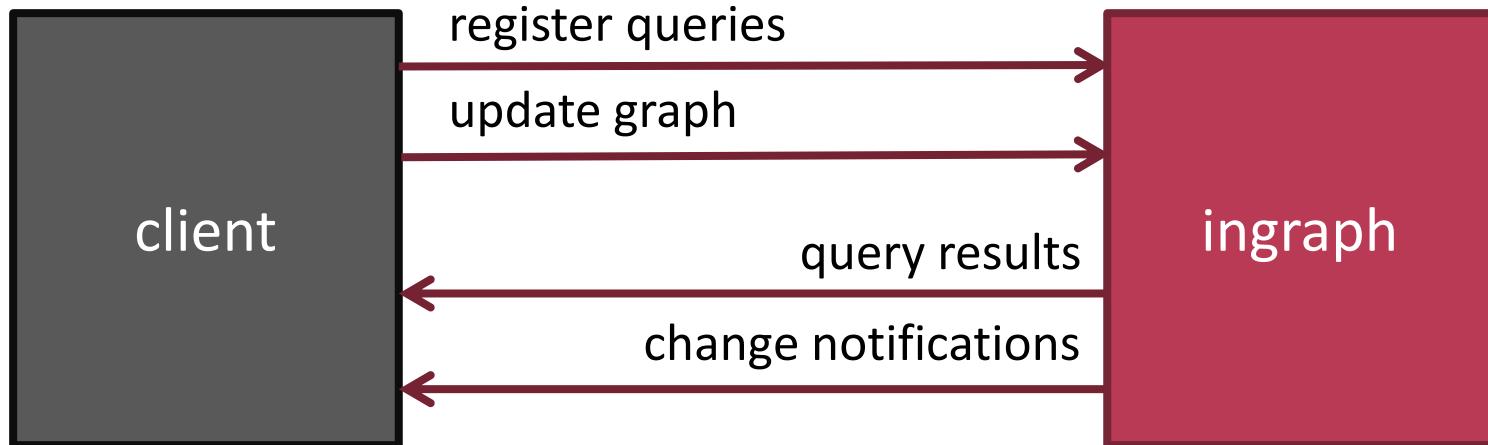
# ingraph

- An incremental, in-memory graph query engine



# ingraph

- An incremental, in-memory graph query engine



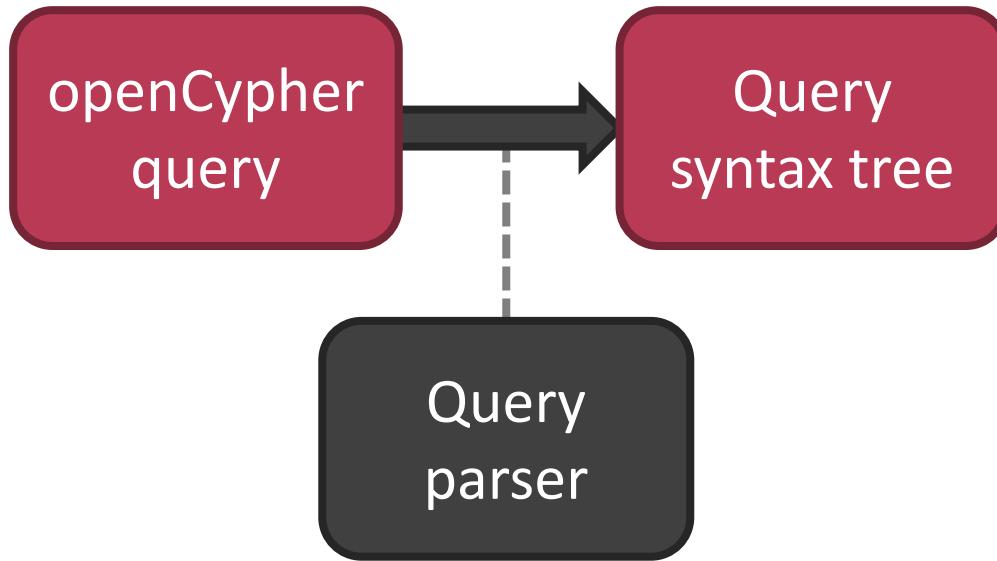
# openCypher query

```
MATCH (t:Train)-[:ON]->(seg:Segment)
      <-[:STRAIGHT]- (sw:Switch)
WHERE sw.position = 'diverging'
RETURN t.number, sw
```

openCypher  
query

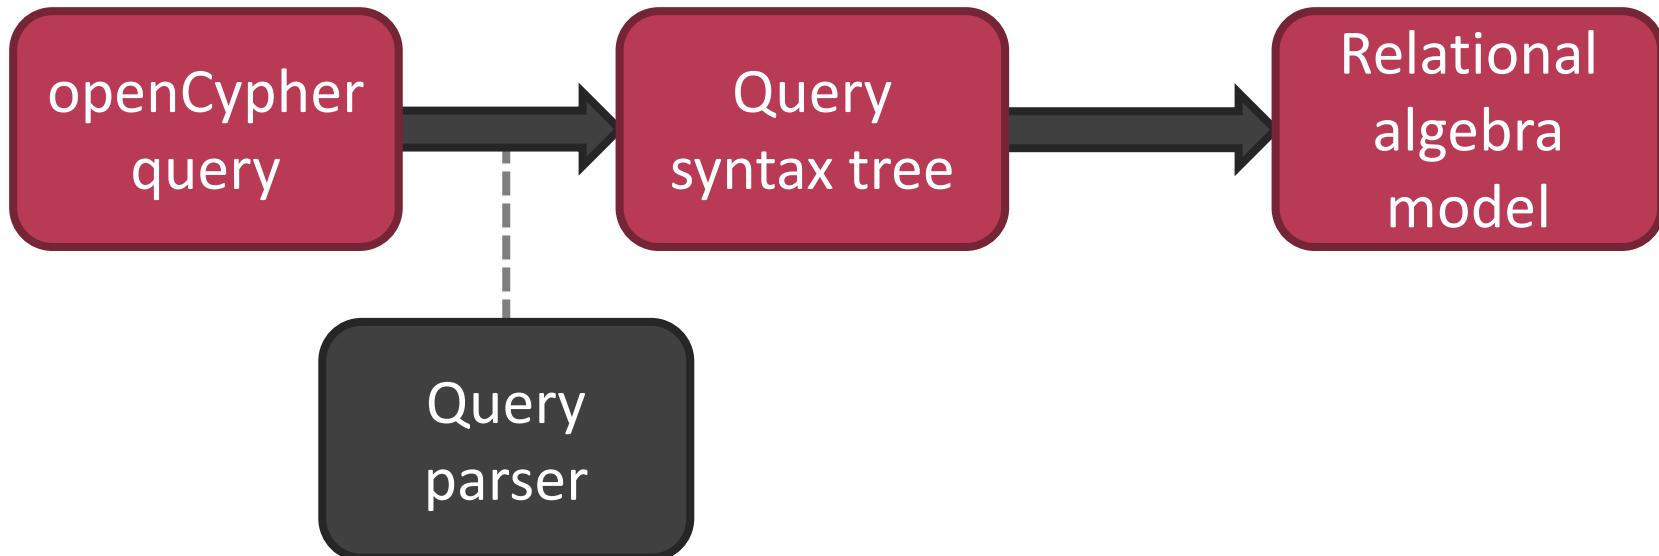
Query  
syntax tree

```
MATCH (t:Train)-[:ON]->(seg:Segment)
      <-[ :STRAIGHT ]-( sw:Switch )
WHERE sw.position = 'diverging'
RETURN t.number, sw
```



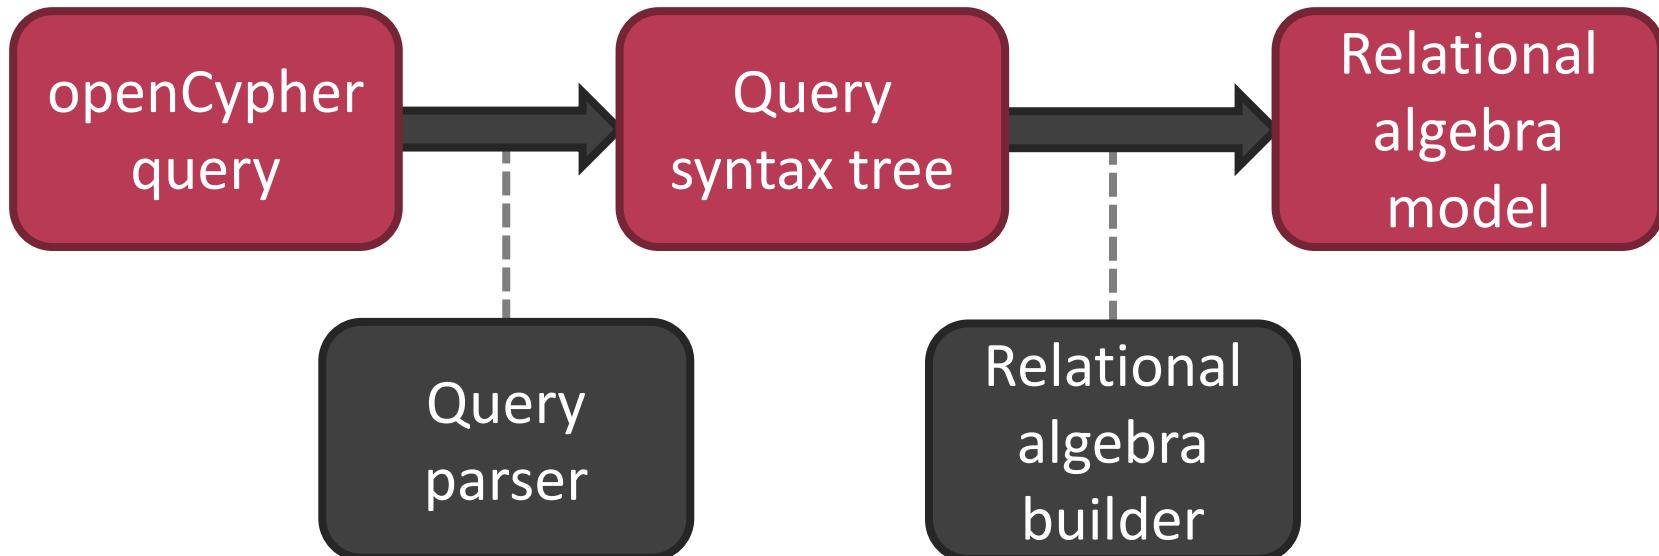
sliza  
Xtext

```
MATCH (t:Train)-[:ON]->(seg:Segment)
      <-[:STRAIGHT]->(sw:Switch)
WHERE sw.position = 'diverging'
RETURN t.number, sw
```



sliza  
Xtext

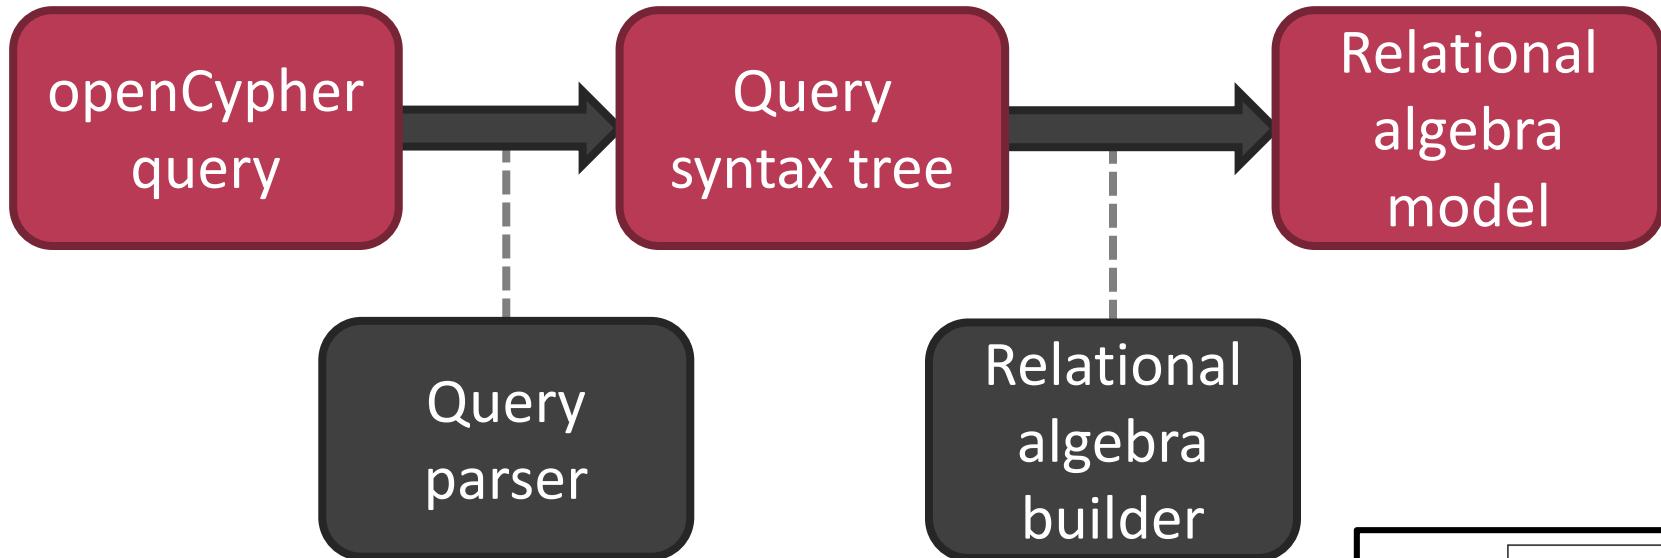
```
MATCH (t:Train)-[:ON]->(seg:Segment)
      <-[:STRAIGHT]->(sw:Switch)
WHERE sw.position = 'diverging'
RETURN t.number, sw
```



**sliza**  
**Xtext**

**Xtend**

```
MATCH (t:Train)-[:ON]->(seg:Segment)
      <-[:STRAIGHT]->(sw:Switch)
WHERE sw.position = 'diverging'
RETURN t.number, sw
```

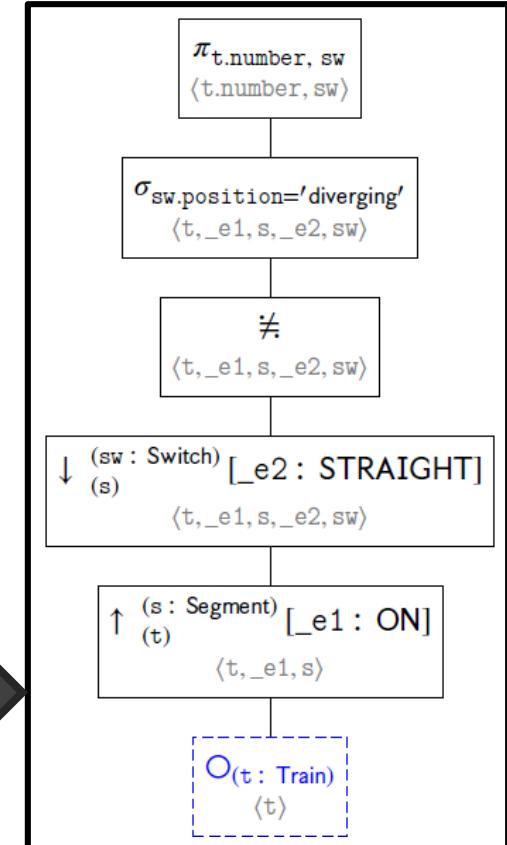


**slizaXtend**

```

MATCH (t:Train)-[:ON]->(seg:Segment)
  <-[:STRAIGHT]- (sw:Switch)
WHERE sw.position = 'diverging'
RETURN t.number, sw
  
```

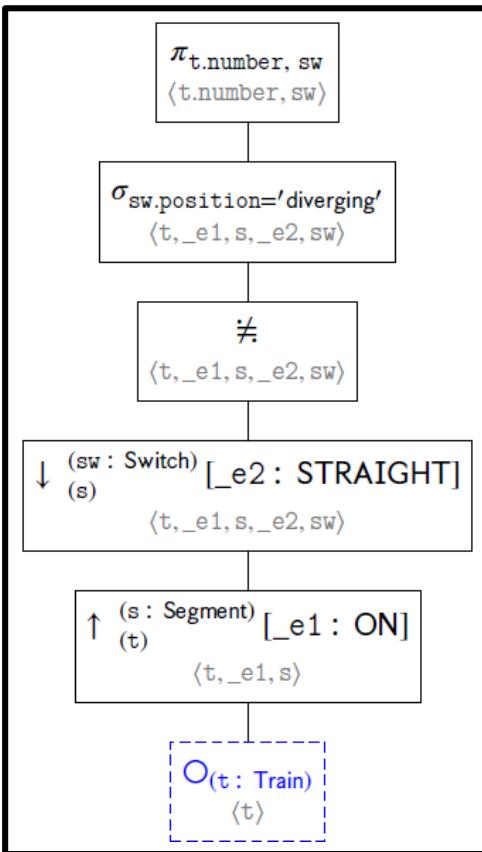
**Xtend**



Relational  
algebra  
model

Rete  
network  
model

Rete  
network

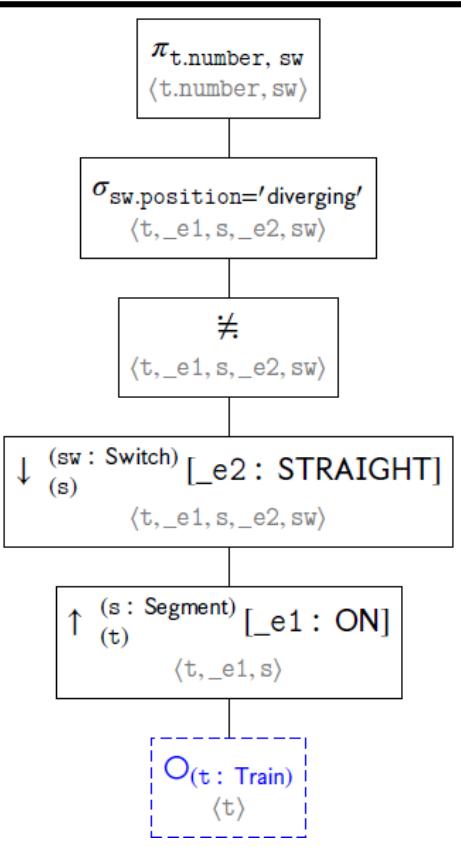


Relational  
algebra  
model

Rete  
network  
model

Rete  
network

Transformer  
and optimizer



VIATRA

Relational  
algebra  
model

Rete  
network  
model

Rete  
network

Transformer  
and optimizer

$\pi_{t.number, sw}$   
 $\langle t.number, sw \rangle$

$\sigma_{sw.position='diverging'}$   
 $\langle t, _e1, s, _e2, sw \rangle$

$\not\equiv$

$\downarrow$   
( $sw : Switch$ ) [ $_e2 : STRAIGHT$ ]  
 $\langle t, _e1, s, _e2, sw \rangle$

$\uparrow$   
( $s : Segment$ ) [ $_e1 : ON$ ]  
 $\langle t, _e1, s \rangle$

$O(t : Train)$   
 $\langle t \rangle$



VIATRA

$\pi_{t.number, sw}$   
 $\langle t.number, sw \rangle$

$\sigma_{sw.position='diverging'}$   
 $\langle t, _e7, s, sw, _e8 \rangle$

$\not\equiv$

$\bowtie$   
 $\langle t, _e7, s, sw, _e8 \rangle$

$\uparrow$   
( $s : Segment$ ) [ $_e7 : ON$ ]  
 $\langle t, _e7, s \rangle$

$\uparrow$   
( $s : Segment$ ) [ $_e8 : STRAIGHT$ ]  
 $\langle sw, _e8, s \rangle$

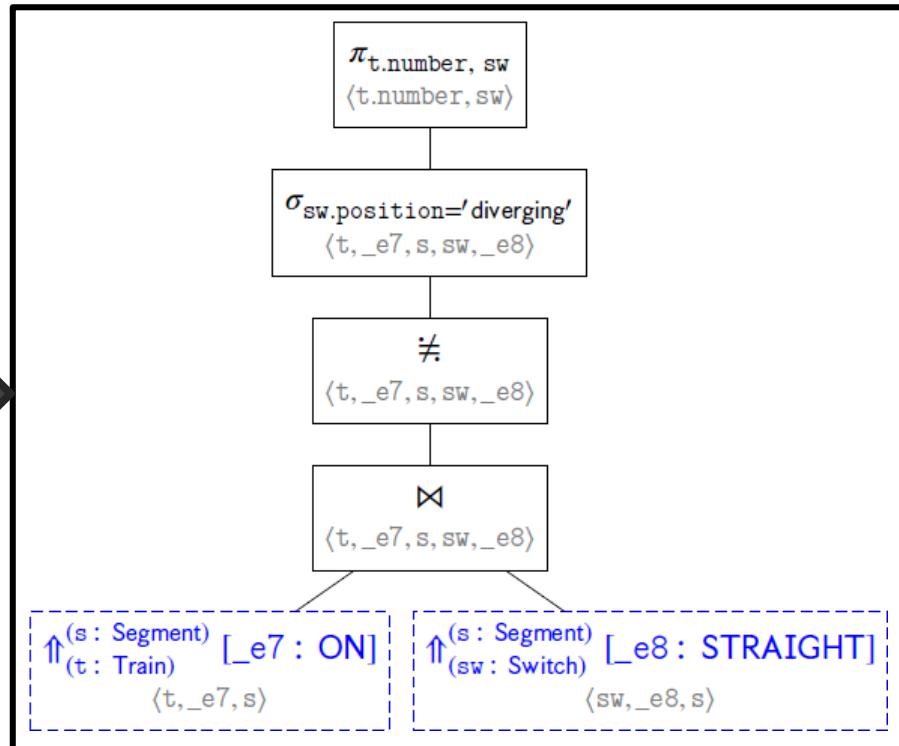
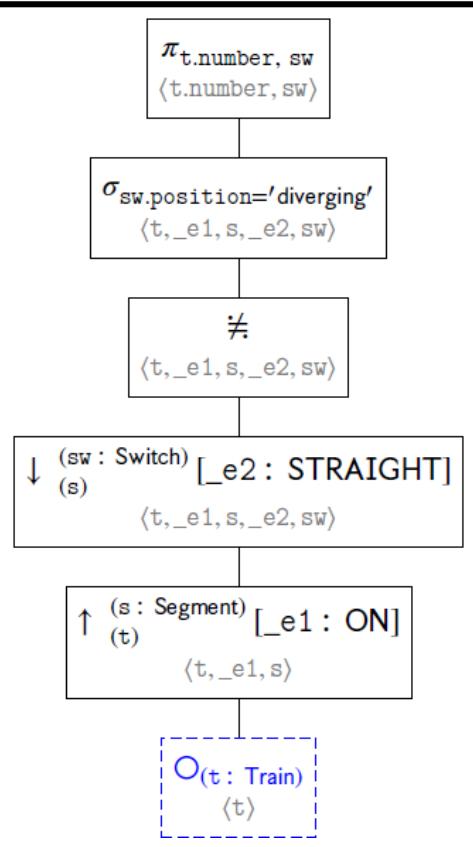
Relational  
algebra  
model

Rete  
network  
model

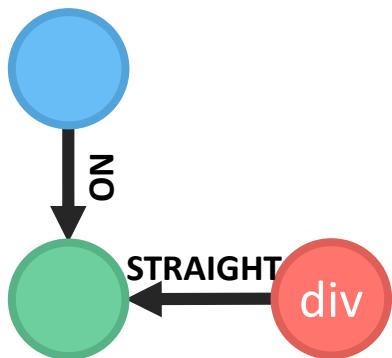
Rete  
network

Transformer  
and optimizer

Query  
deployer



## Trailing the switch



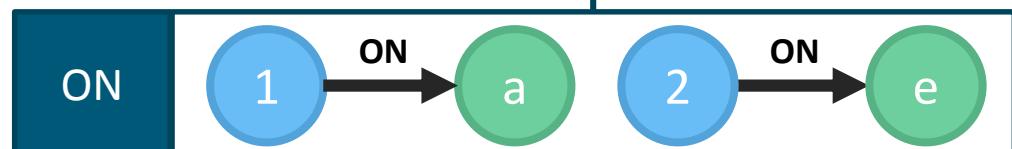
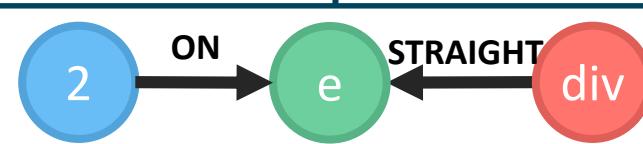
$\pi_{t.number, sw}$



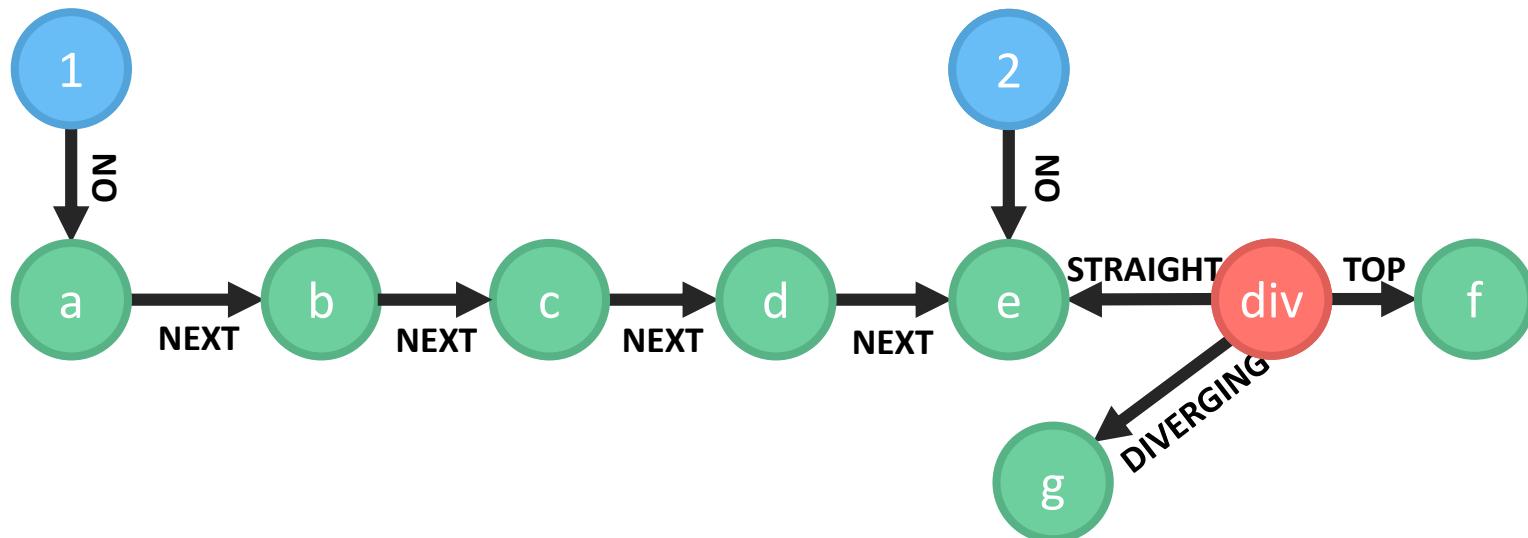
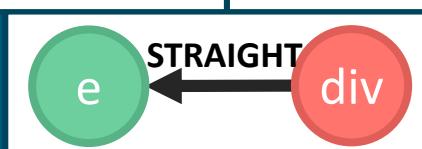
$\sigma_{sw.position} = 'diverging'$



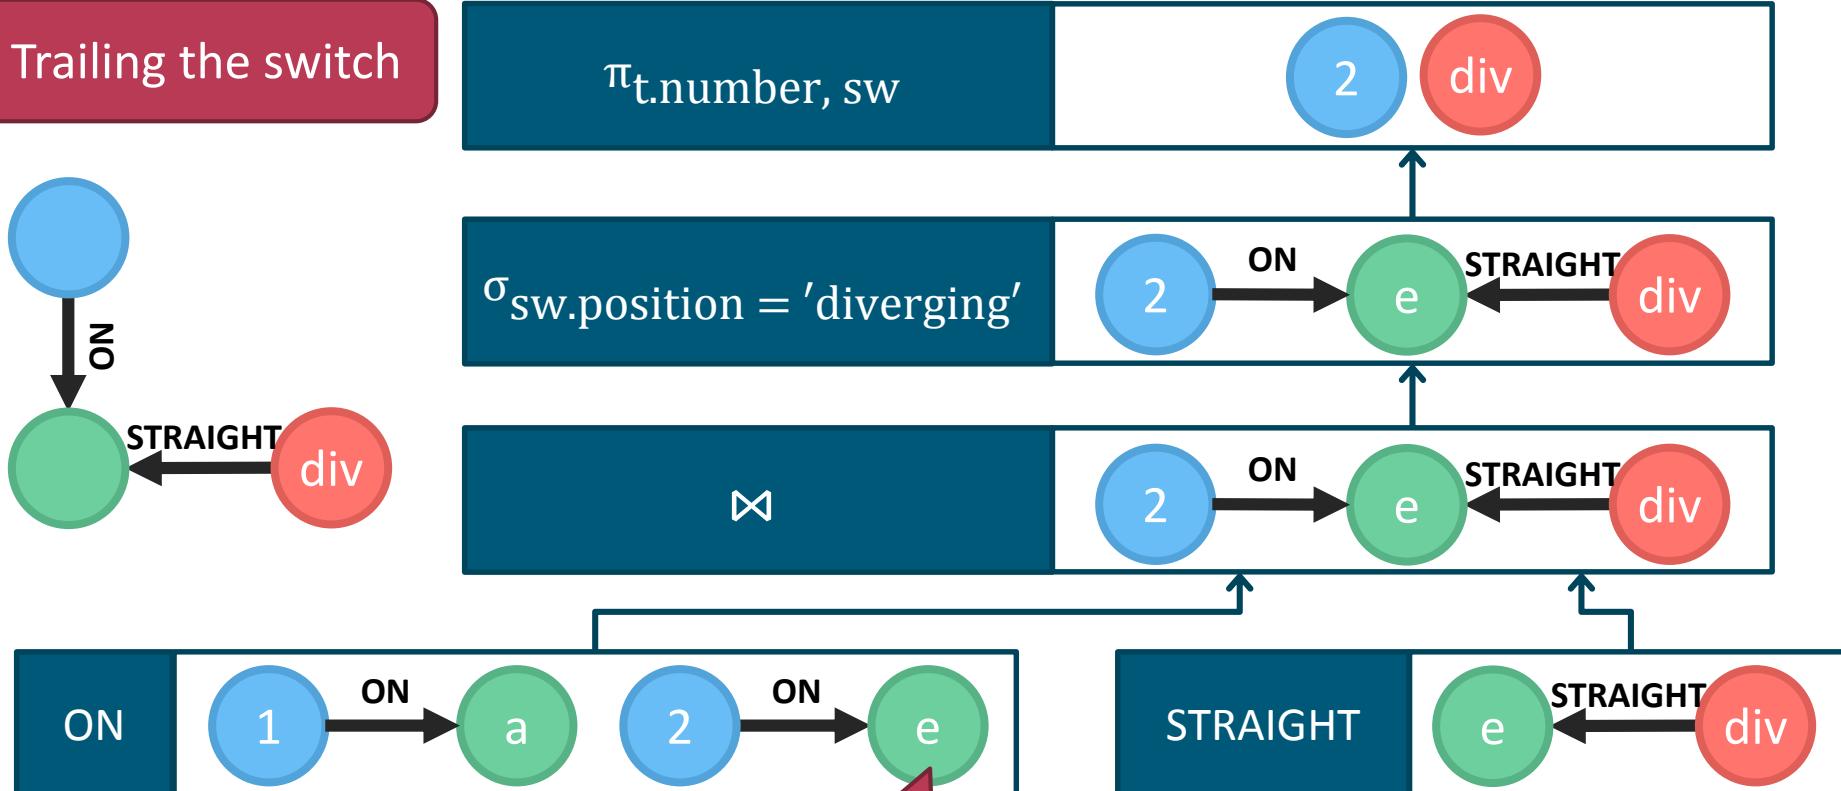
$\bowtie$



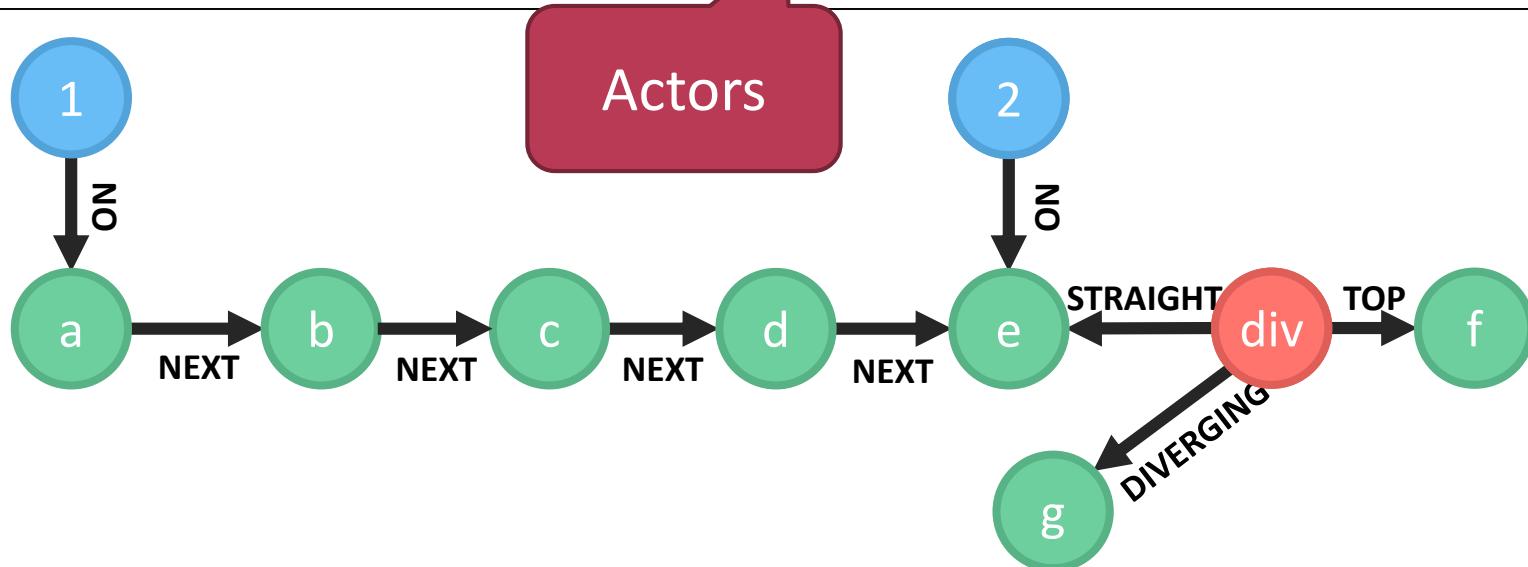
STRAIGHT



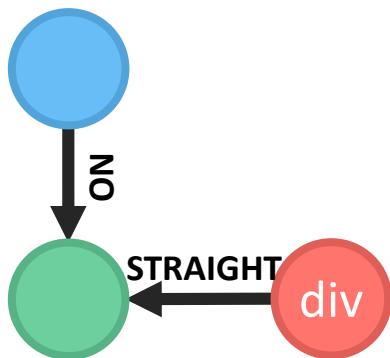
## Trailing the switch



## Actors



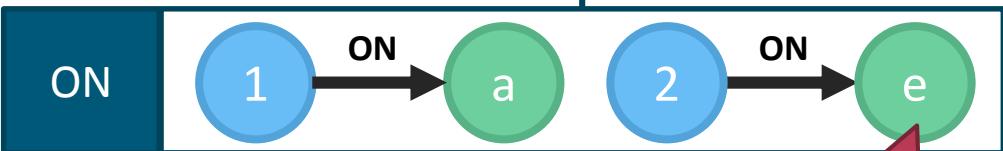
## Trailing the switch



$\pi_{t.number, sw}$

2 div

$\sigma_{sw.position} = 'diverging'$



Actors

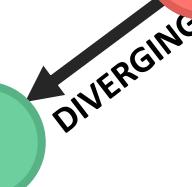
NEXT

NEXT

NEXT

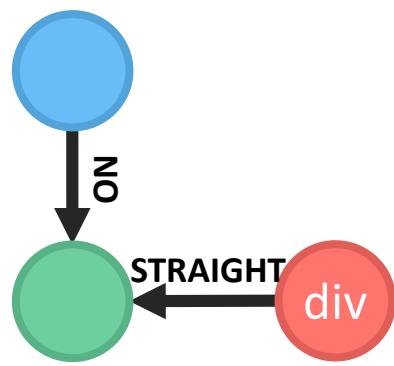
NEXT

NEXT



STRAIGHT → div → TOP → f

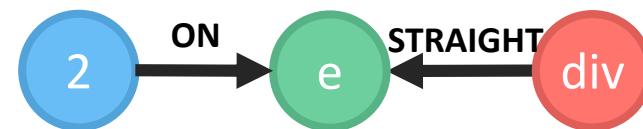
## Trailing the switch



$\pi_{t.number, sw}$



$\sigma_{sw.position} = 'diverging'$



Challenge:  
properties



Async messages



Actors

NEXT

NEXT

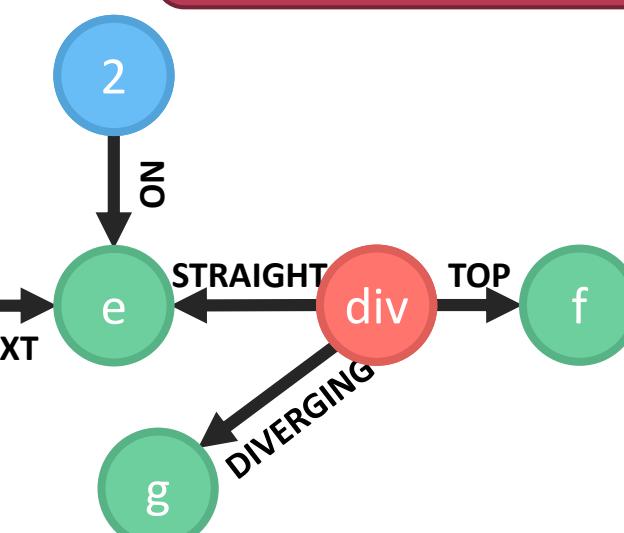
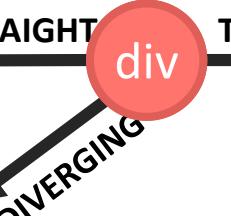
NEXT

NEXT

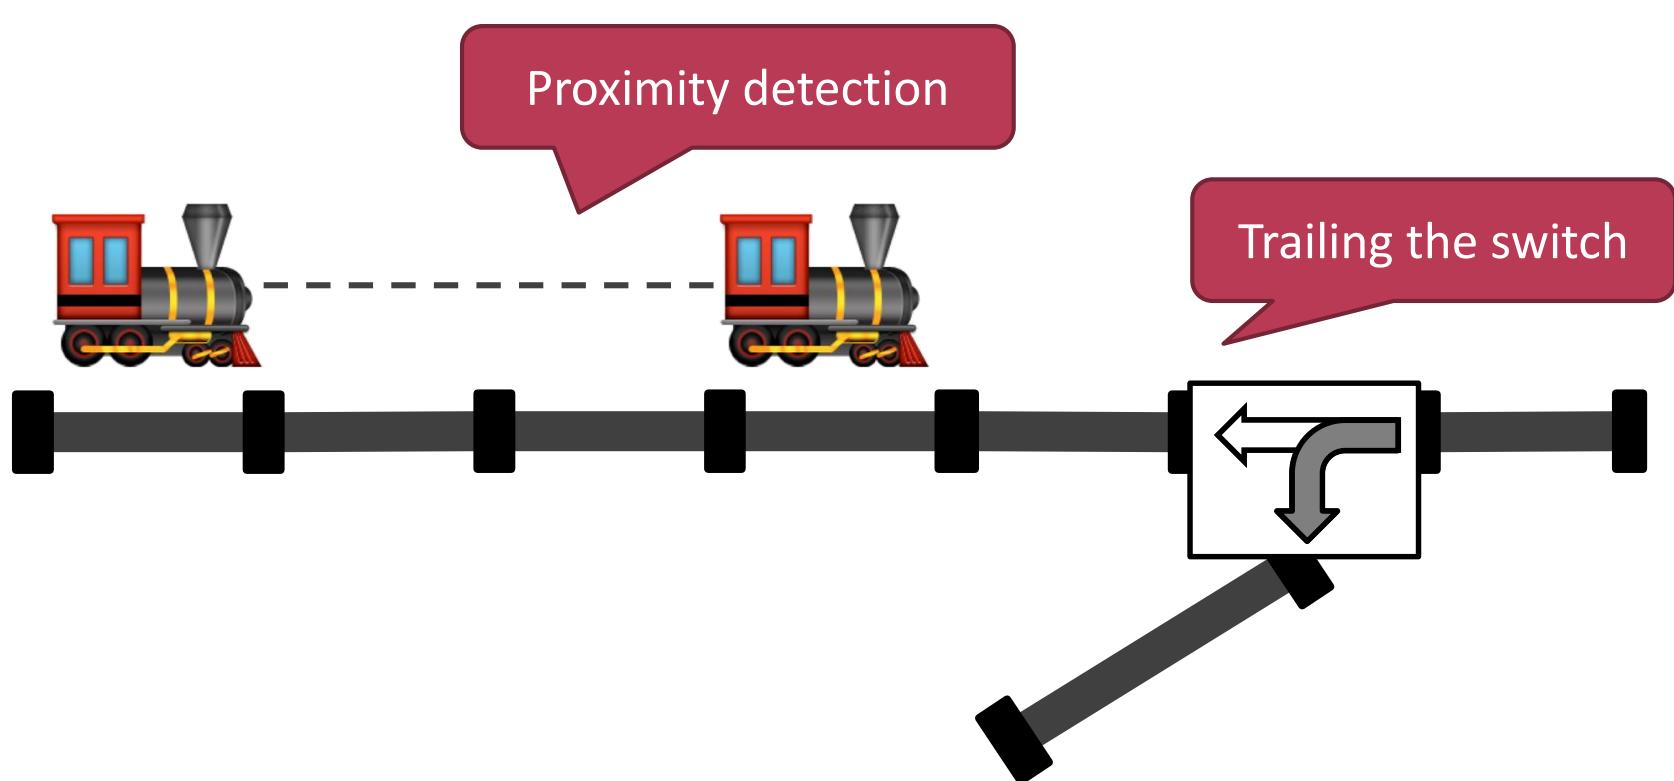
2

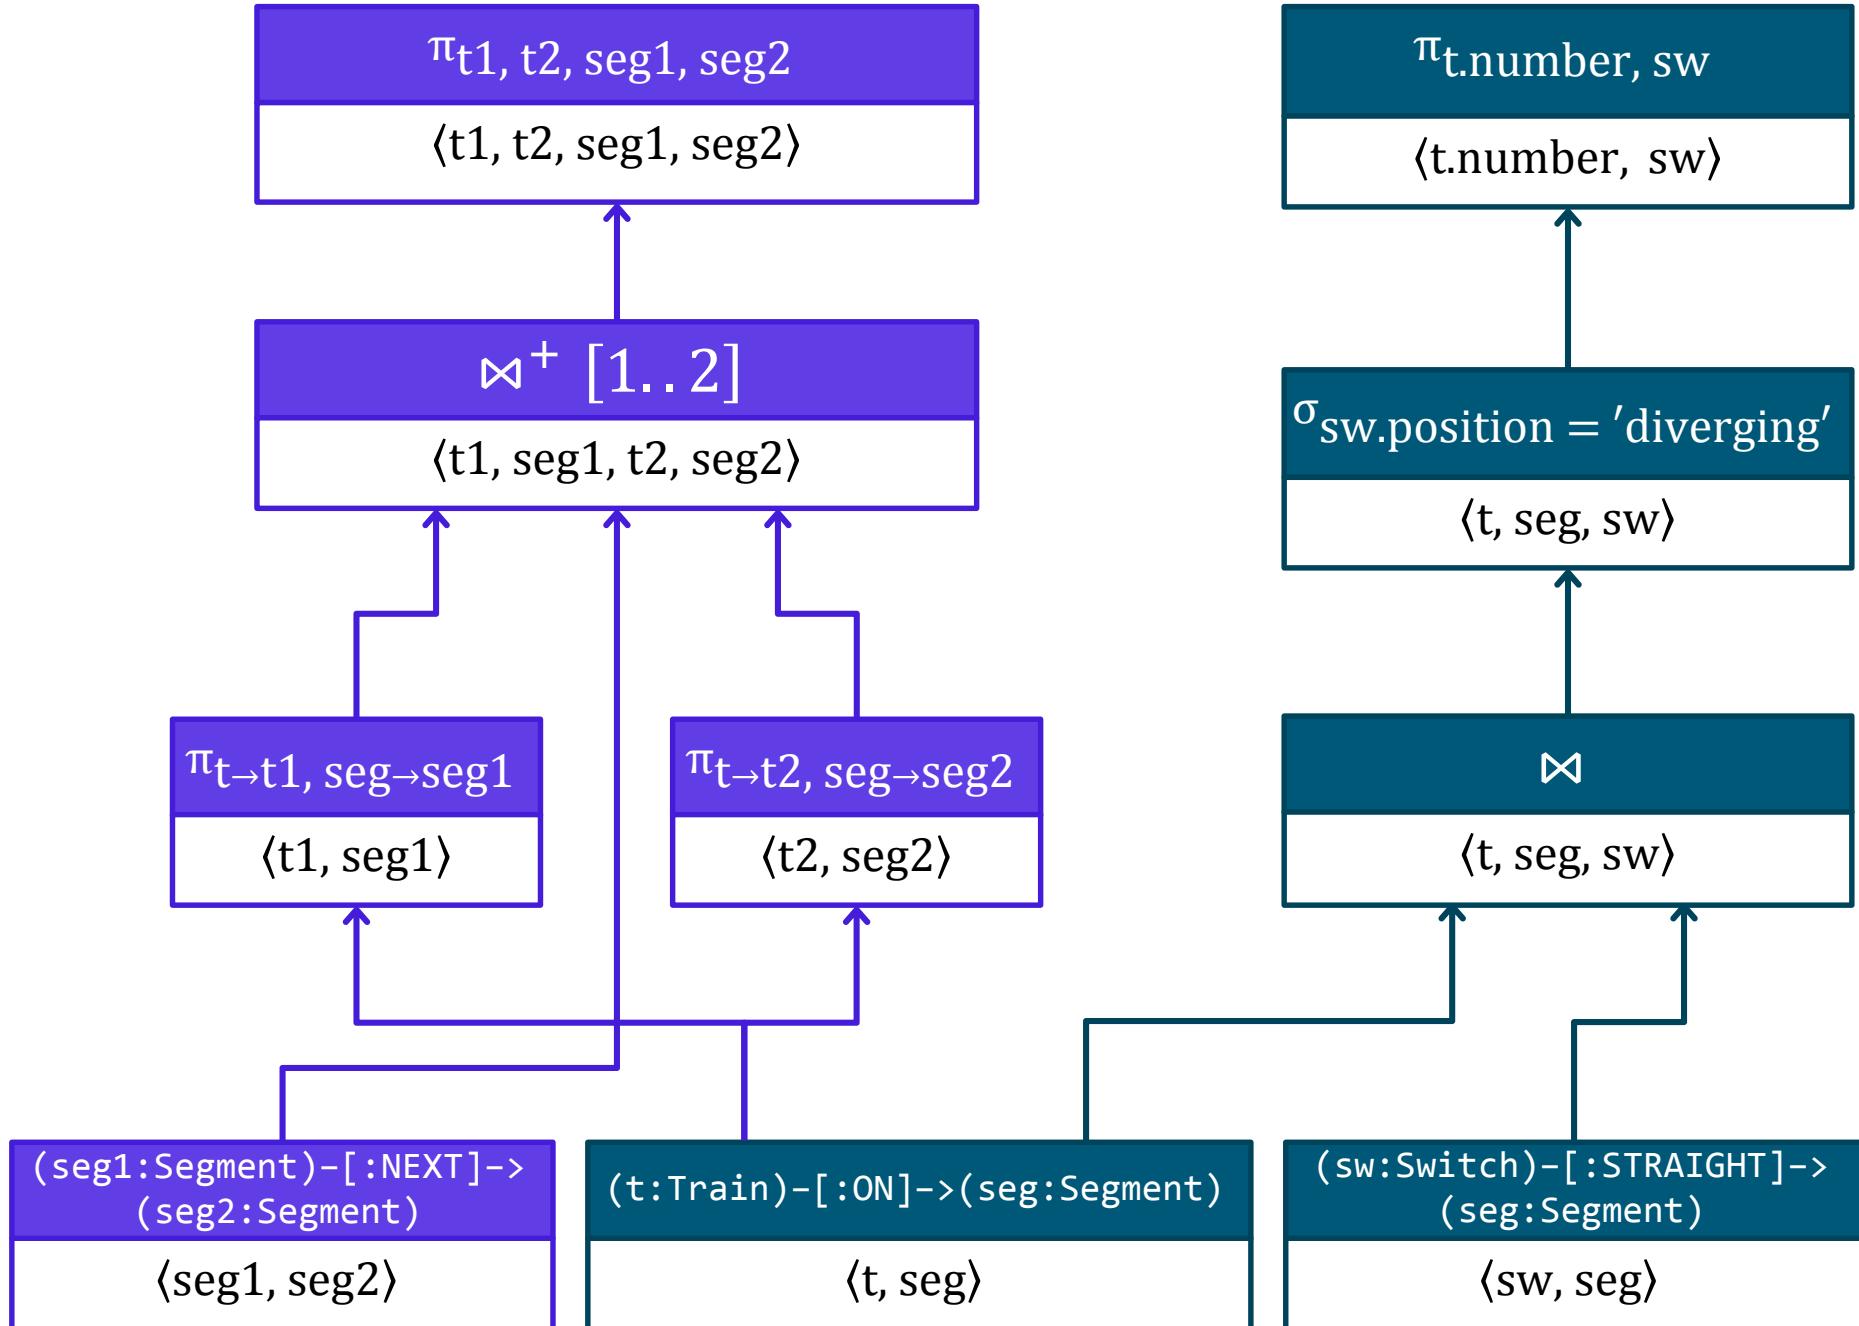
ON

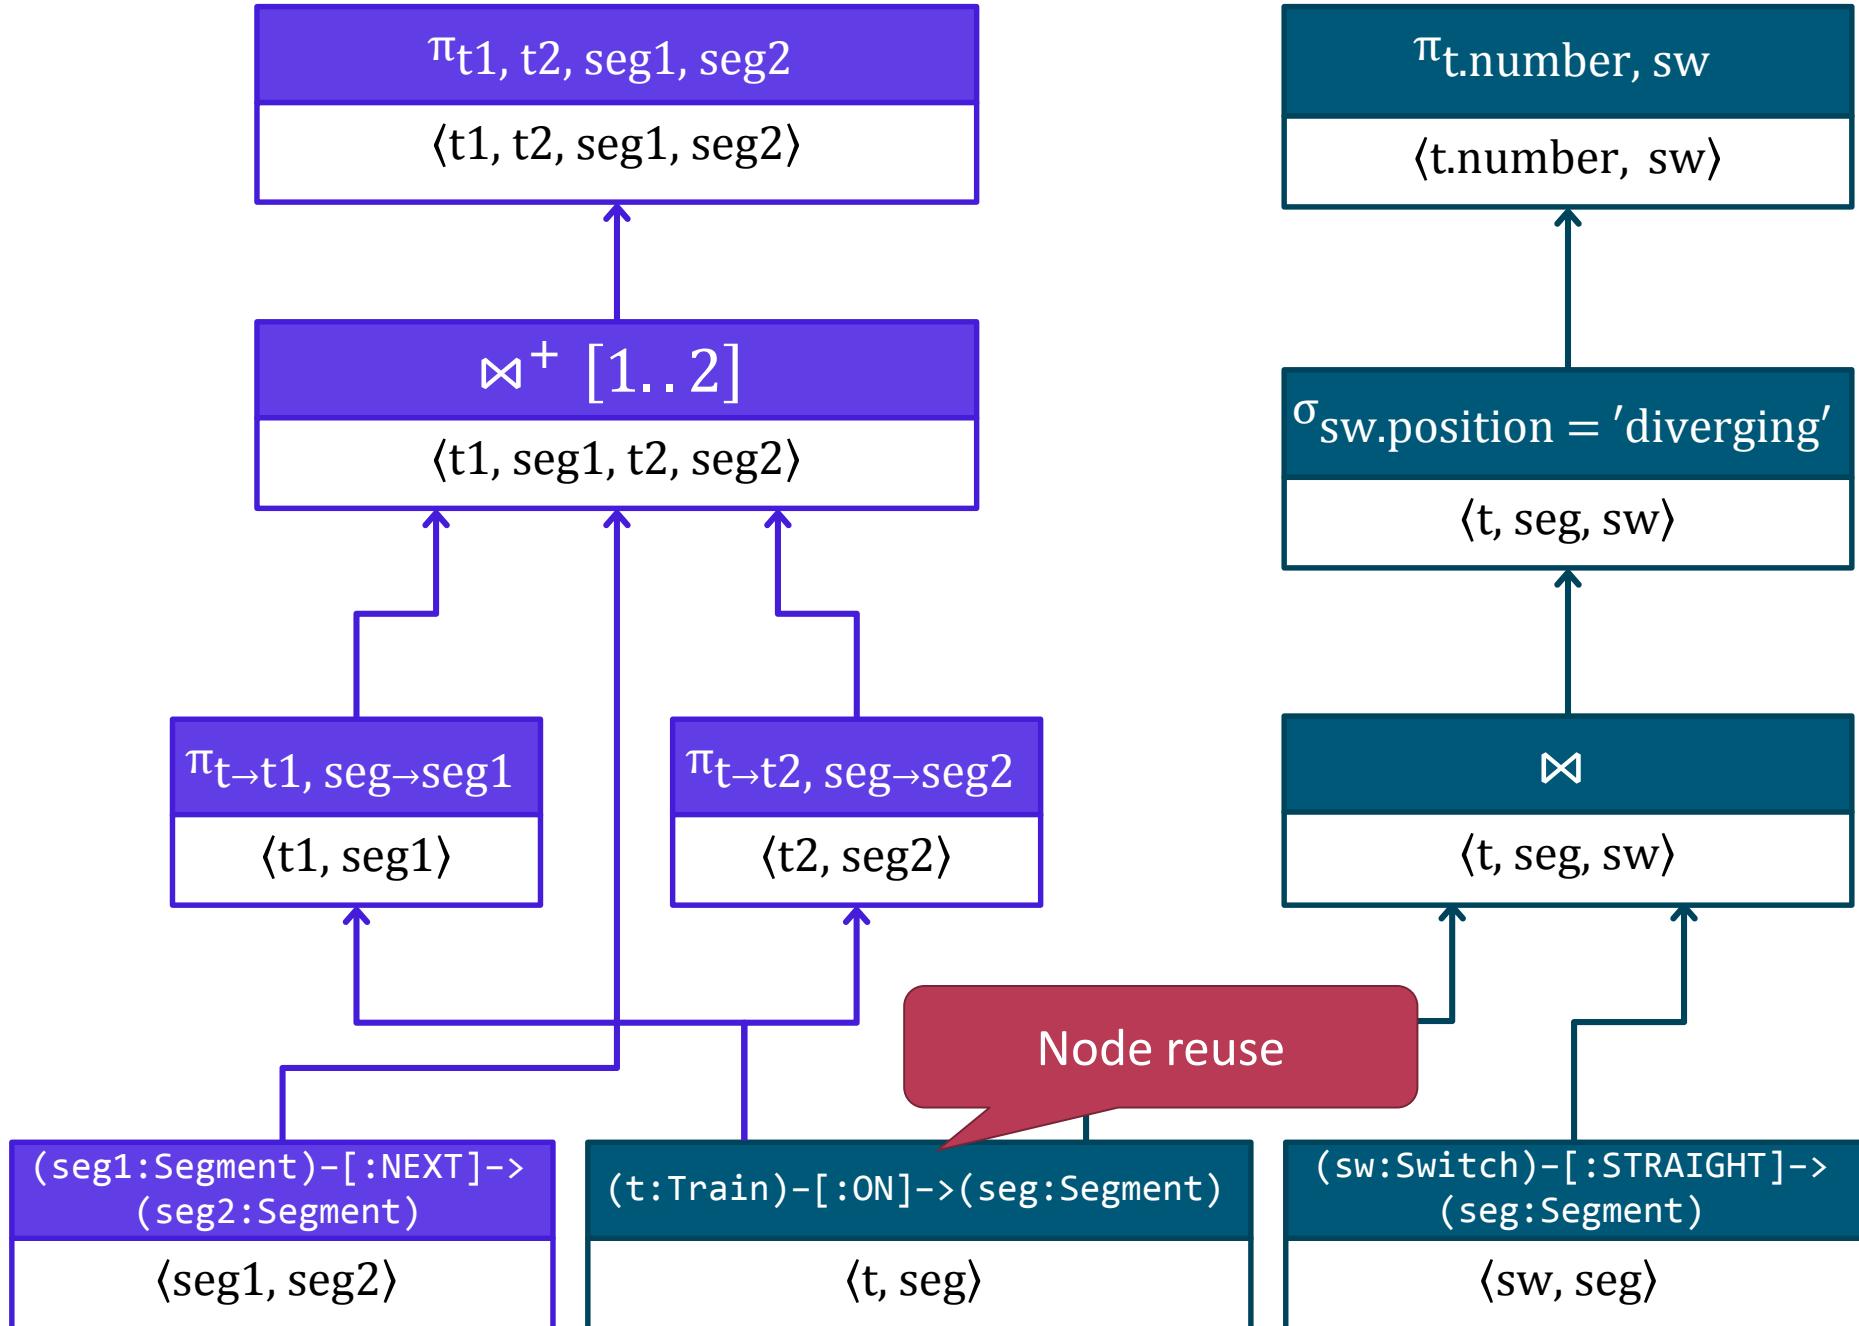
g

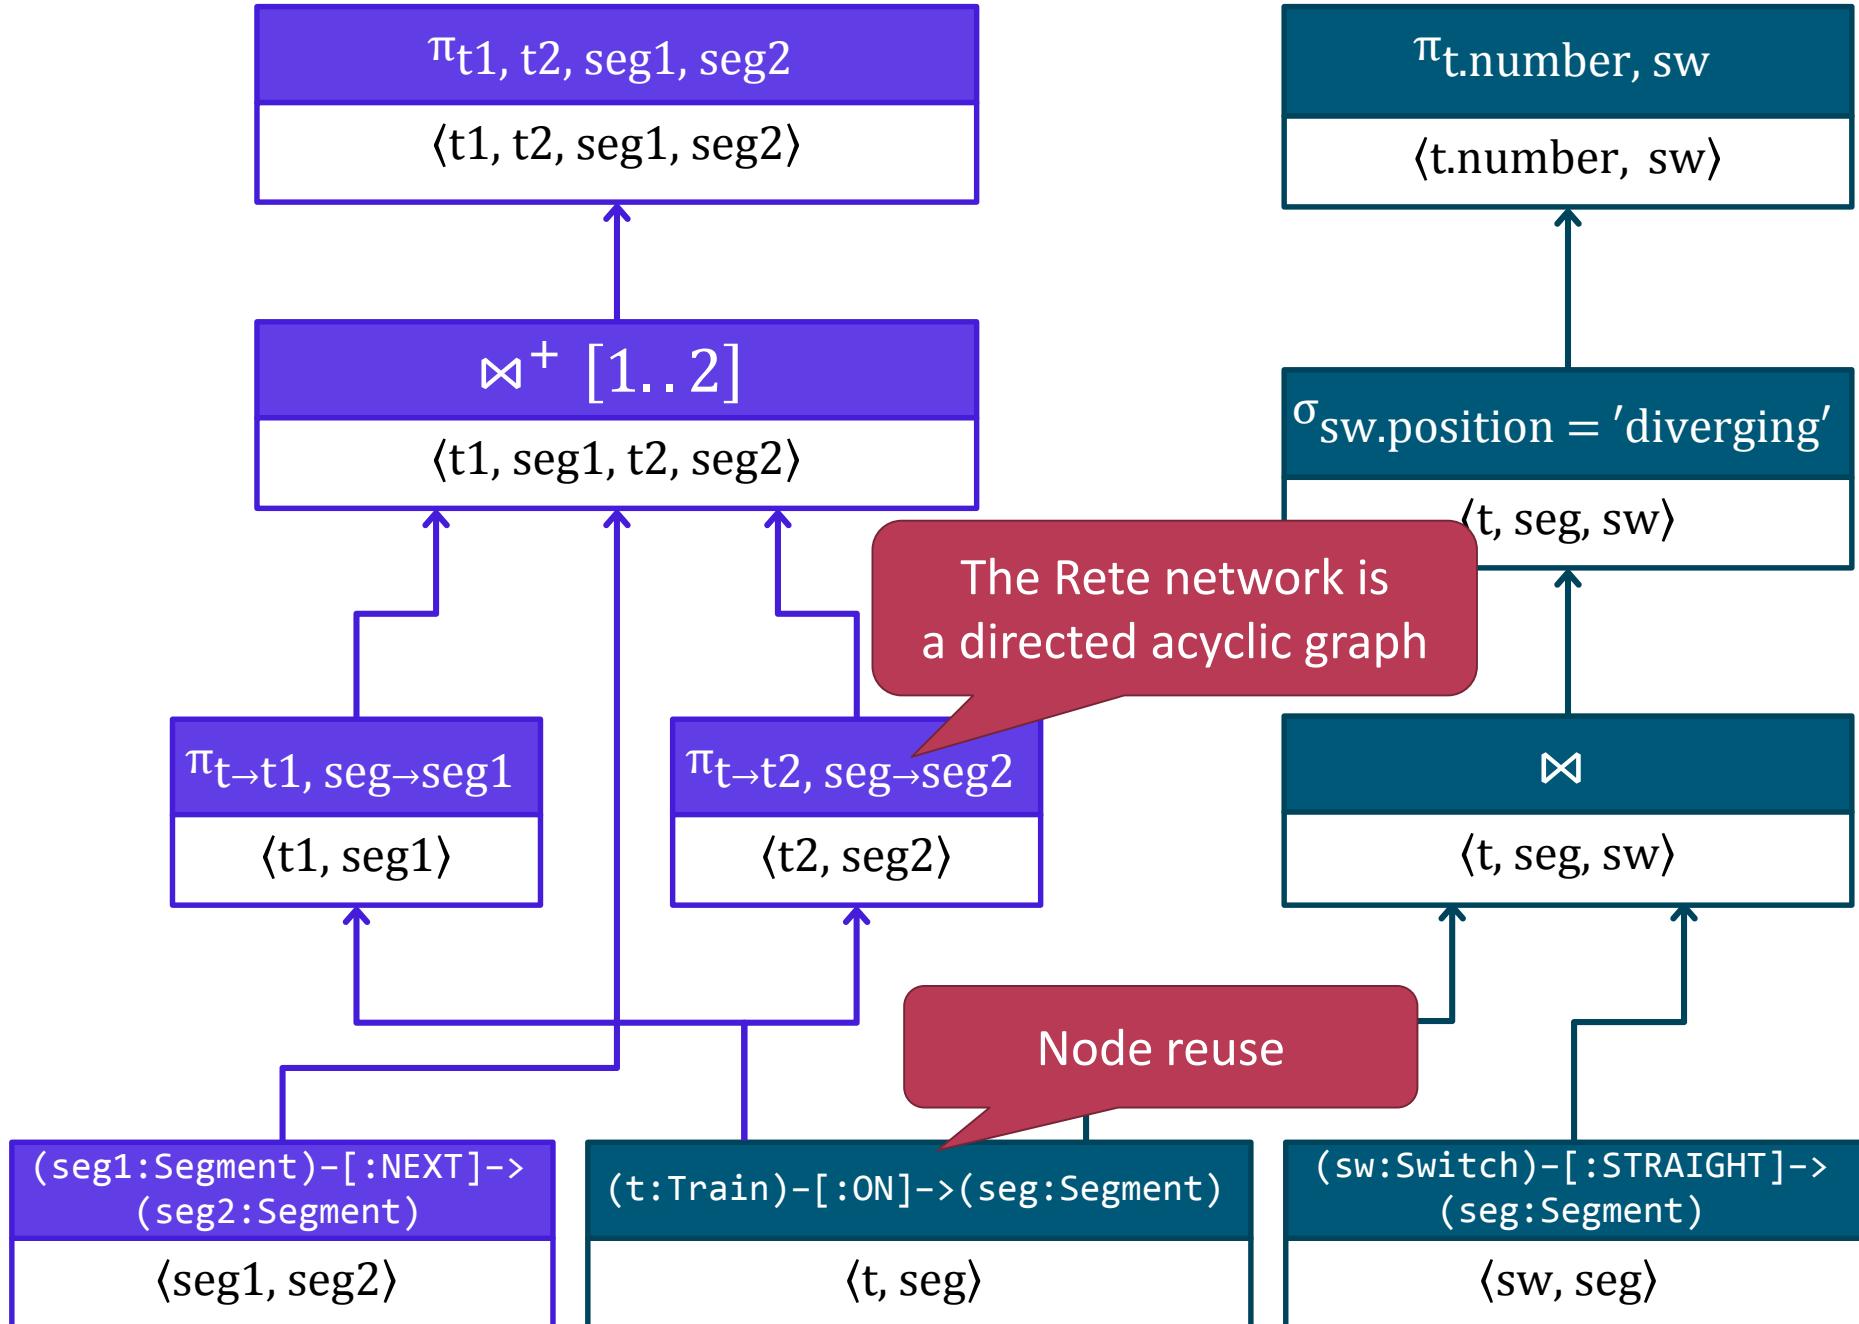


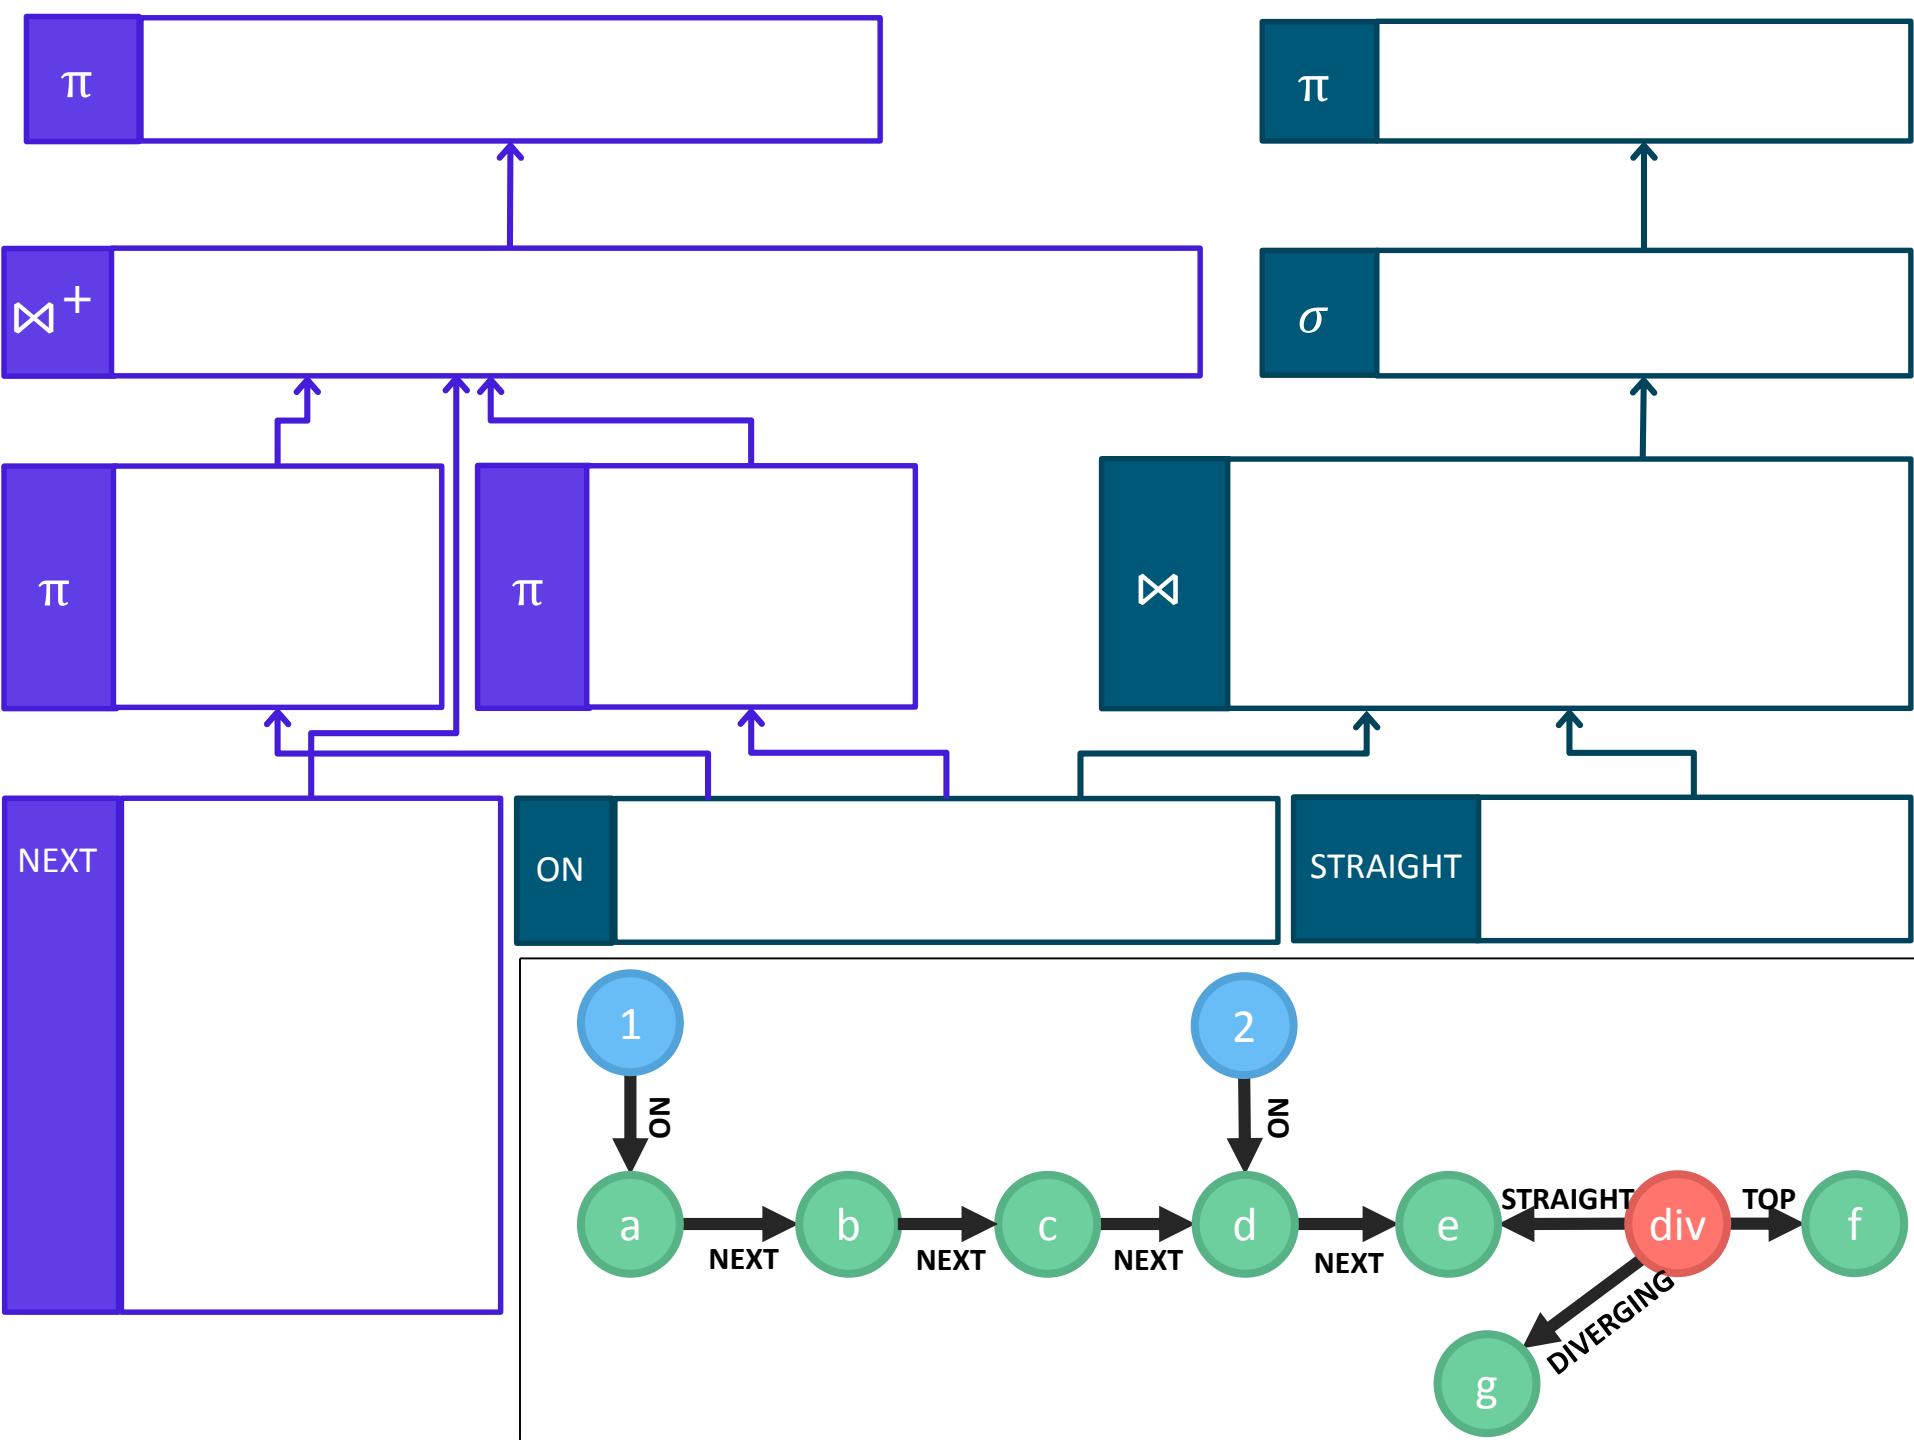
# Live railway model (revisited)

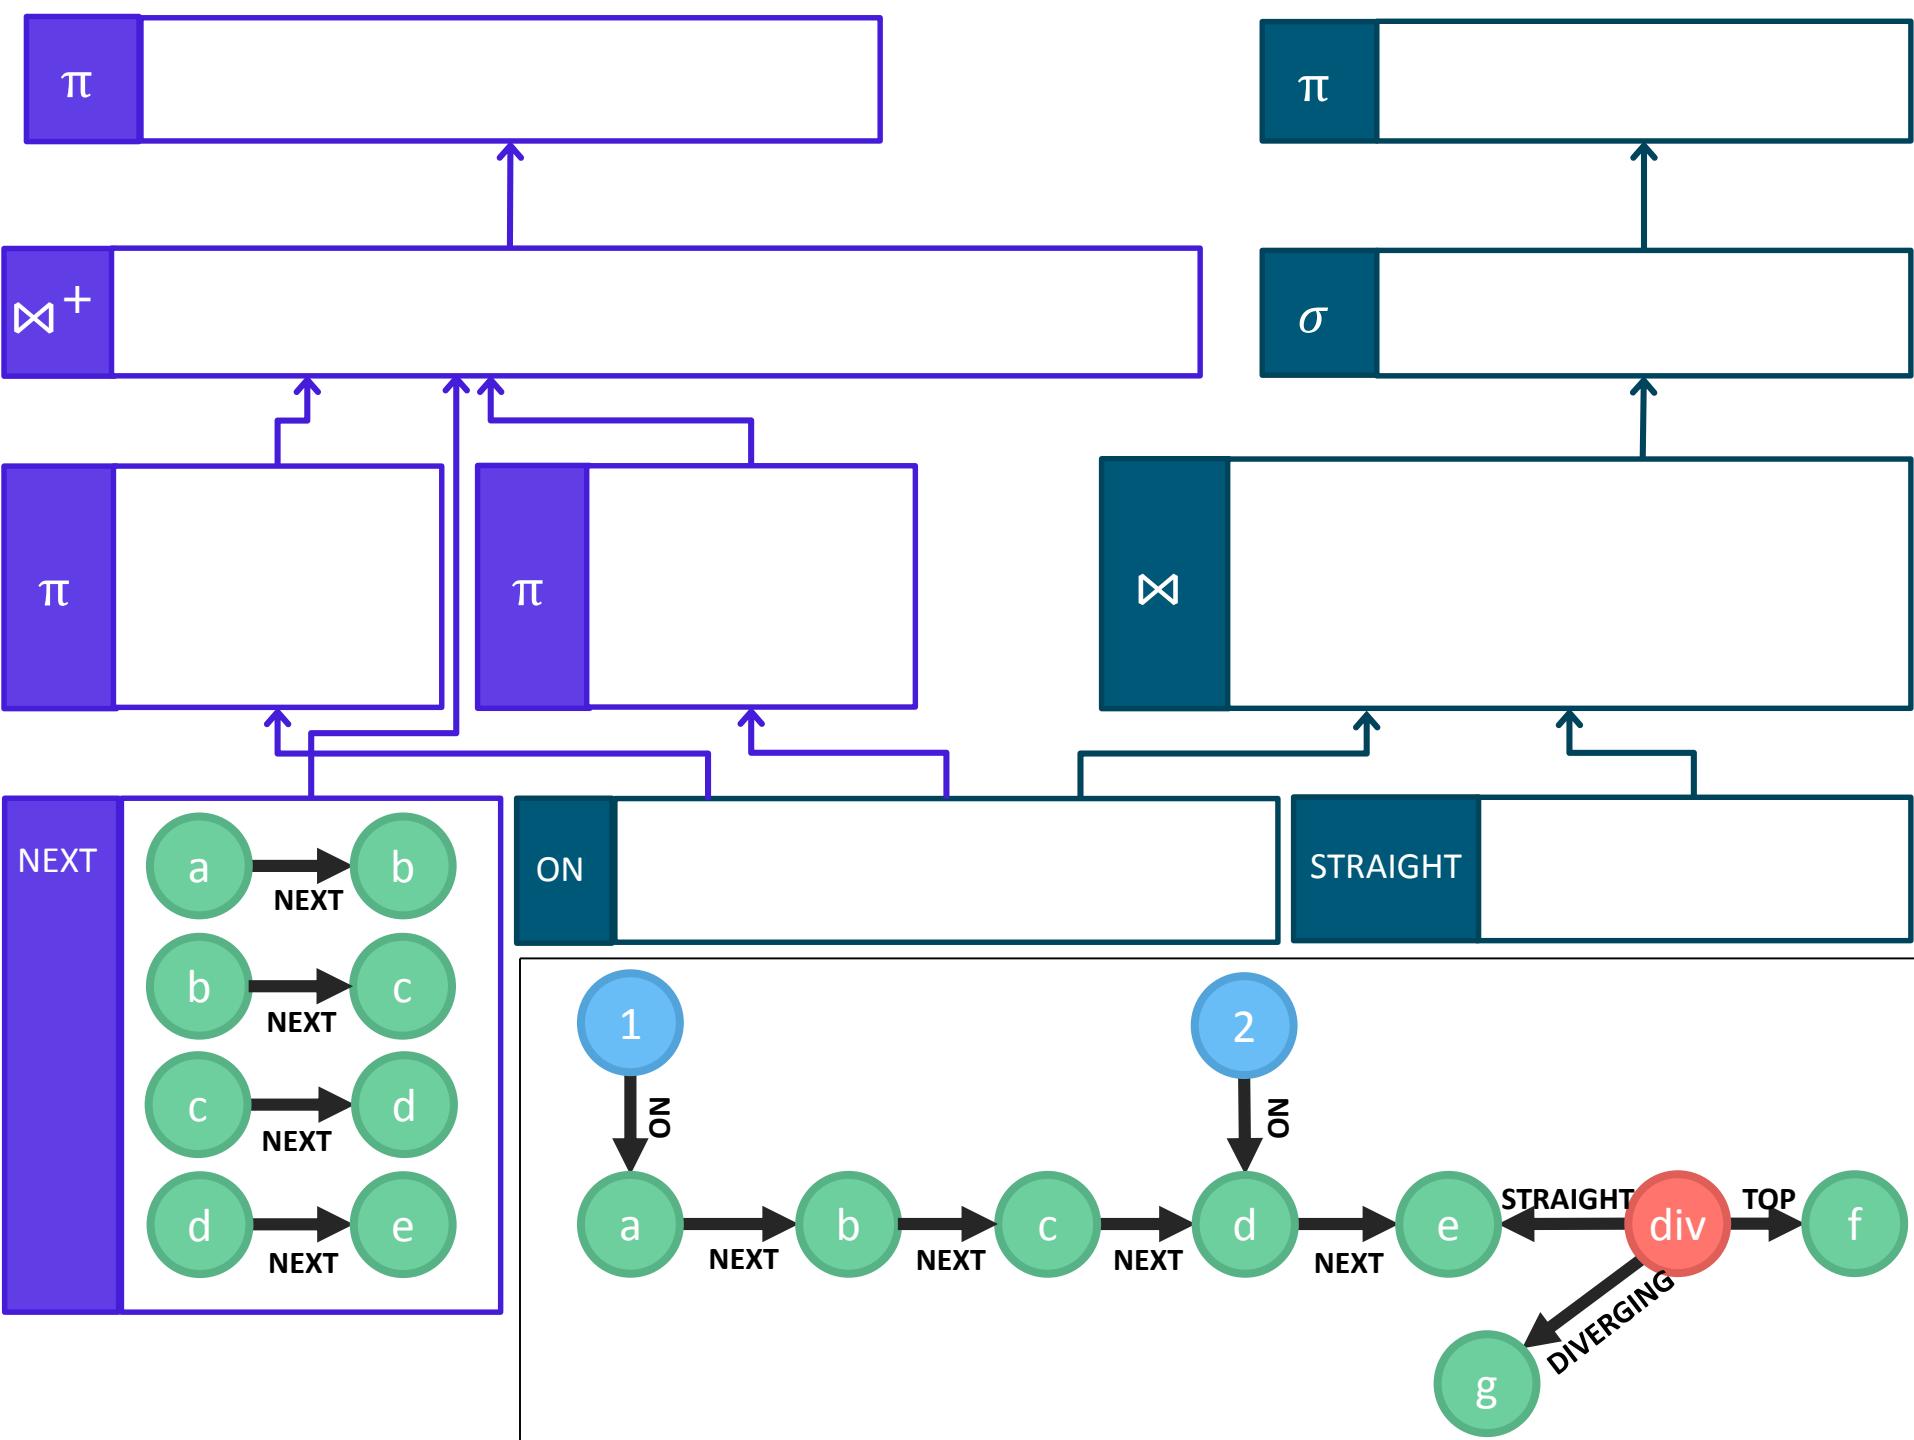


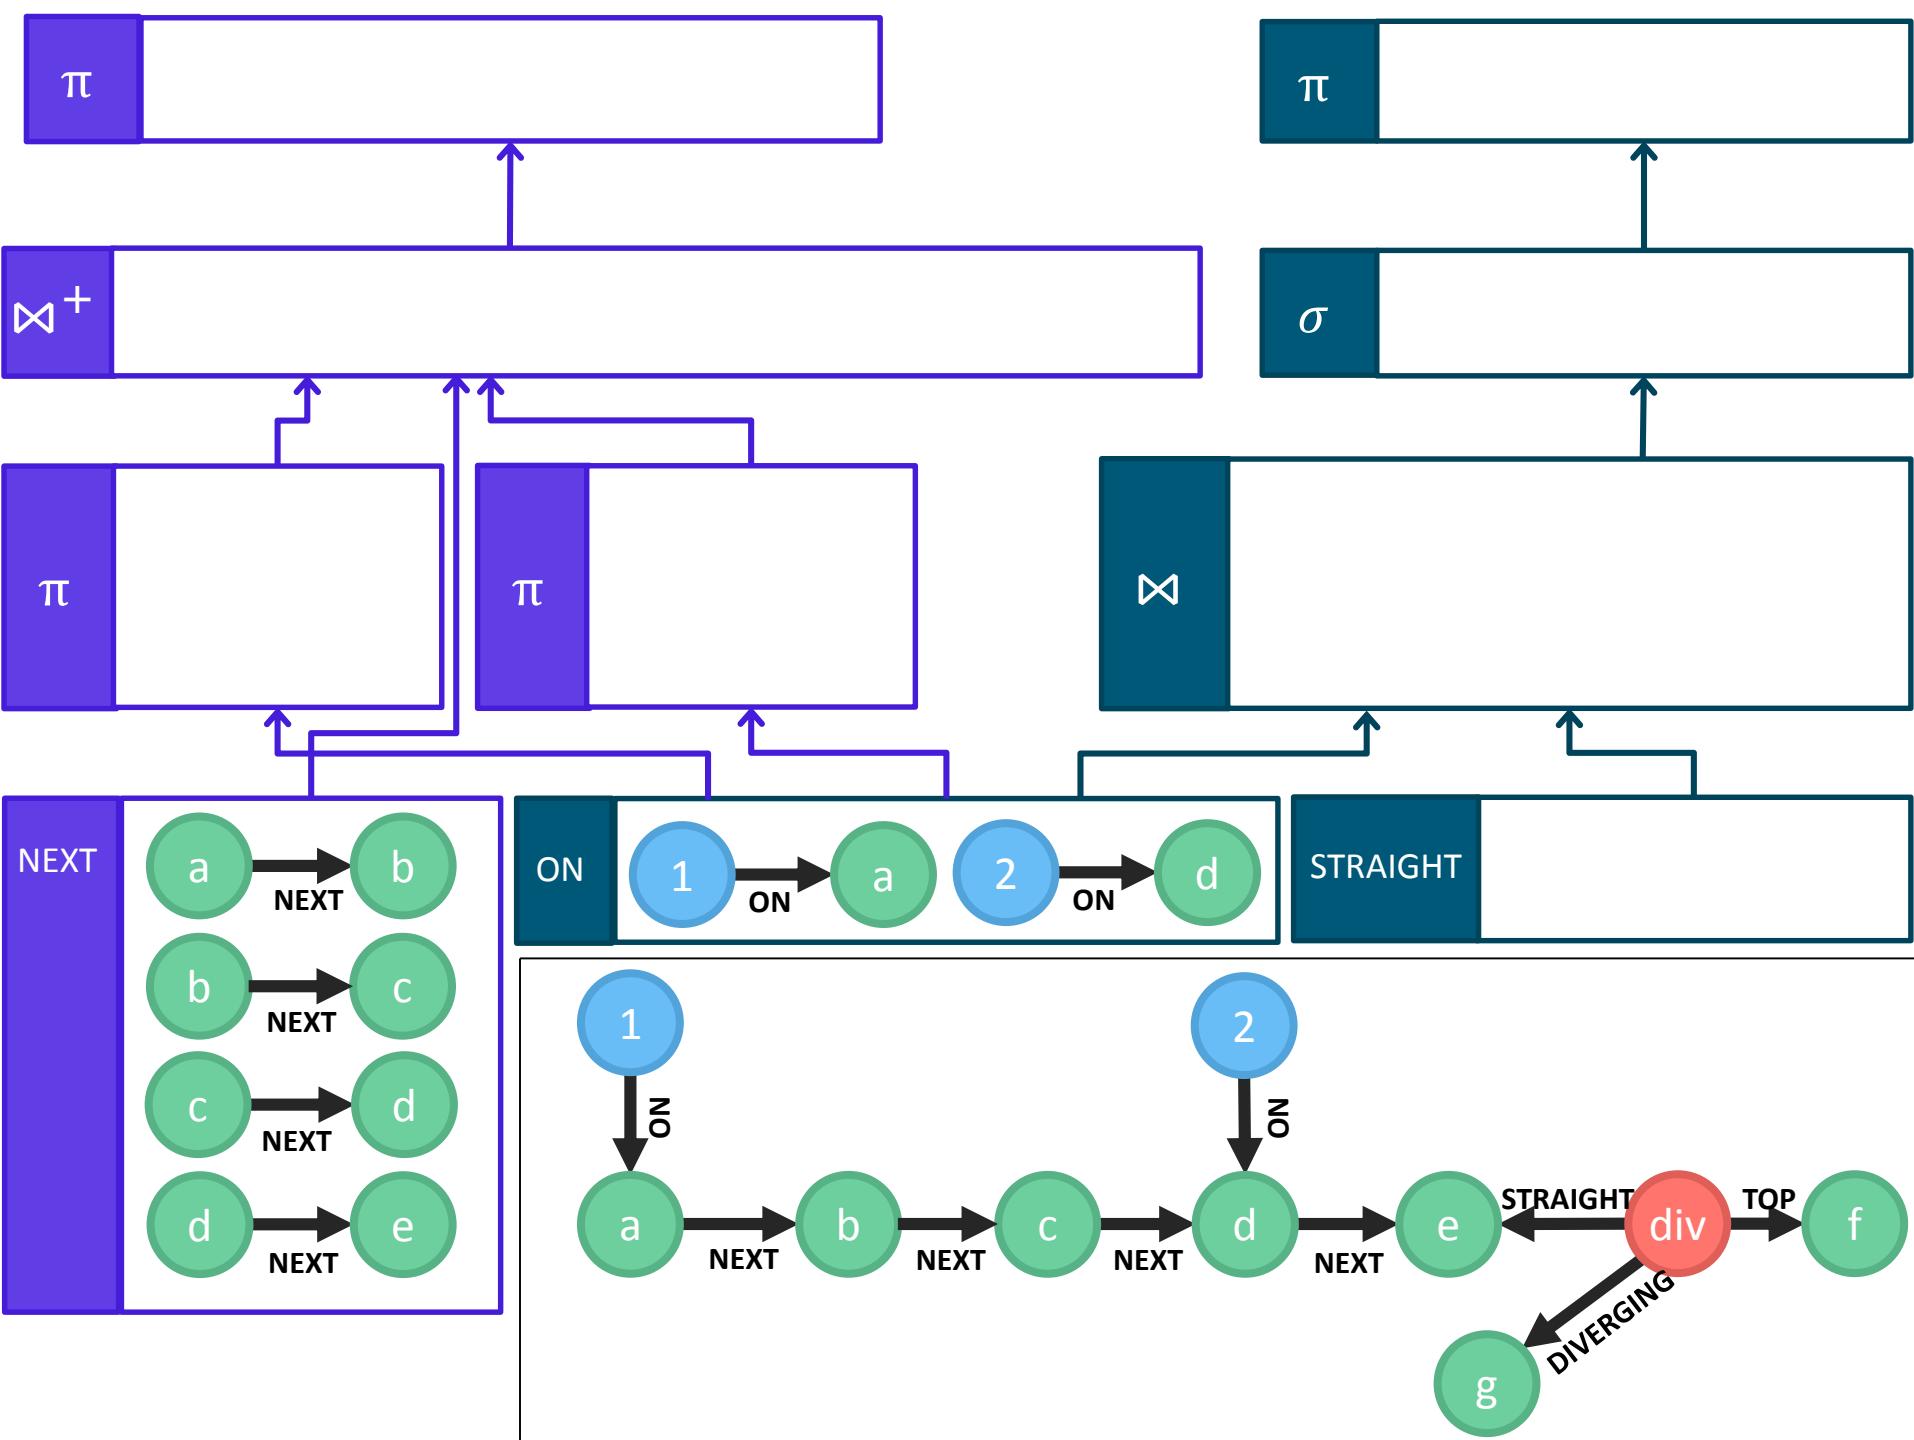


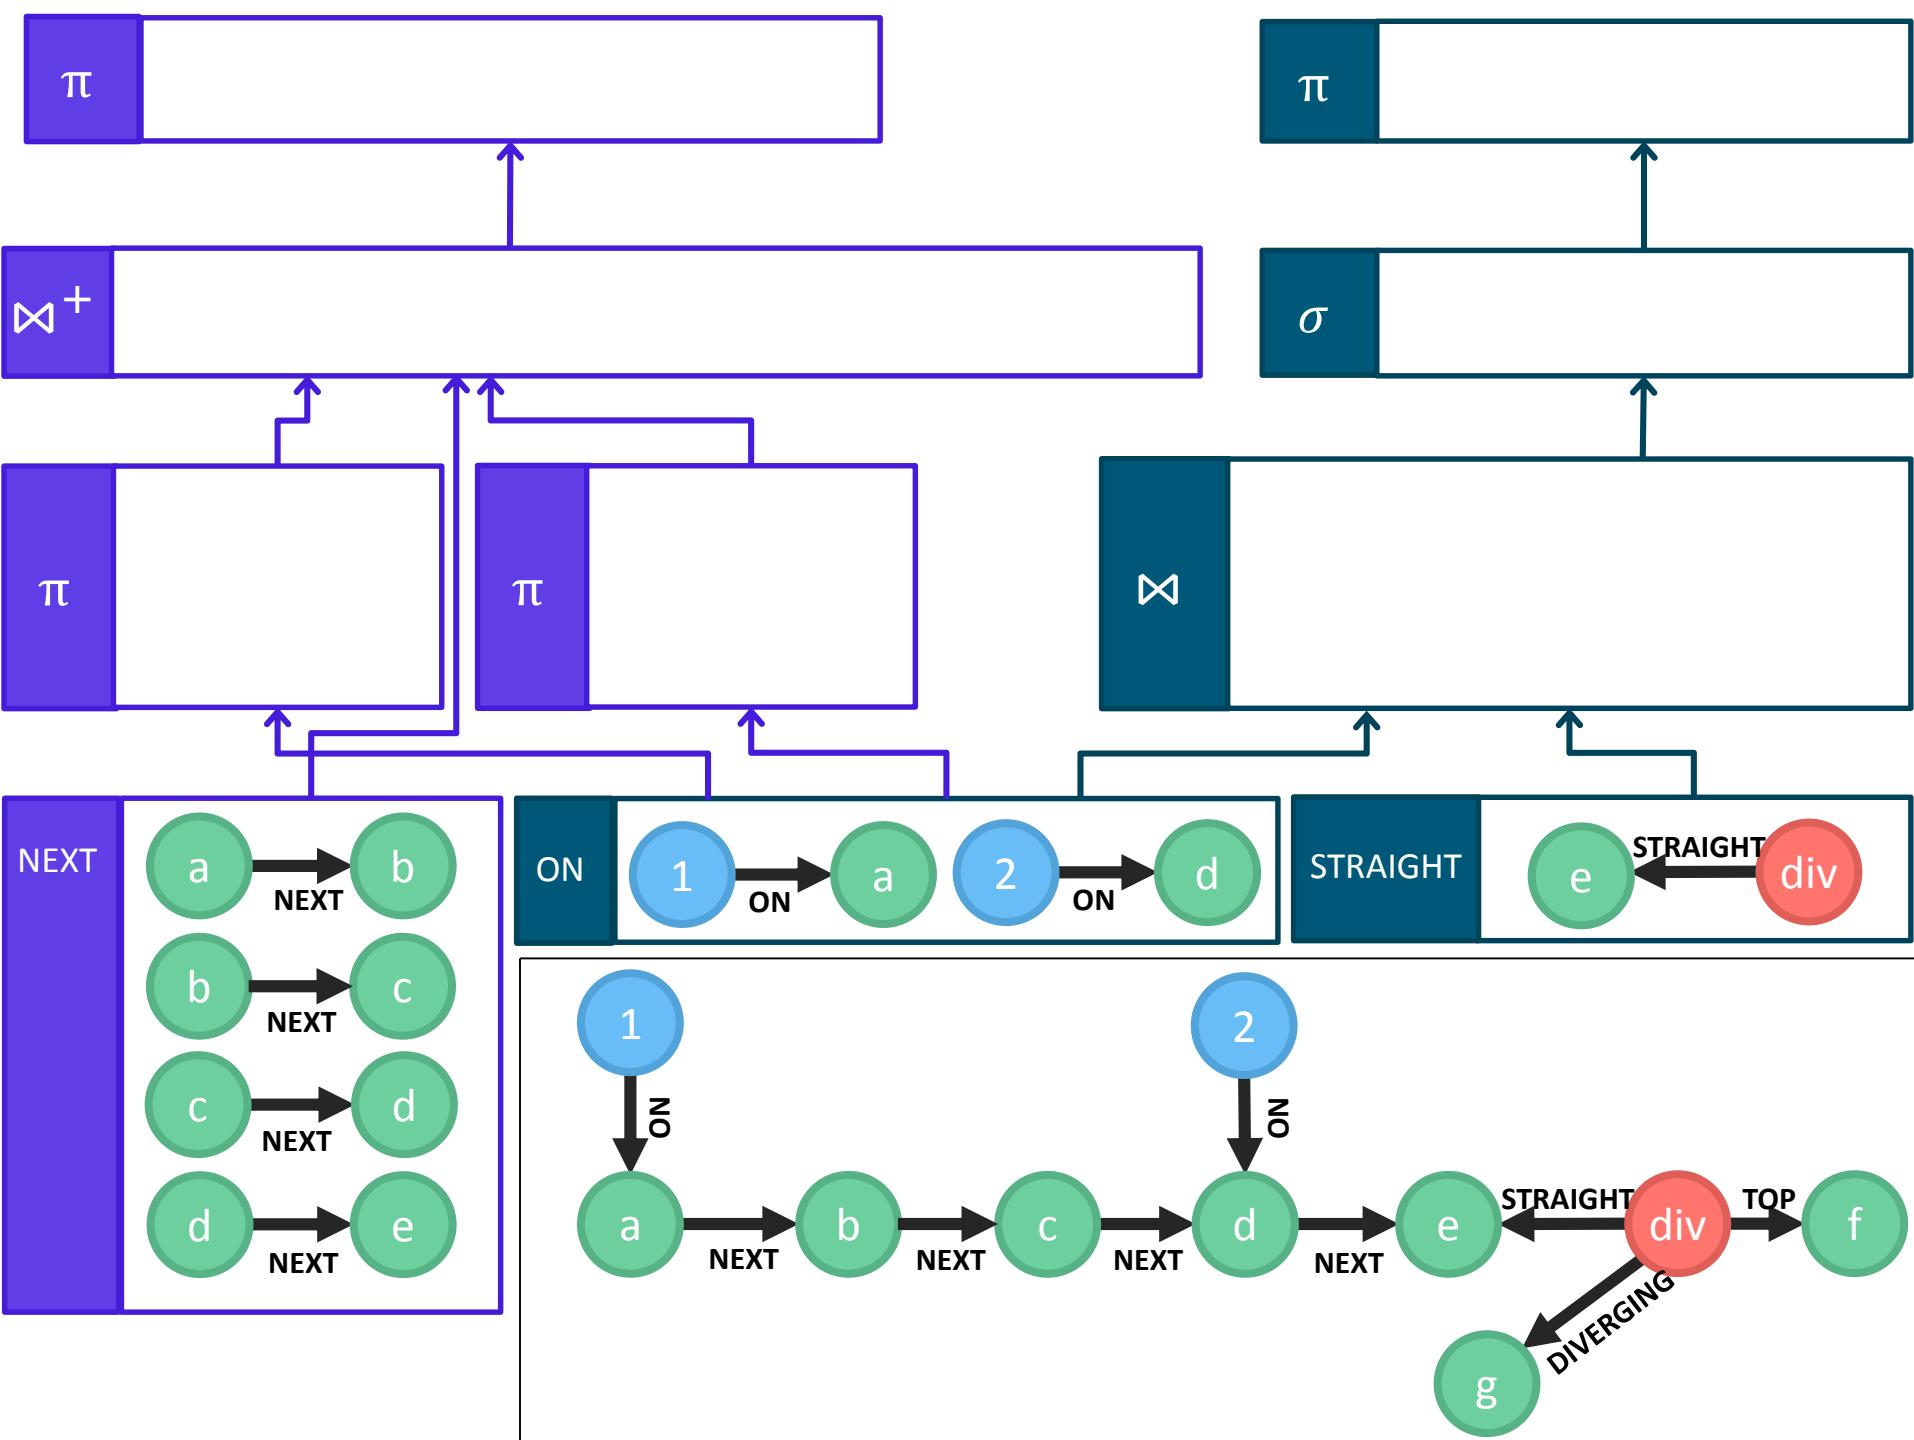


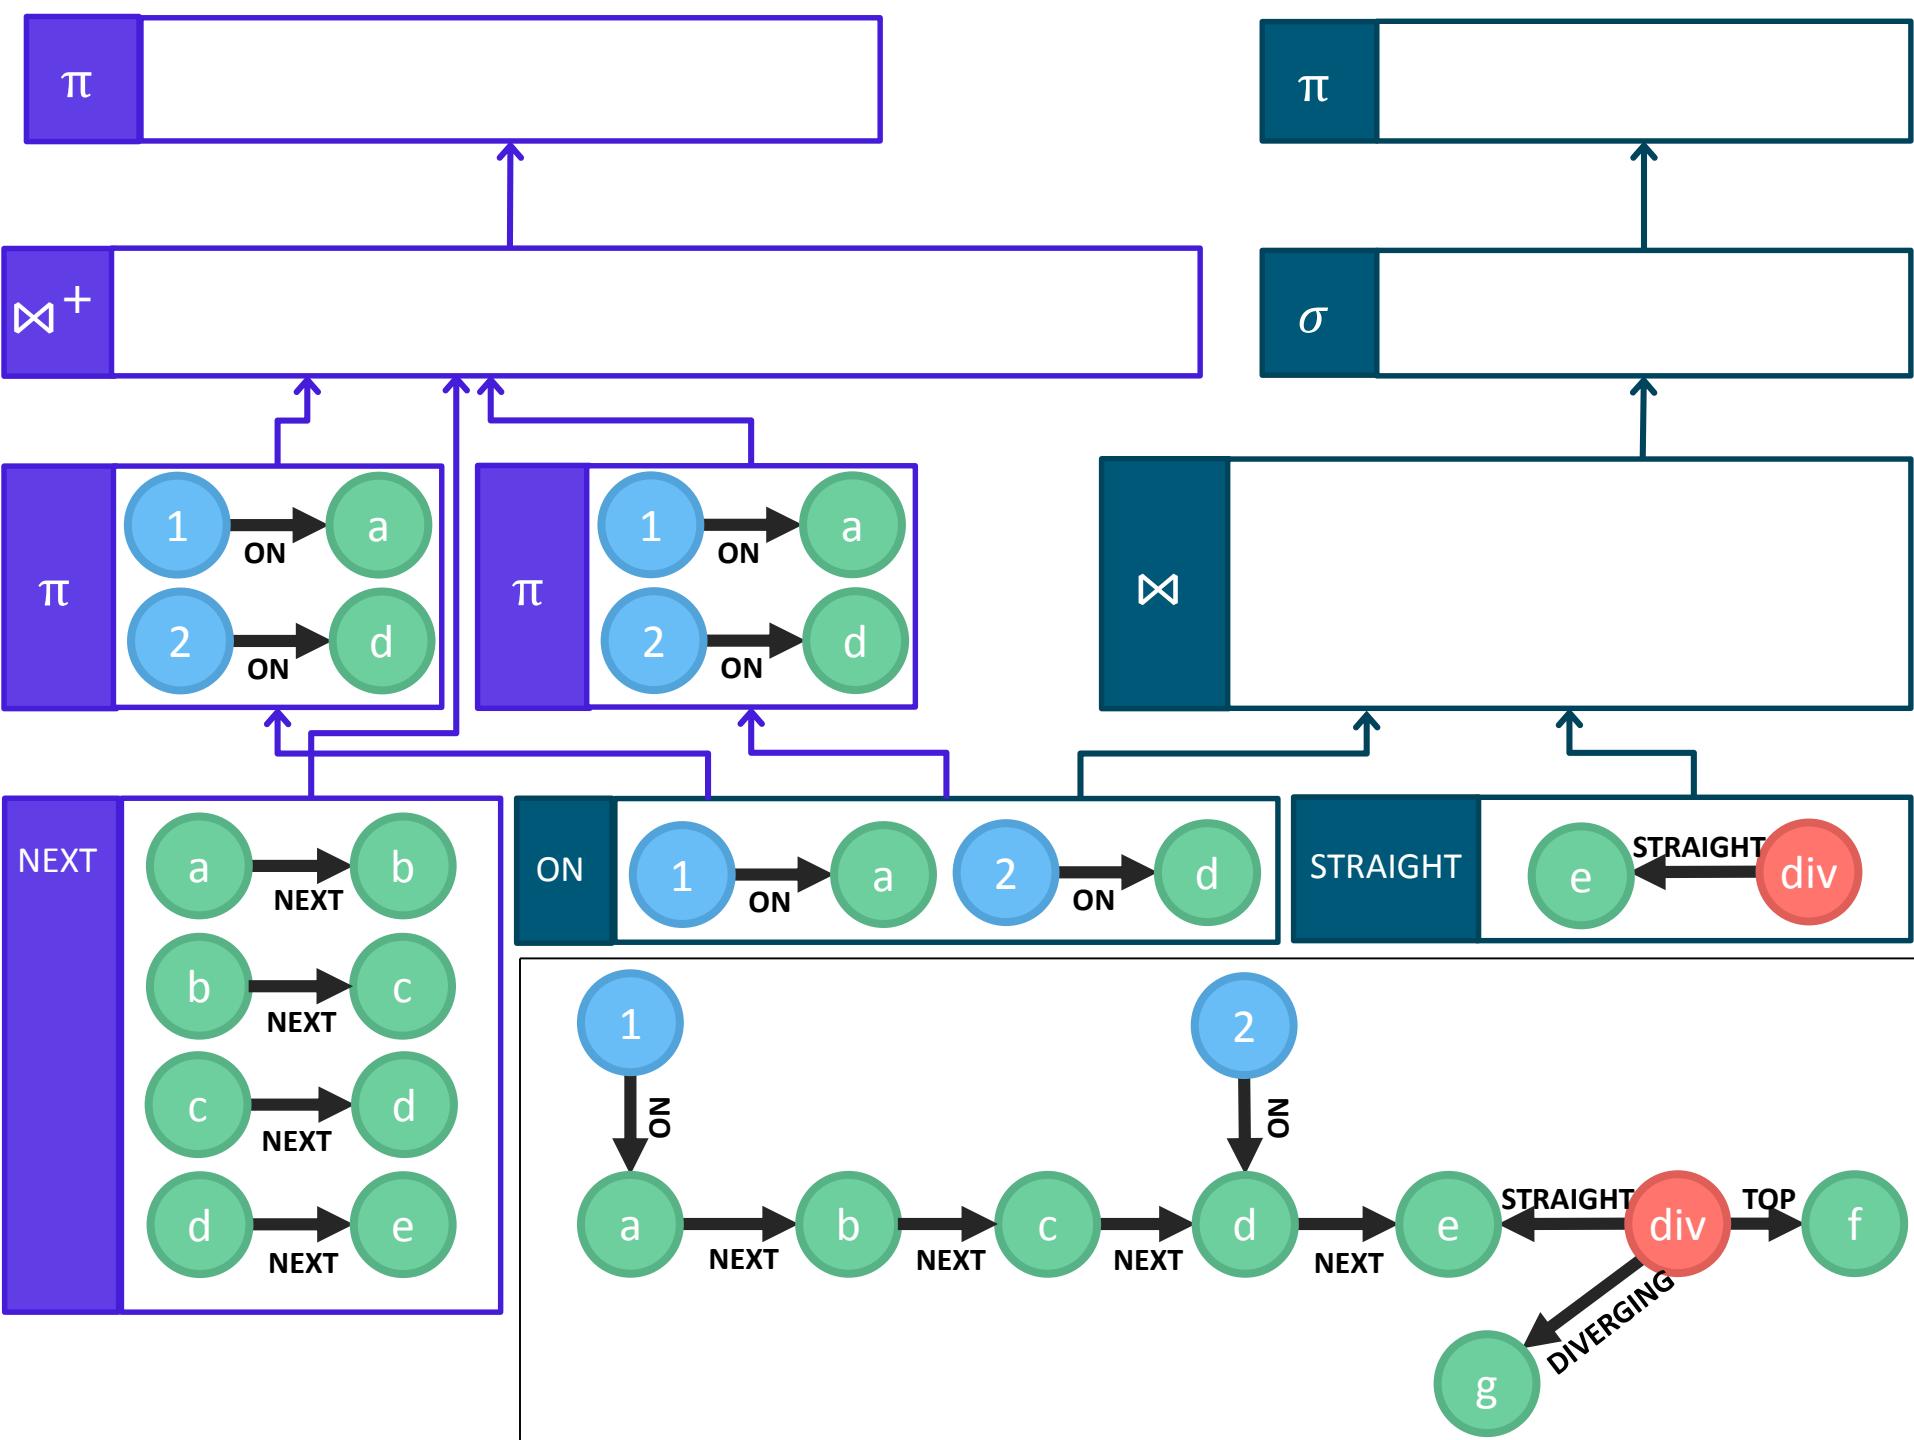


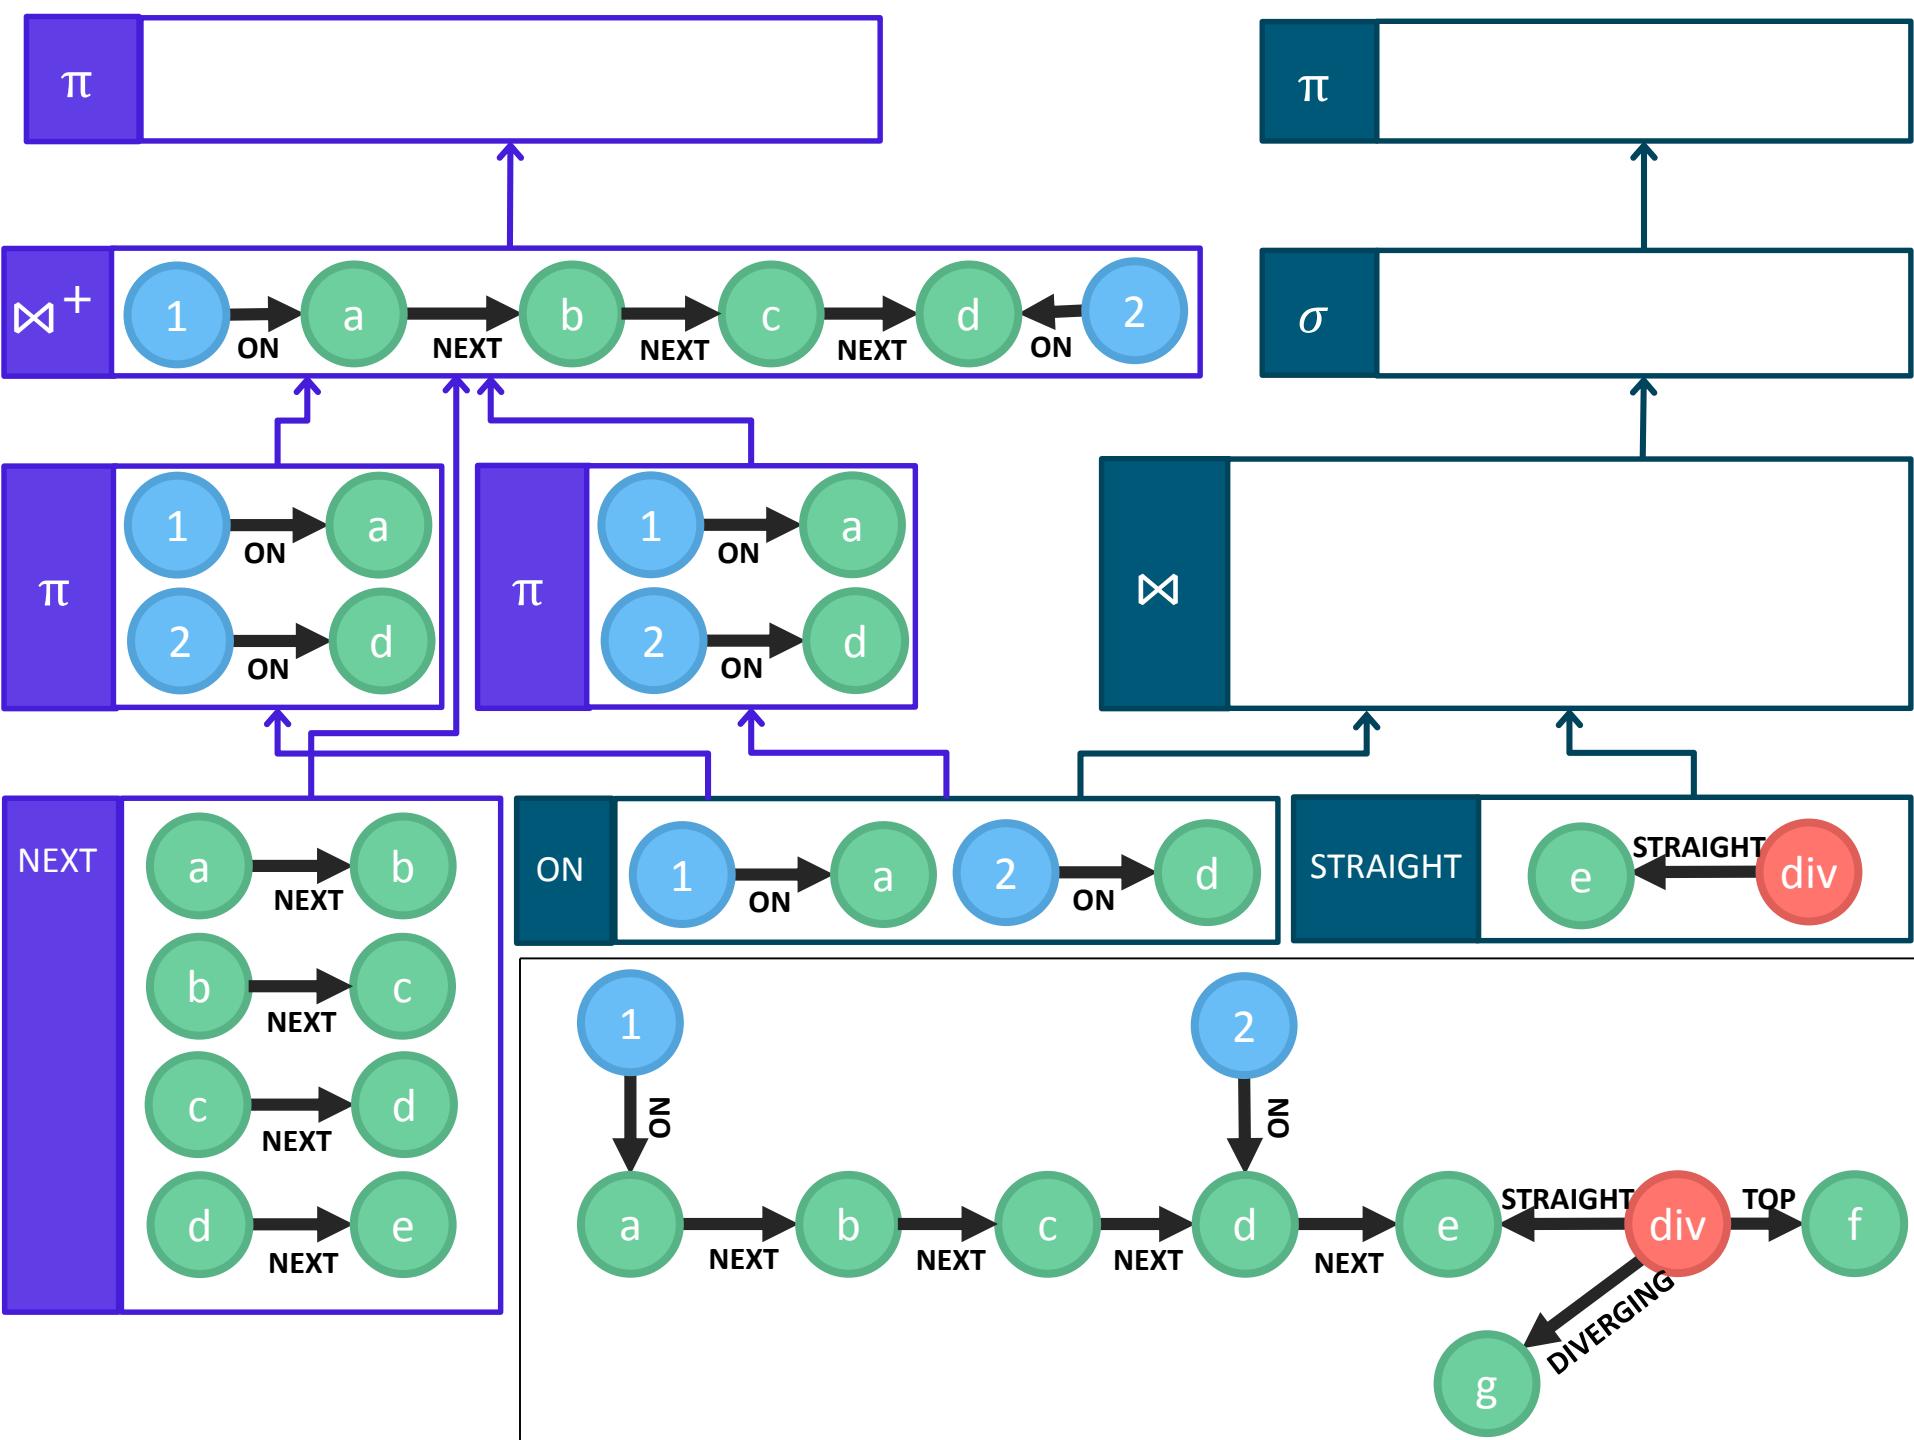


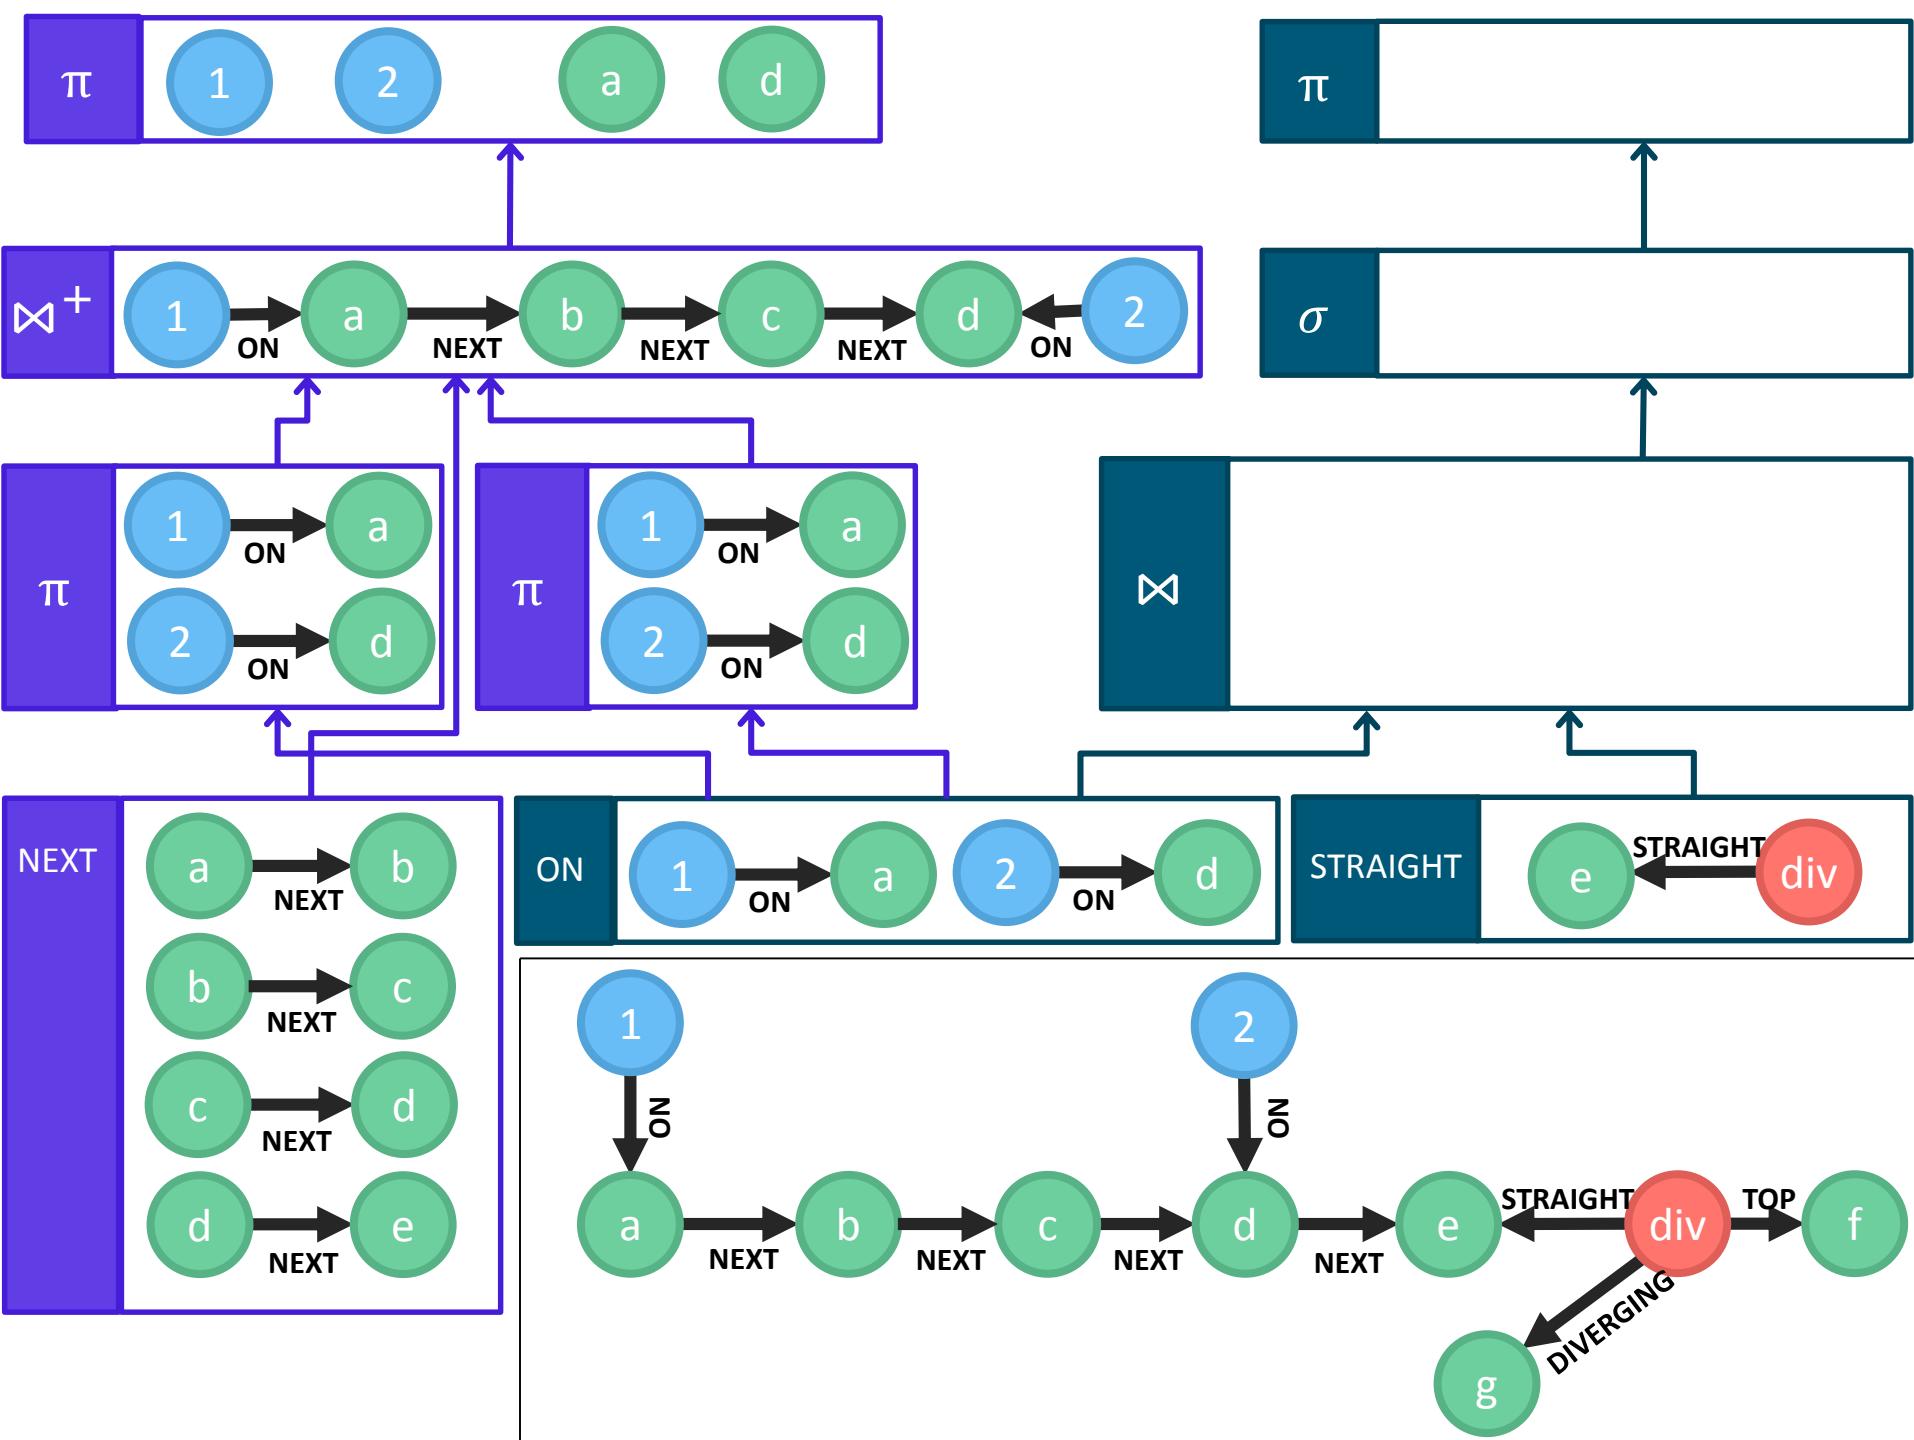


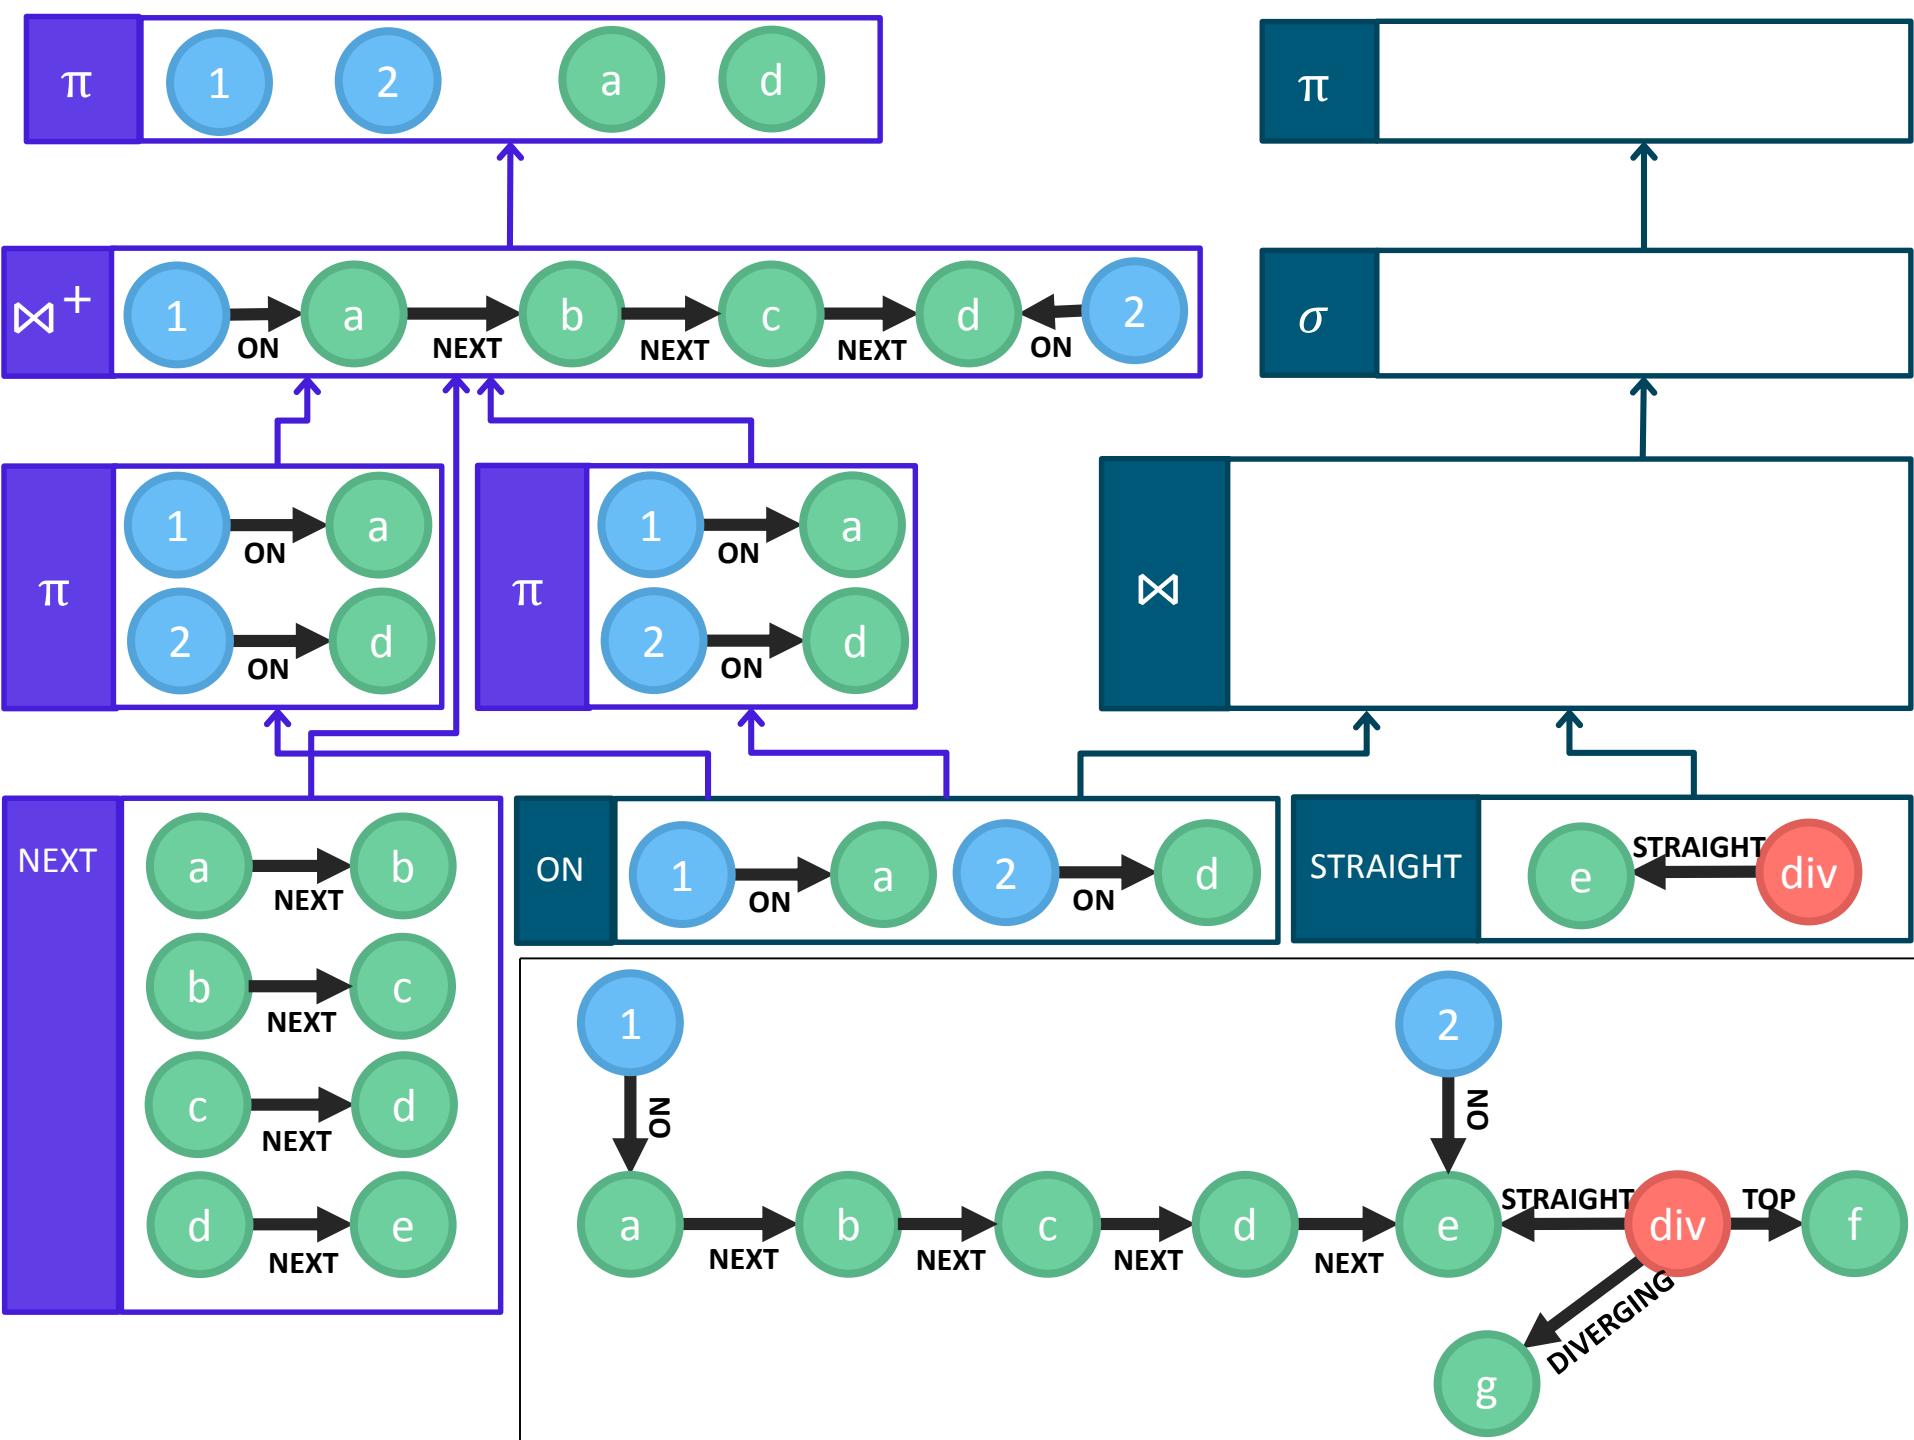


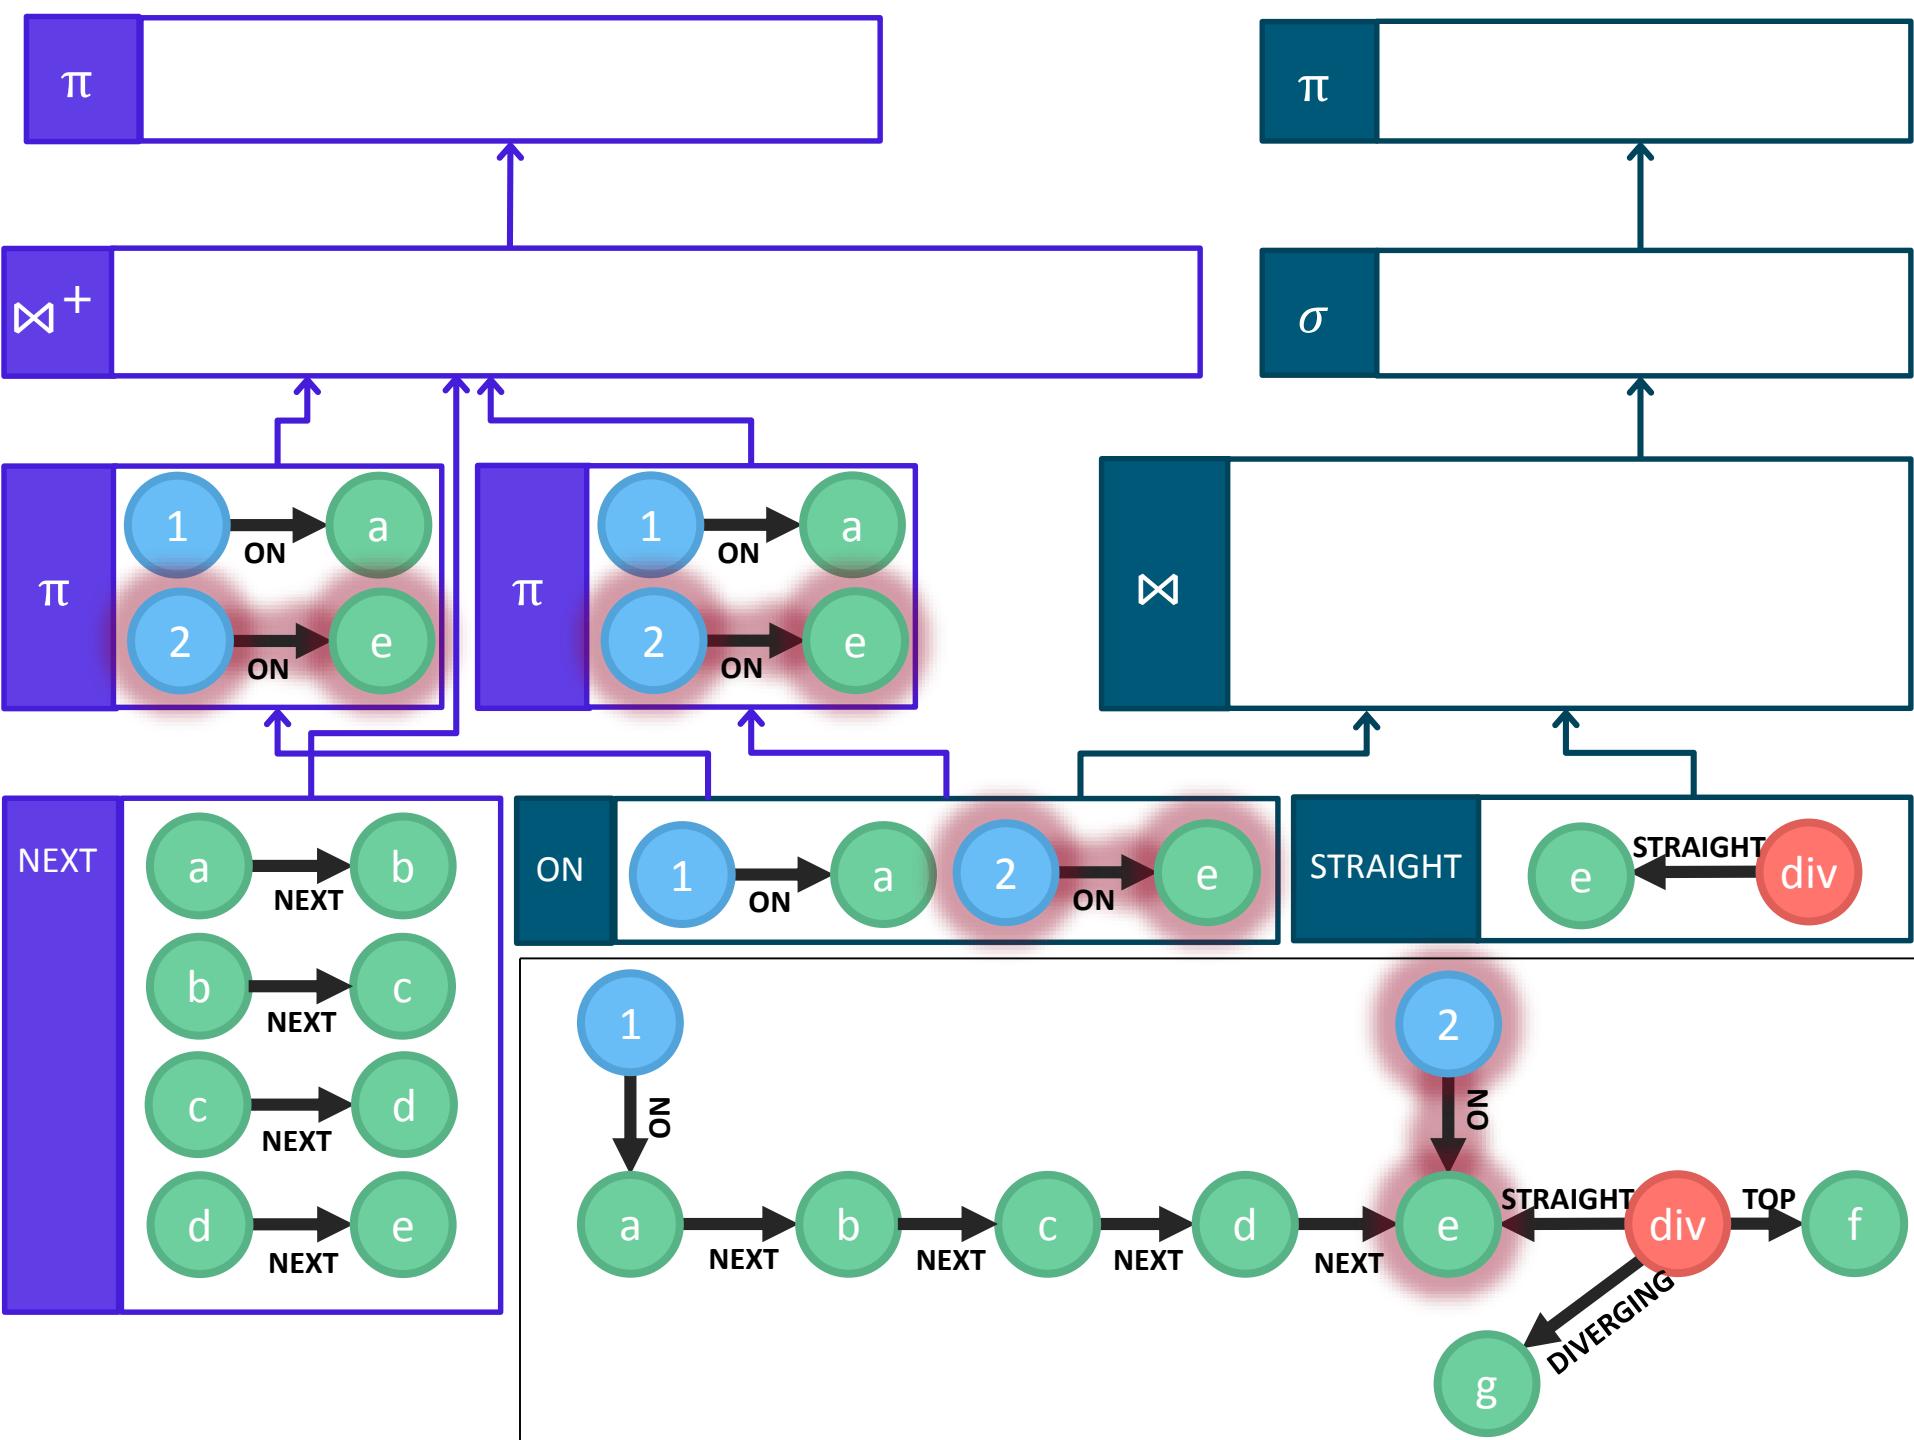


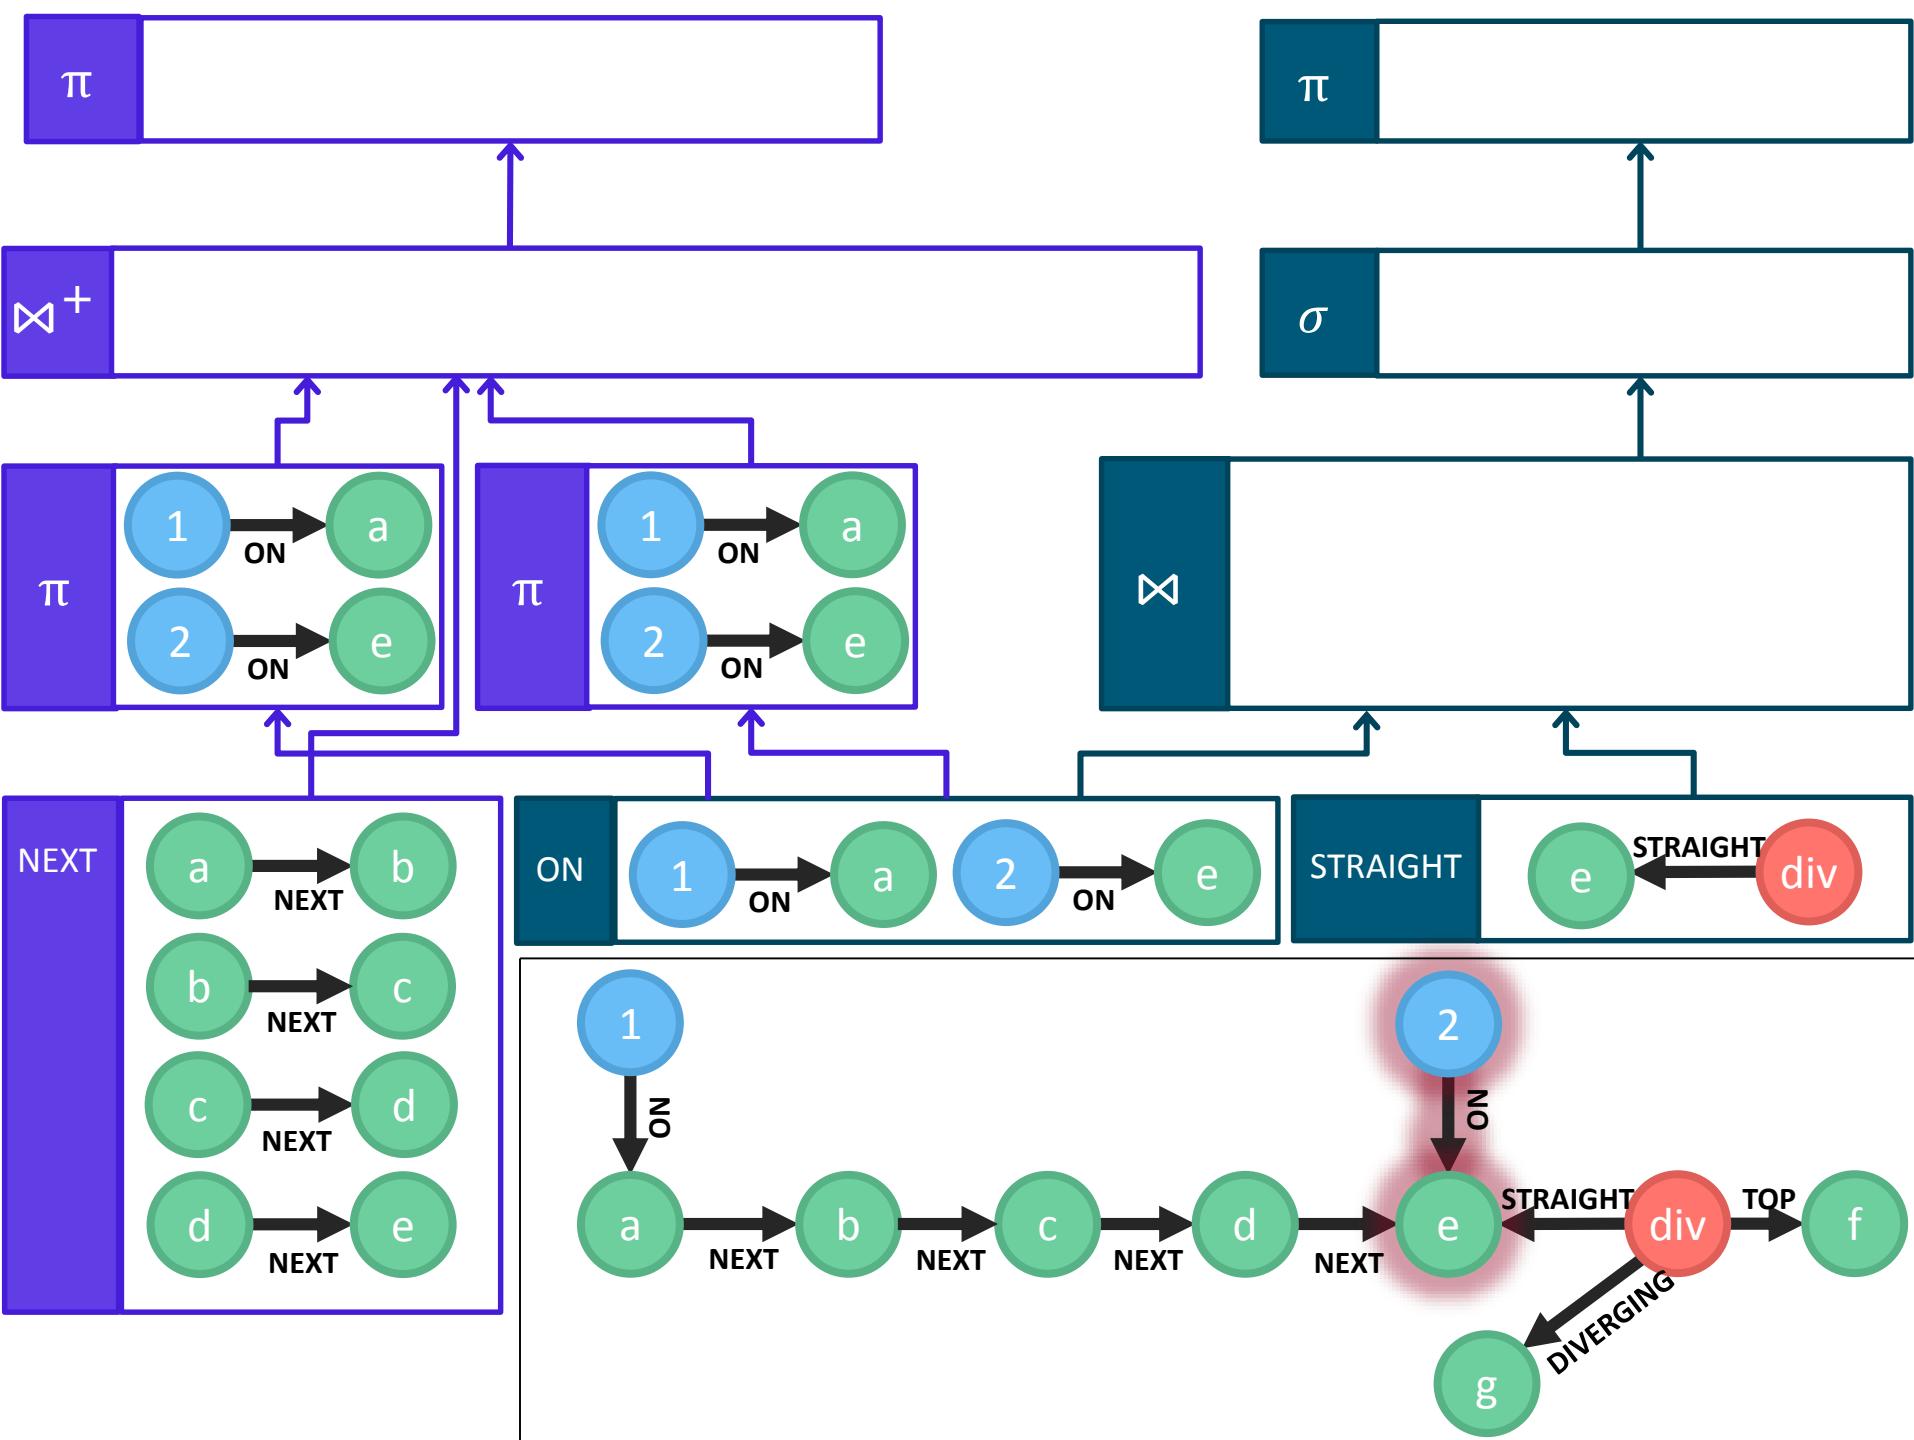


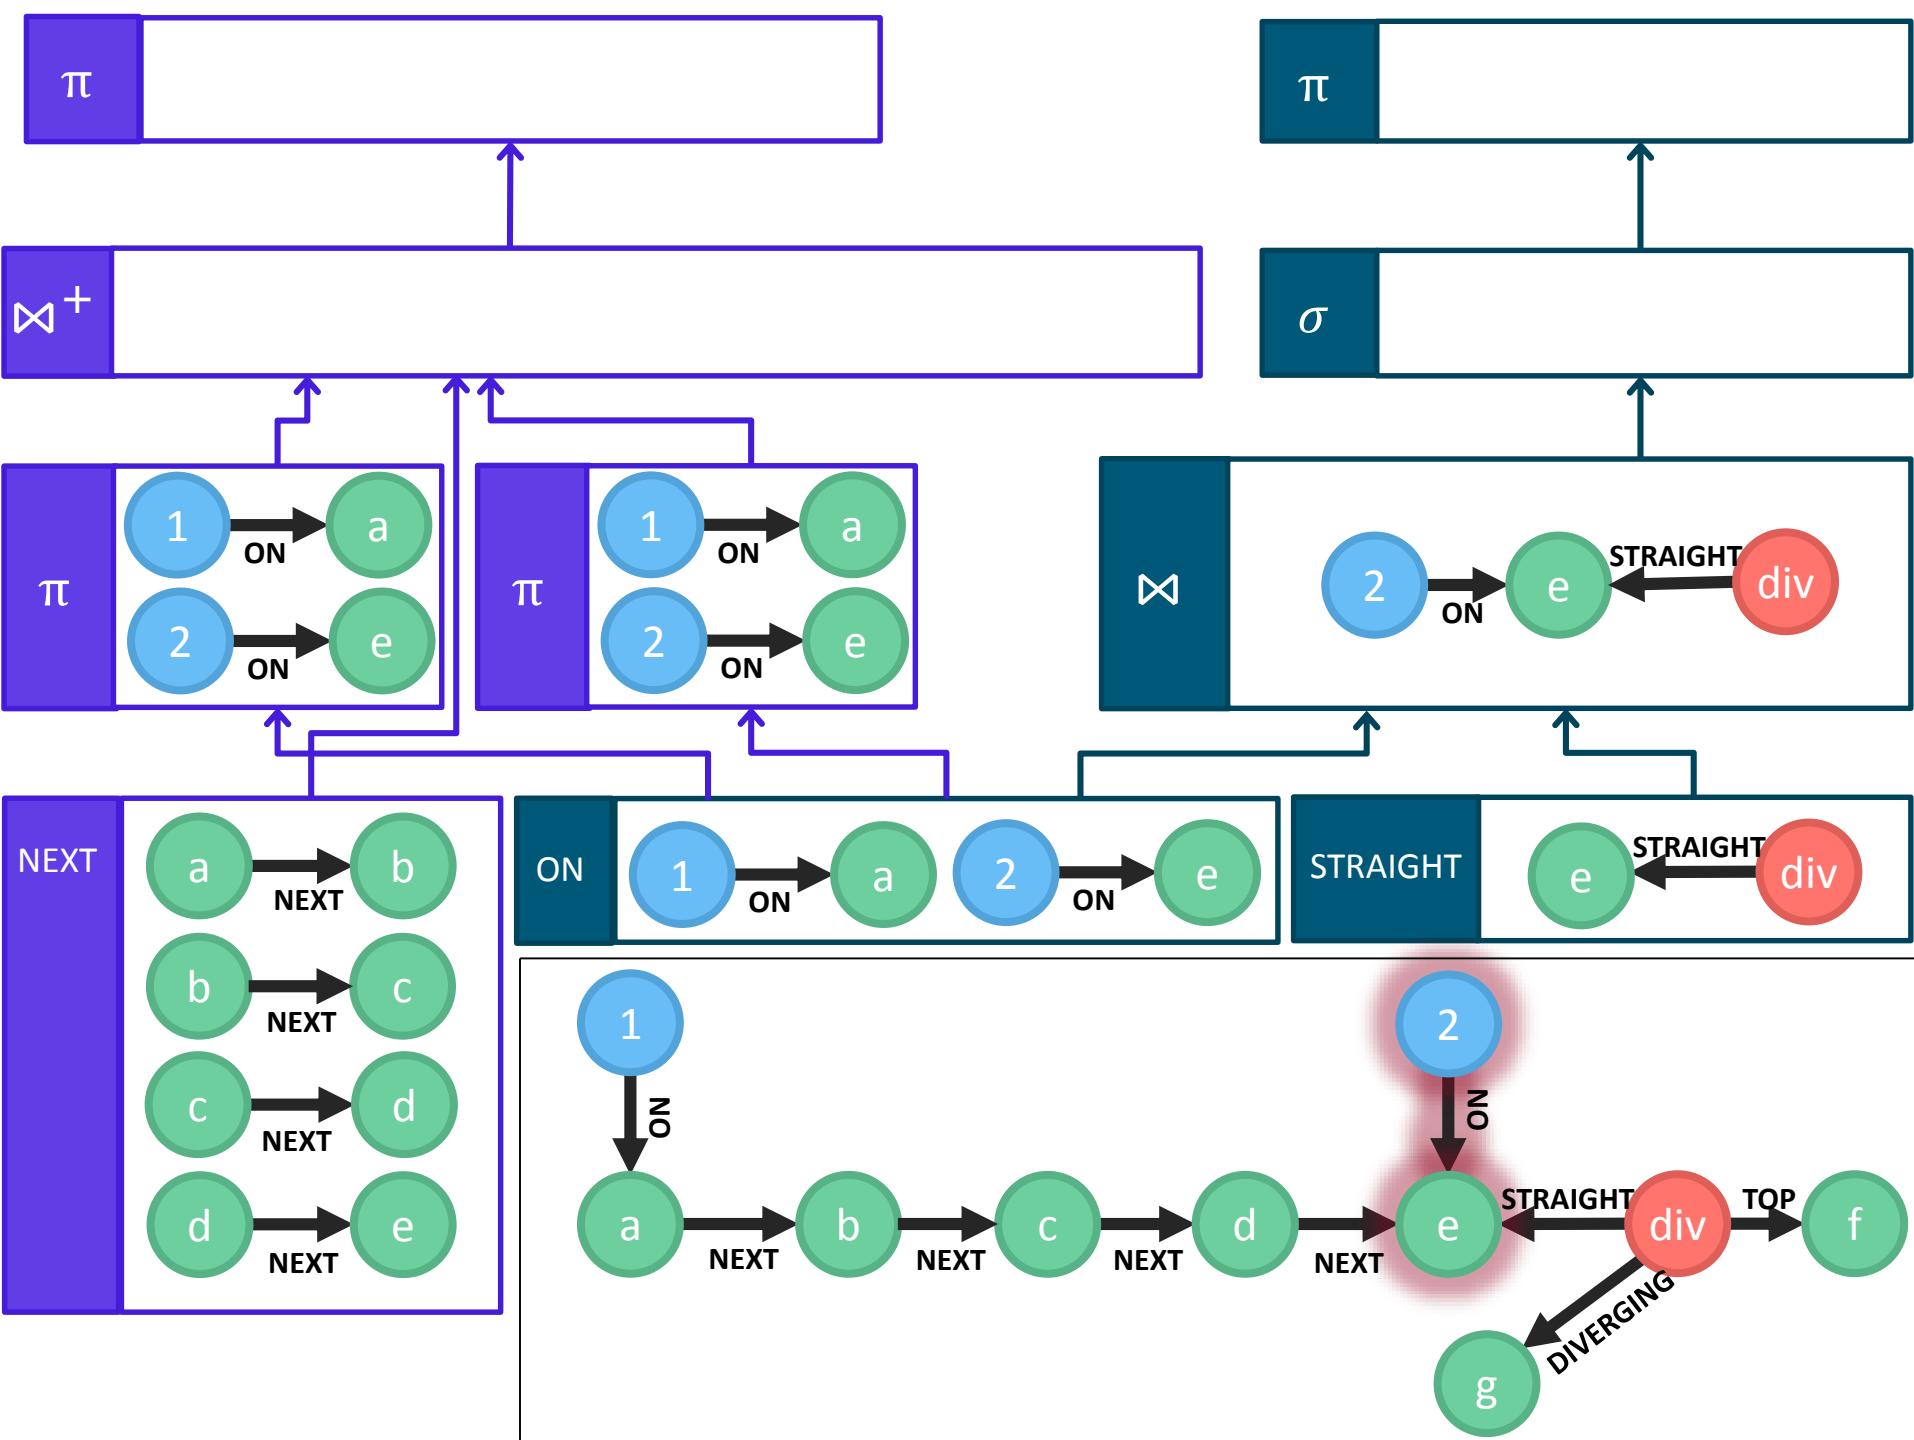


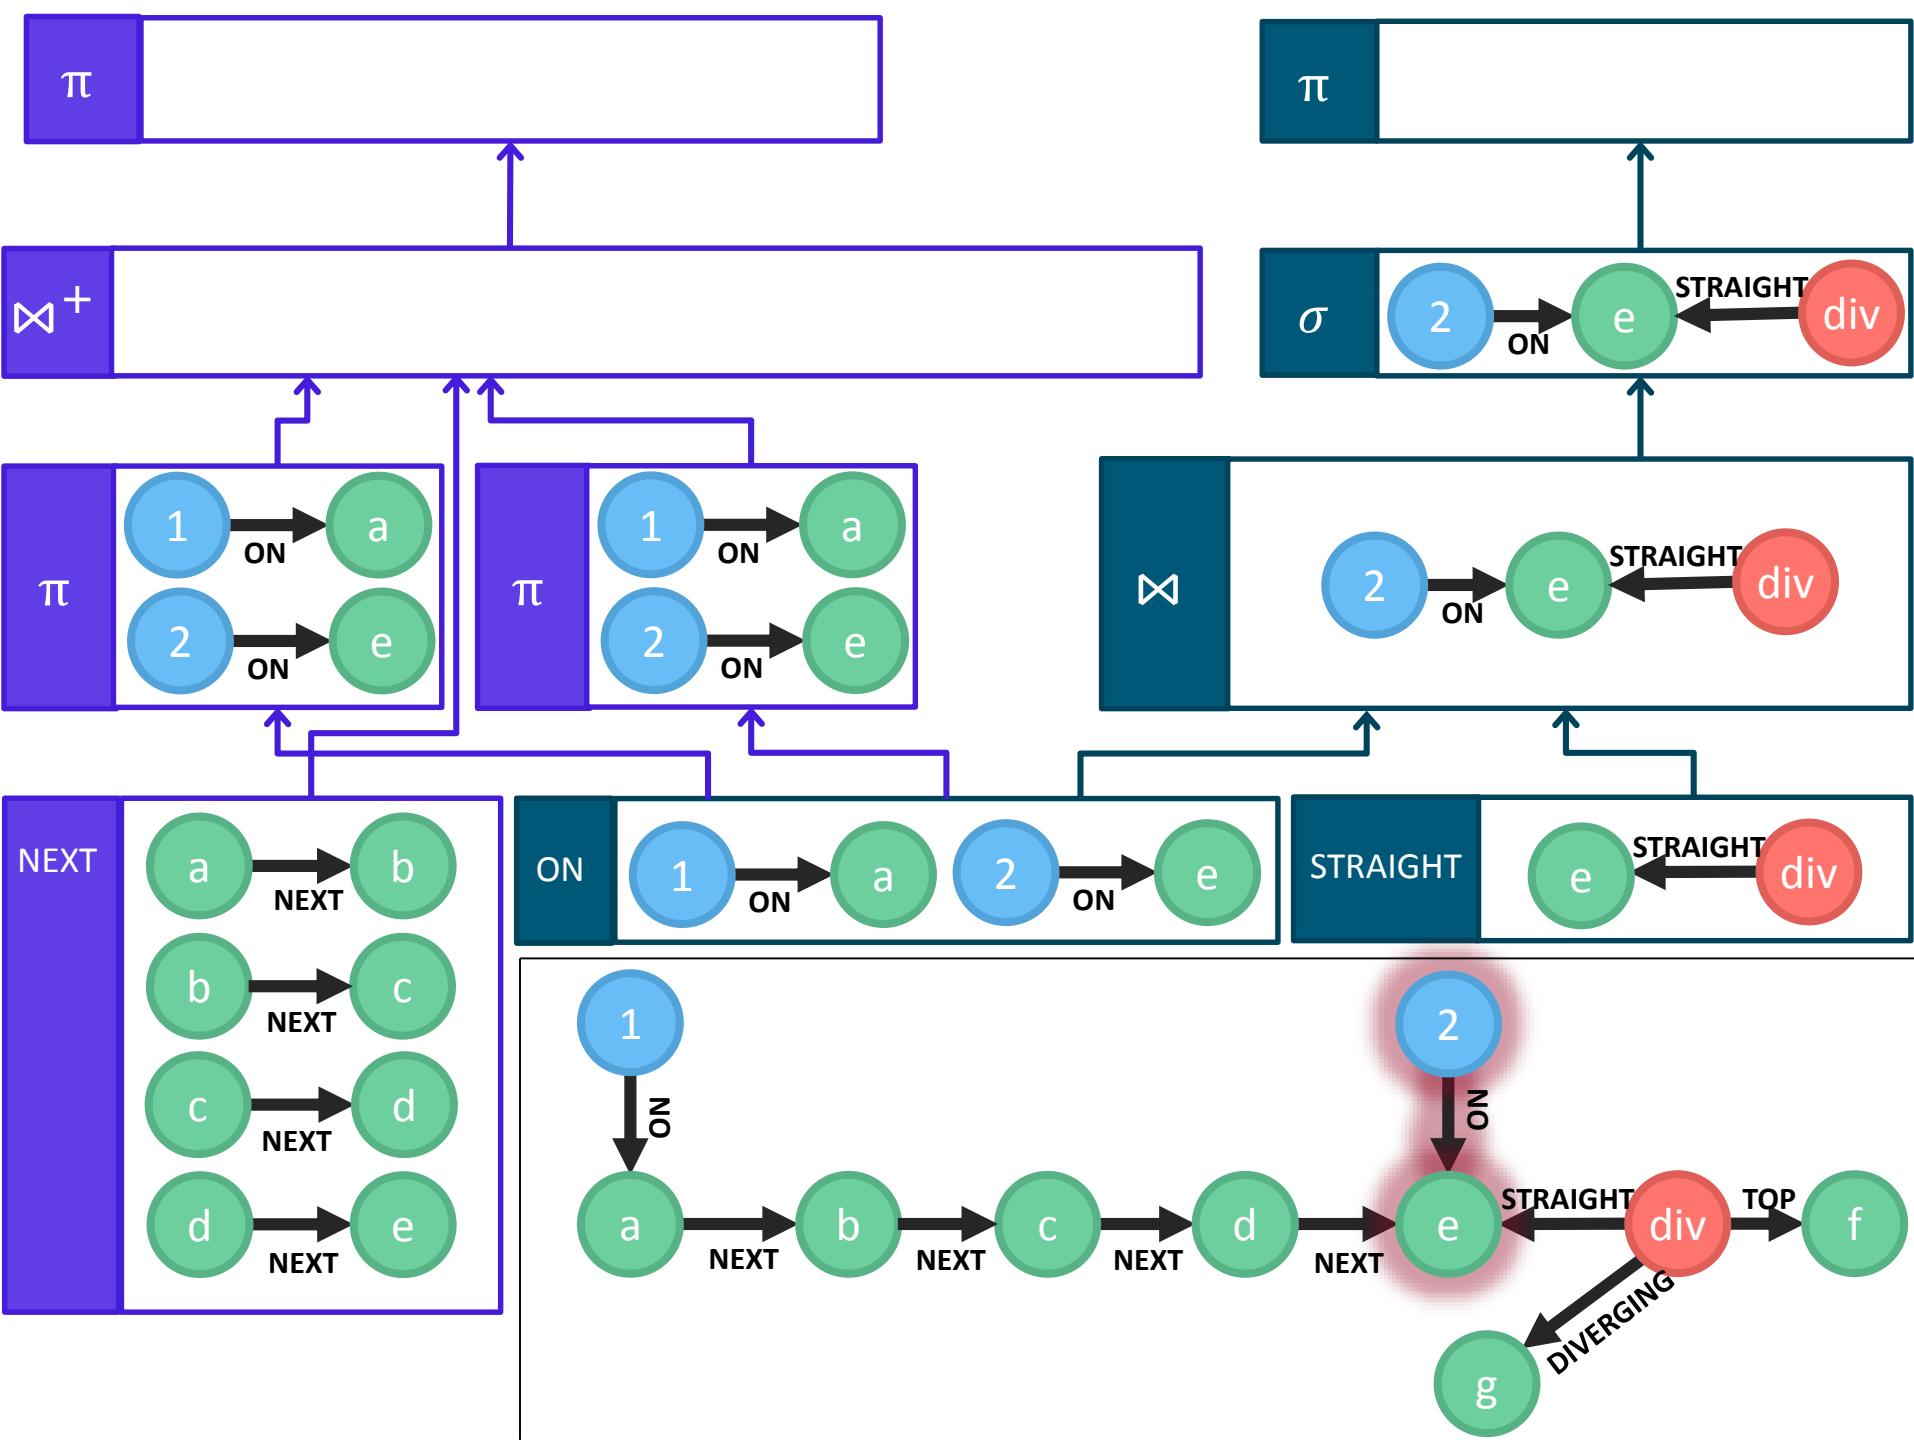


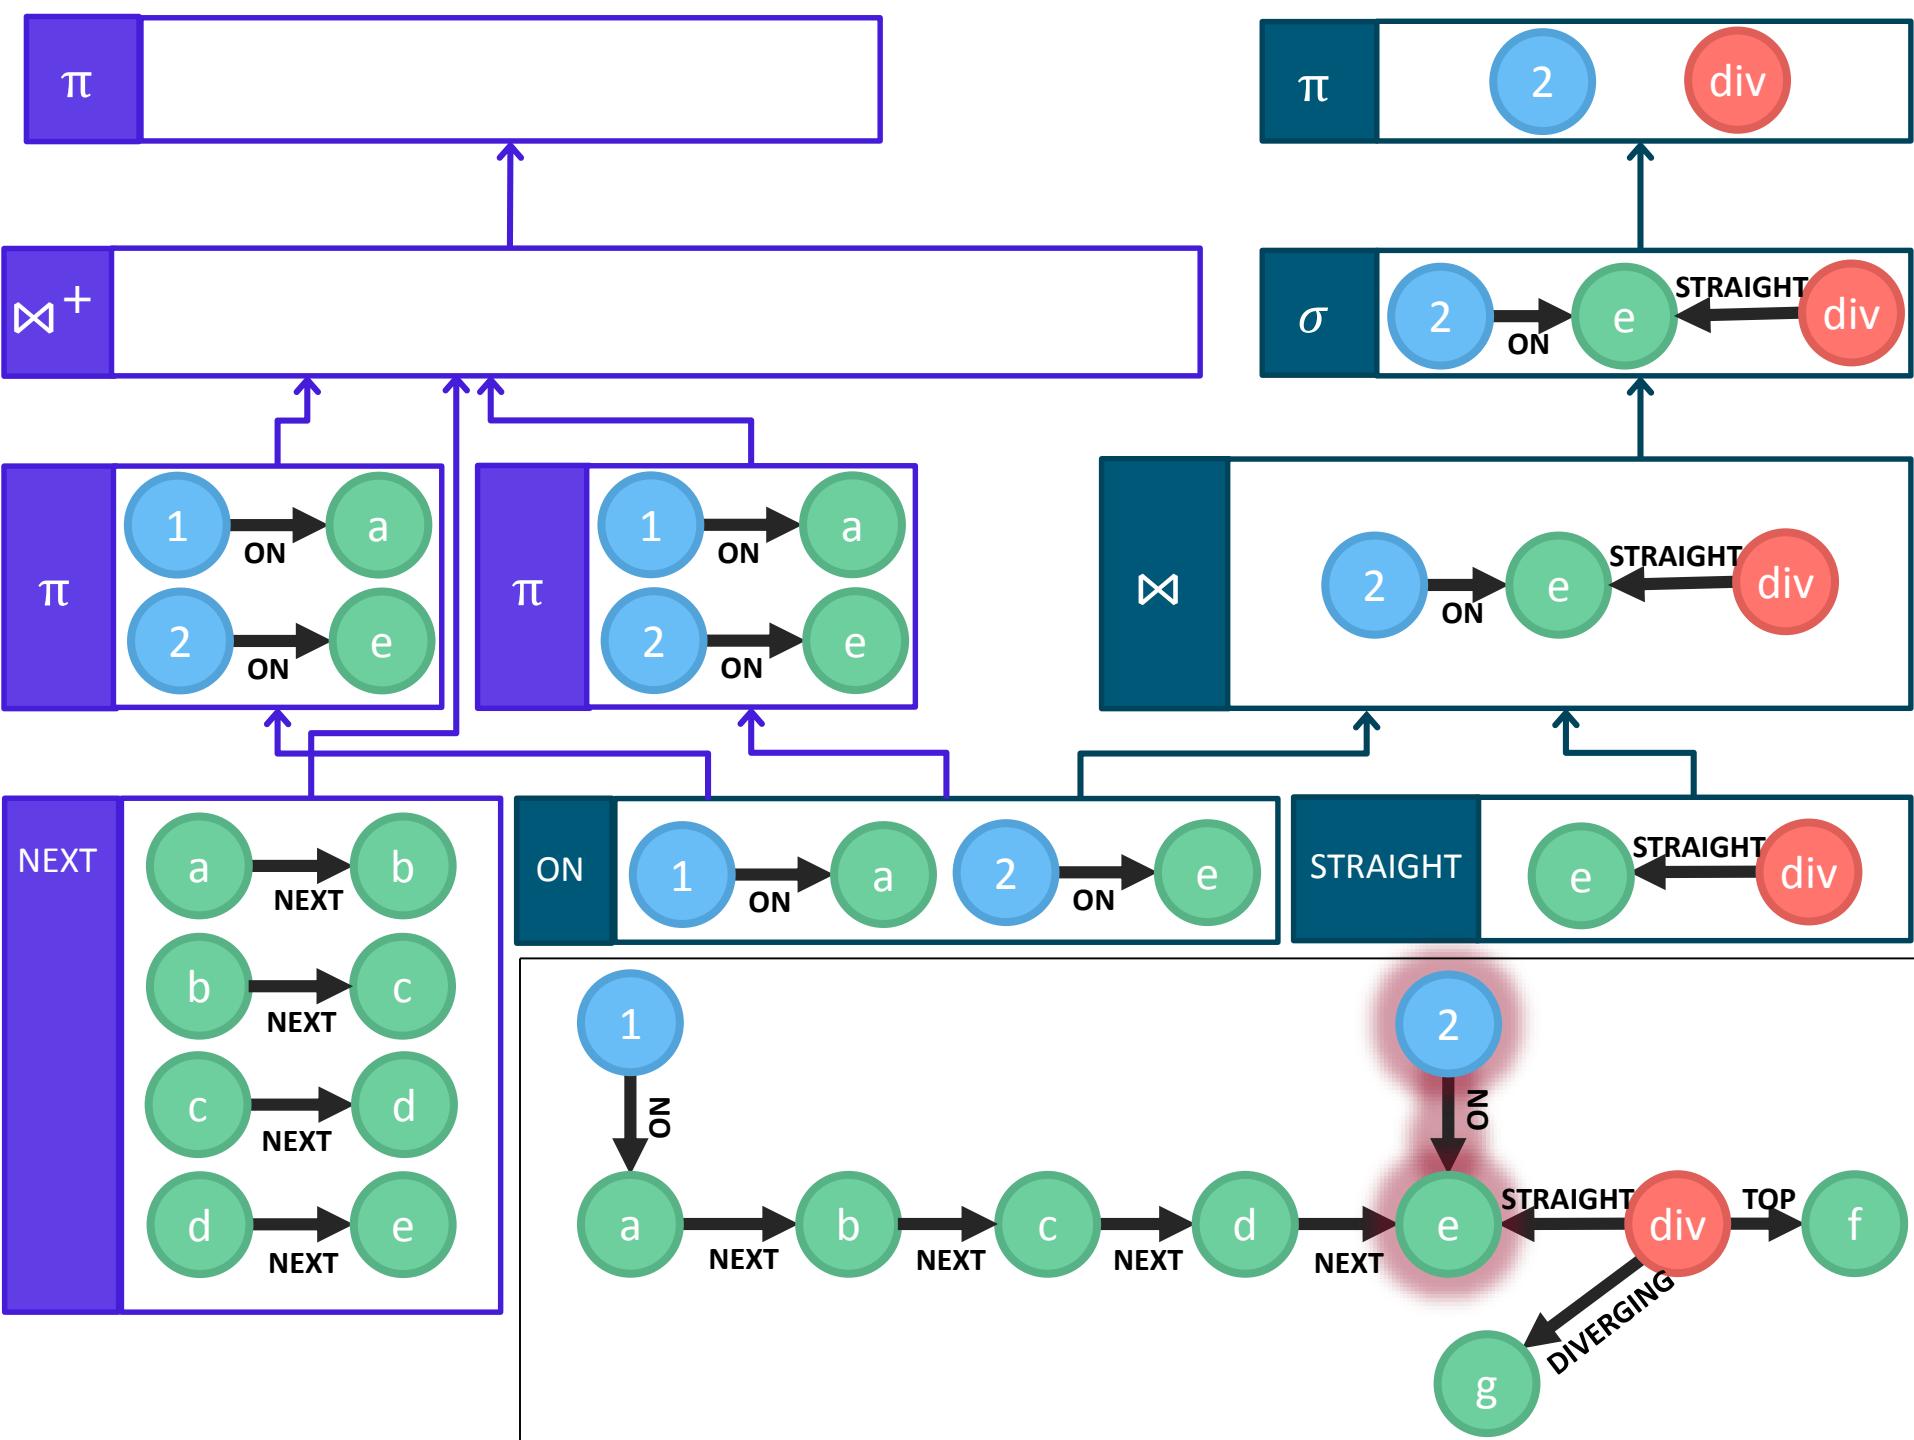












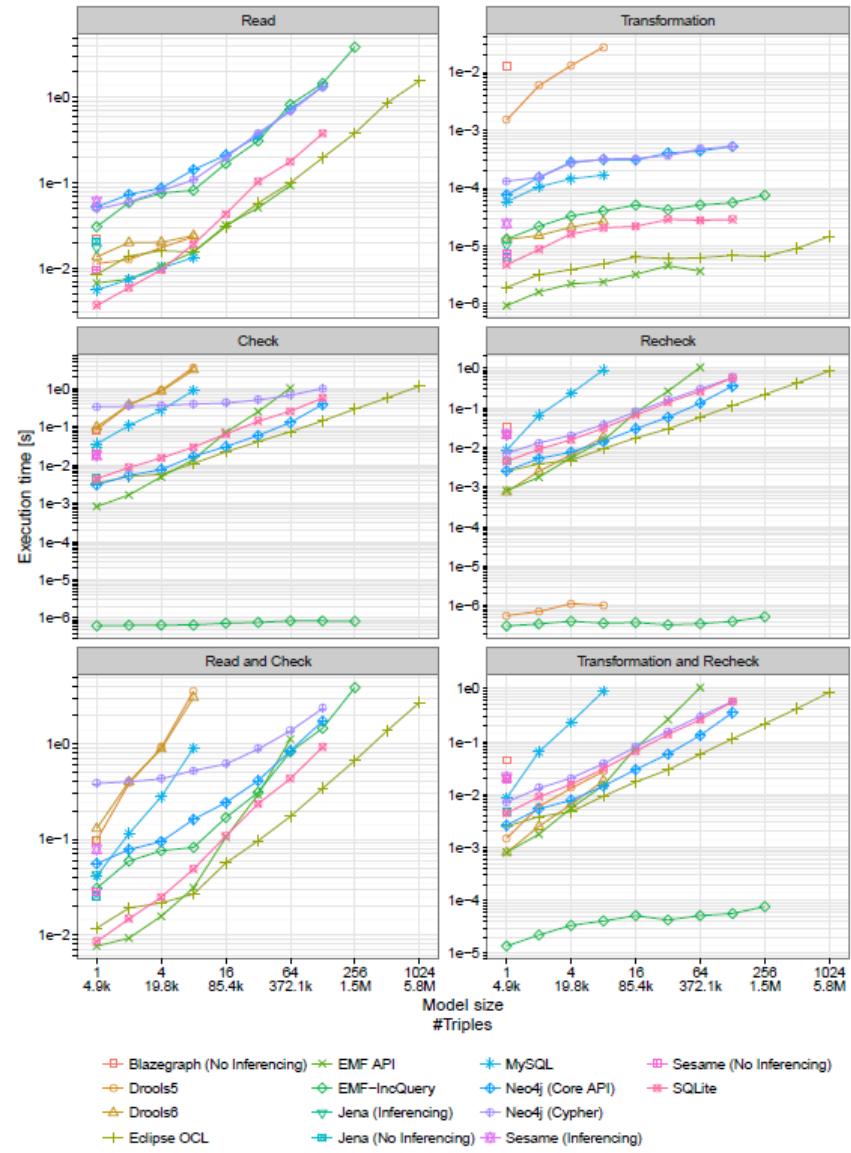
# Graph use cases

Standing queries on large & quickly changing graphs

- Model validation: The Train Benchmark
- Static analysis of JavaScript source code
- Fraud detection
- IT infrastructure monitoring

# Model validation: the Train Benchmark

- Scalable graph generator
  - EMF
  - Property graph
  - RDF
  - SQL
- Realistic workload
  - 6 validation queries
  - 12 transformations
  - Implemented for 12+ tools
- Visualization and reporting

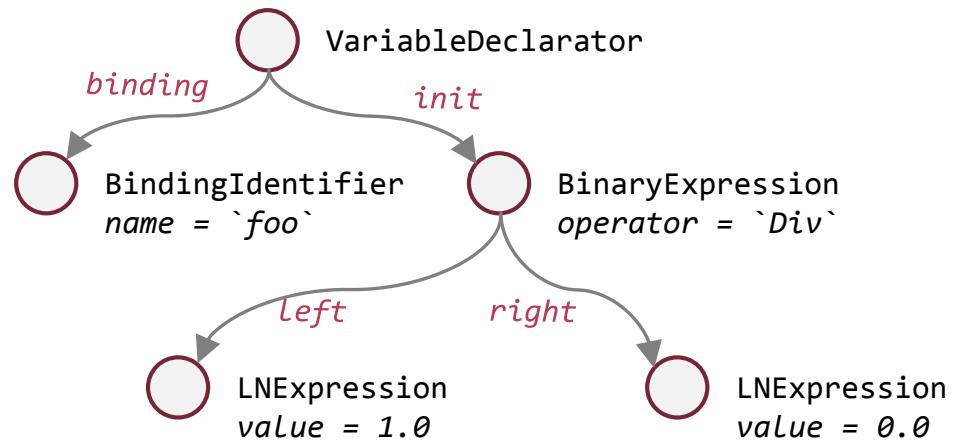


# Static analysis of JavaScript

```
var foo = 1 / 0;
```

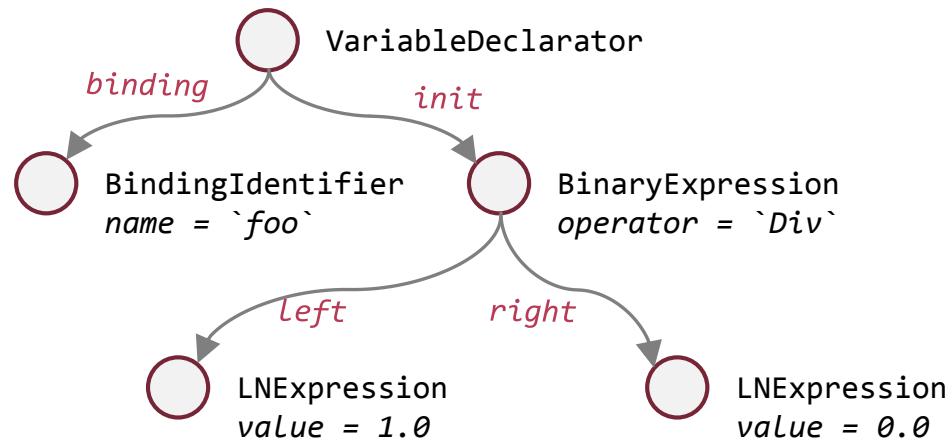
# Static analysis of JavaScript

```
var foo = 1 / 0;
```



# Static analysis of JavaScript

```
var foo = 1 / 0;
```



MATCH

```
(binding:BindingIdentifier)<-[:binding]-()-->
(be:BinaryExpression)-[:right]->
(right:LNEExpression)
```

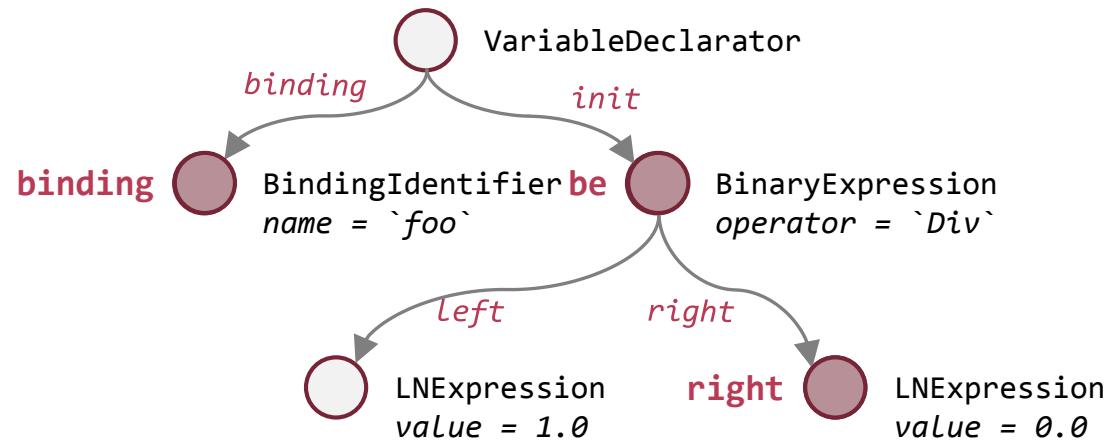
WHERE be.operator = 'Div'

AND right.value = 0.0

RETURN binding

# Static analysis of JavaScript

```
var foo = 1 / 0;
```



MATCH

```
(binding:BindingIdentifier)<-[:binding]-()-->
(be:BinaryExpression)-[:right]->
(right:LNExpression)
```

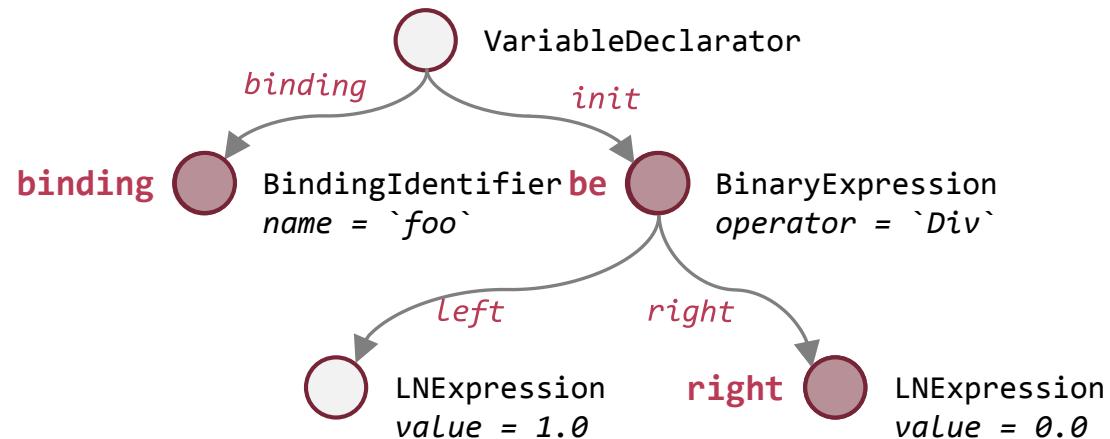
WHERE be.operator = 'Div'

AND right.value = 0.0

RETURN binding

# Static analysis of JavaScript

```
var foo = 1 / 0;
```



MATCH

```
(binding:BindingIdentifier)<- [:binding]-()-->  
(be:BinaryExpression)-[:right]->  
(right:LNExpression)
```

WHERE `be.operator = 'Div'`

AND `right.value = 0.0`

RETURN `binding`

- Dead code detection
- Type inferencing
- Test generation

# Project roadmap

# Project roadmap

- Until now
  - formalization for most standard openCypher constructs  
**[technical report]**
  - research prototype for model validation queries

# Project roadmap

- Until now
  - formalization for most standard openCypher constructs  
**[technical report]**
  - research prototype for model validation queries
- 2017 Q1
  - Use TCK tests for feature coverage & debugging

# Project roadmap

- Until now
  - formalization for most standard openCypher constructs  
**[technical report]**
  - research prototype for model validation queries
- 2017 Q1
  - Use TCK tests for feature coverage & debugging
- 2017 Q2
  - prototype for core openCypher constructs
  - implement optimizer
  - publish benchmark results **[conference paper]**

# Project roadmap

- Until now
  - formalization for most standard openCypher constructs  
**[technical report]**
  - research prototype for model validation queries
- 2017 Q1
  - Use TCK tests for feature coverage & debugging
- 2017 Q2
  - prototype for core operators **Spark Catalyst or Apache Calcite**
  - implement optimizer
  - publish benchmark results **[conference paper]**

# Open-Source Projects

**Incremental Graph Engine:**

<https://github.com/ftsrg/ingraph>

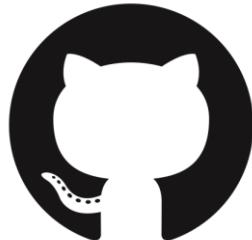
**Train Benchmark:**

<https://github.com/ftsrg/trainbenchmark>

**BME-MODES3:**

<https://github.com/ftsrg/bme-modes3>

Available under EPL v1.0.



$\Omega$ 

MŰEGYETEM 1782

