

Inferring Likely Deterministic Specifications for Multithreaded Programs

Jacob Burnim

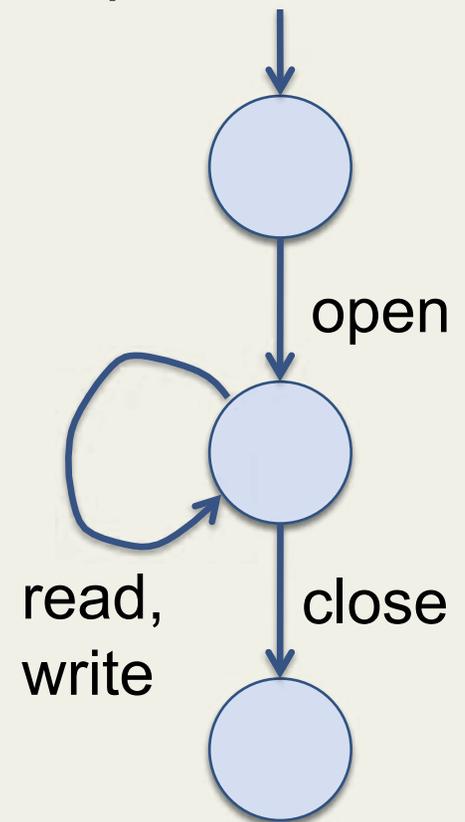
Koushik Sen

Parallel Computing Laboratory
University of California, Berkeley

Specification Inference

- Inferring **temporal** specs (protocols):

- *Ammons, et al, POPL 02*
- *Whaley, et al, ISSTA 02*
- *Livshits, Zimmermann, FSE 05*
- *Perrecotta: Yang, et al, ICSE 06*
- *JADET: Wasylkowski, et al, FSE 07*
- *Acharya, et al, FSE 07*
- *JAVERT: Gabel, Su, ICSE 08*
- many others ...



Specification Inference

- Inferring likely *invariants*:

$$0 \leq i < N \quad \text{sum} = \sum_{j=0}^{i-1} A[j] \quad \wedge \quad i = N + 1$$

- DAIKON: *Ernst, et al., ICSE 00*
- DIDUCE: *Hangal, Lam, ICSE 02*
- DySy: *Csallner, et al., ICSE 08*
- *Gulwani, et al., VMCAI 09*
- many others ...

Parallel Programming is Hard

- **Key Culprit: Nondeterminism.**
 - Interleaving of parallel threads
- **Determinism** key to parallel correctness.
 - Same input \implies semantically same output.
 - Parallelism is wrong if some schedules give a correct answer while others don't.
- **Previously:** Help programmers make their parallel code deterministic.
 - Assertion framework to **specify** determinism.

Advantages of Deterministic Specs

Burnim, Sen. *Asserting and Checking Determinism for Multithreaded Programs*. FSE 2009.

- Lightweight spec of parallel correctness
 - Independent of functional specification
 - Decompose correctness efforts
- Useful for documentation
- Can effectively test deterministic specs
 - Combine with testing tools to distinguish harmful from benign data races, etc.

Advantages of Deterministic Specs

Burnim, Sen. *Asserting and Checking Determinism for Multithreaded Programs*. FSE 2009.

Goal: Automatically infer deterministic specifications by observing sample program runs.

- Useful for documentation
- Can effectively test deterministic specs
 - Combine with testing tools to distinguish harmful from benign data races, etc.

Advantages of Deterministic Specs

Burnim, Sen. *Asserting and Checking Determinism for Multithreaded Programs*. FSE 2009.

Goal: Automatically infer deterministic specifications by observing sample program runs.

Useful for documentation

Result: Recover our previous manual specifications for most benchmarks.

Outline

- Motivation and Overview
- **Background: Deterministic Specs**
- Specification Inference Problem
- Inferring Deterministic Specifications
- Experimental Evaluation
- Related Work
- Conclusions

Background: Deterministic Specs

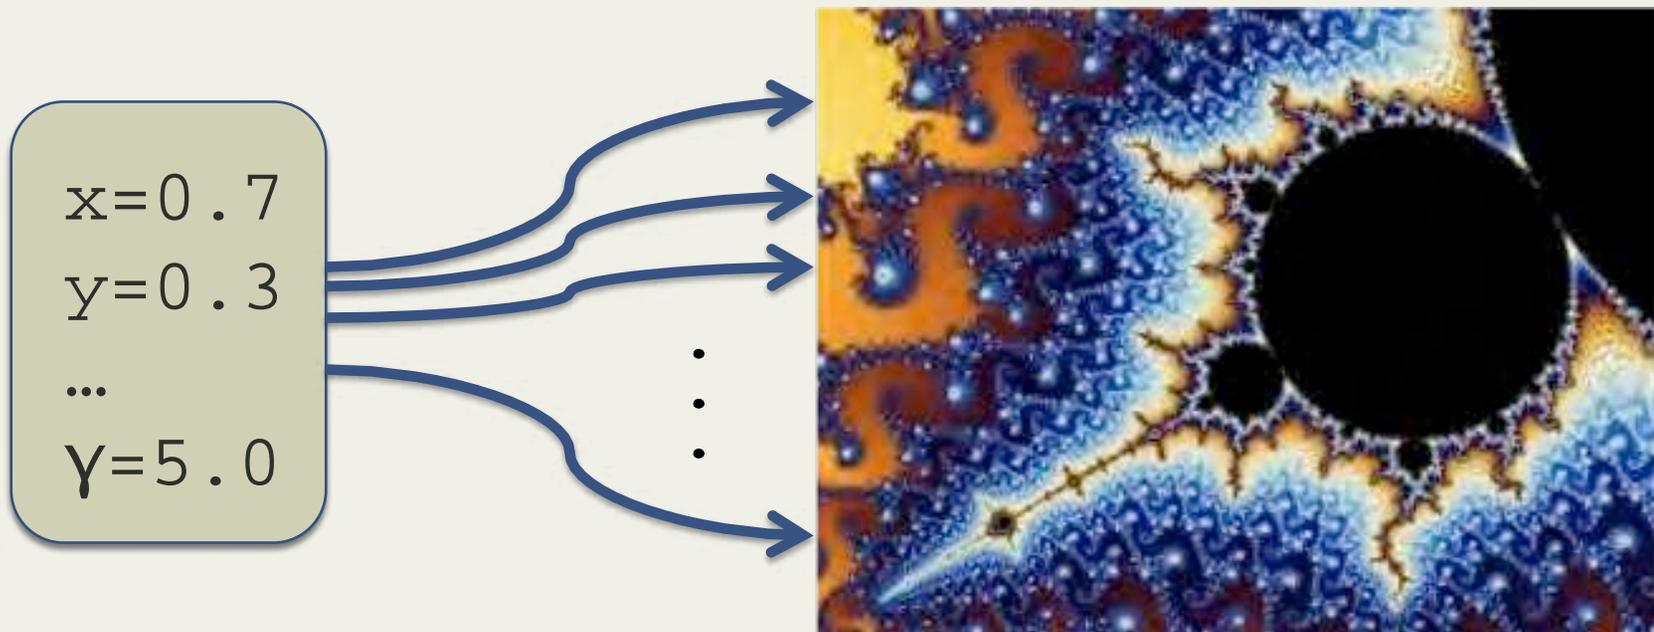
```
// Parallel fractal render.  
mandelbrot(params, img);
```

- Want to assert **parallel correctness**.

Background: Deterministic Specs

```
// Parallel fractal render.  
mandelbrot(params, img);
```

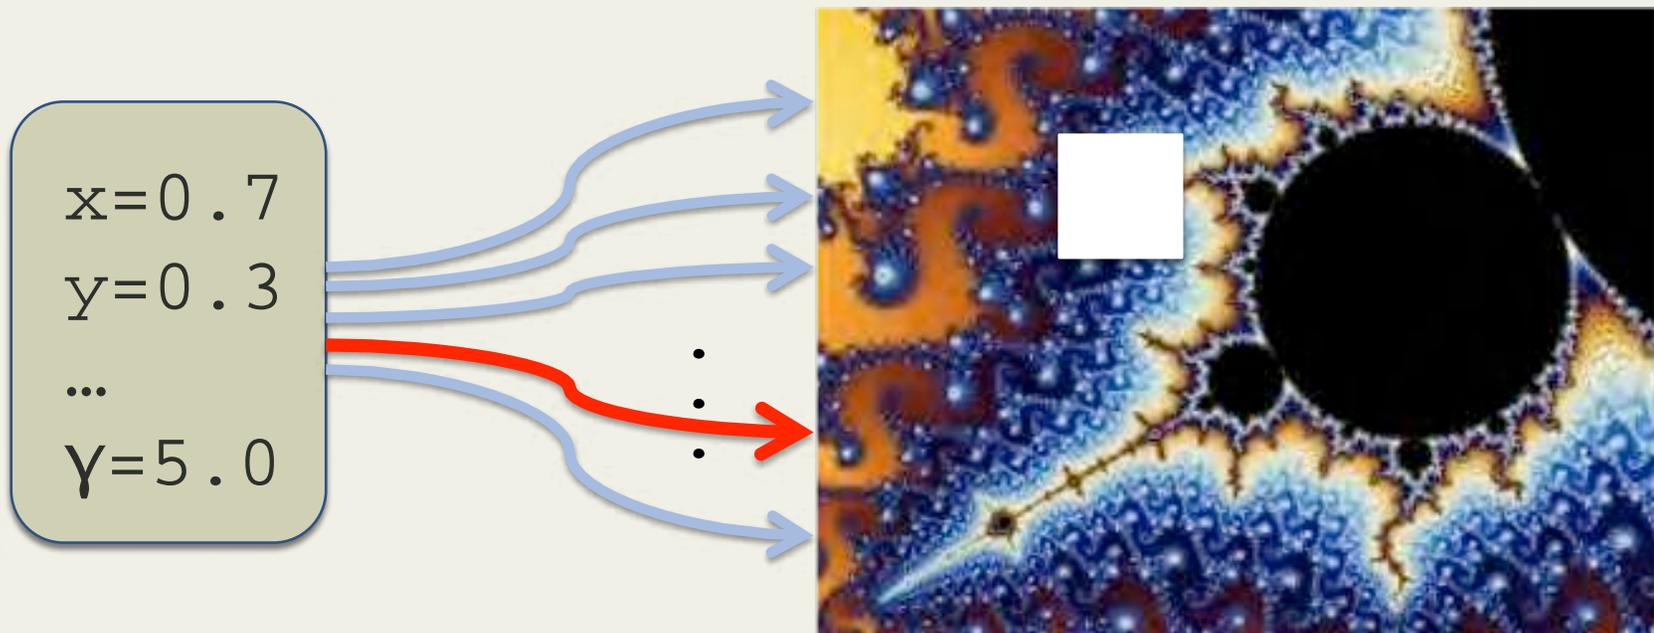
- Want to assert **parallel correctness**.



Background: Deterministic Specs

```
// Parallel fractal render.  
mandelbrot(params, img);
```

- Want to assert **parallel correctness**.



Background: Deterministic Specs

```
deterministic
assume (params == params') {
  // Parallel fractal render.
  mandelbrot(params, img);
} assert (img == img');
```

- Specifies that any two runs on the **same input parameters** must yield the **same output image**.

Background: Deterministic Specs

```
deterministic  
assume (params == params') {  
  // Parallel fractal render.  
  mandelbrot(params, img);  
} assert (img == img');
```

s_0 :

... params ...

s_0' :

... params' ...

Background: Deterministic Specs

```
deterministic  
assume (params == params') {  
  // Parallel fractal render.  
  mandelbrot(params, img);  
} assert (img == img');
```

s_0 :

... params ...

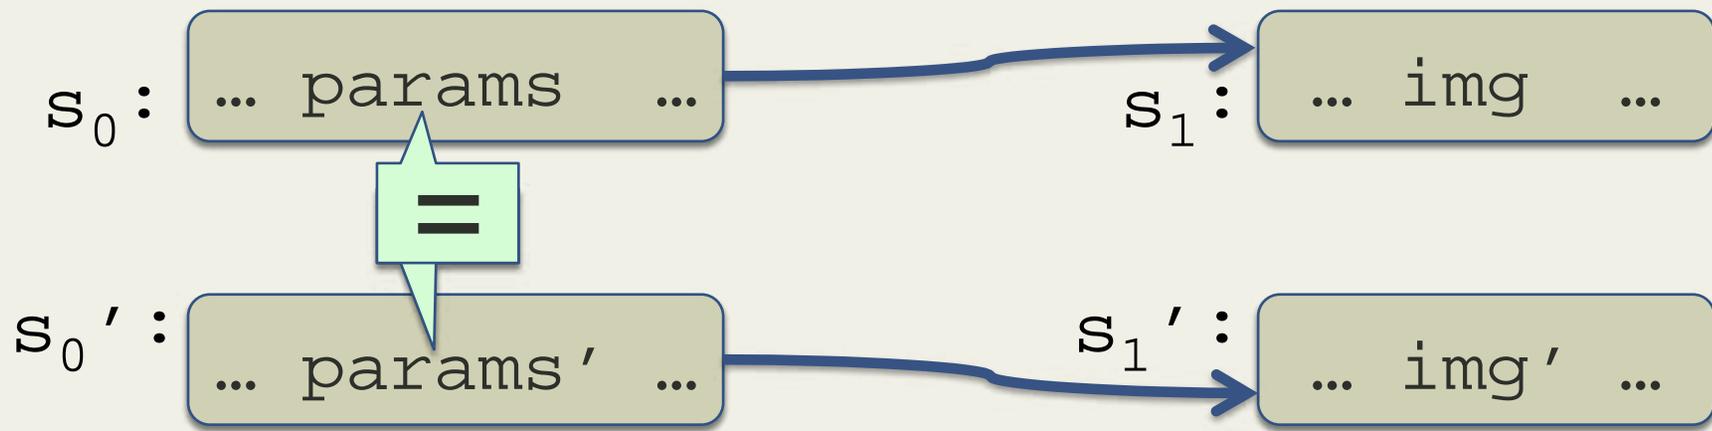
=

s_0' :

... params' ...

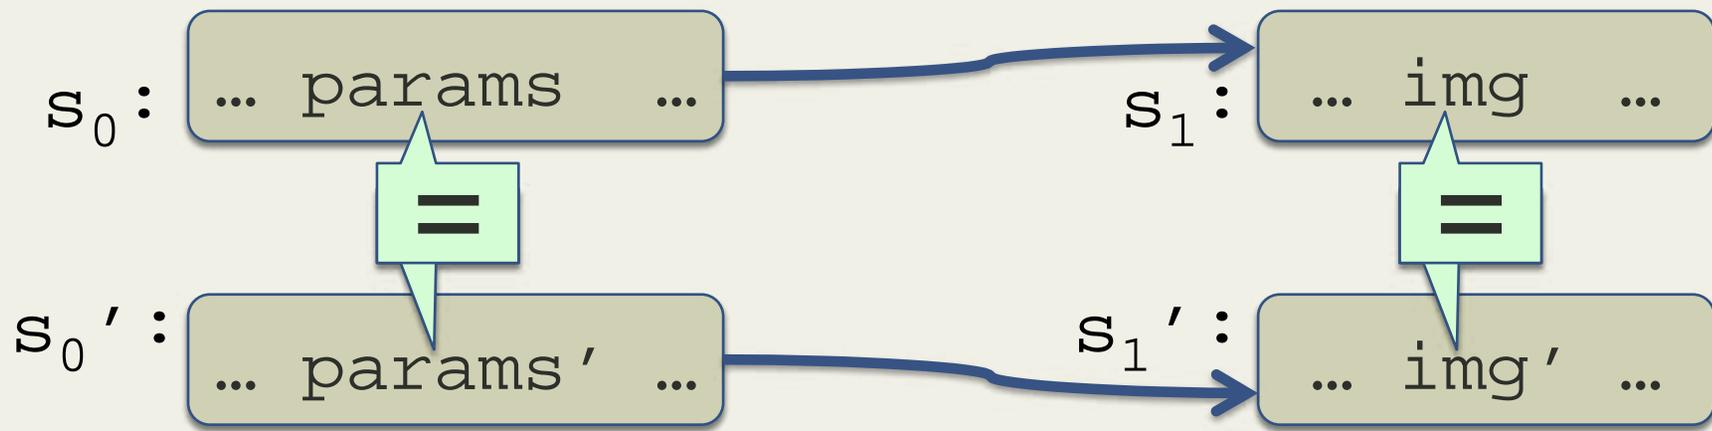
Background: Deterministic Specs

```
deterministic  
assume (params == params') {  
  // Parallel fractal render.  
  mandelbrot(params, img);  
} assert (img == img');
```



Background: Deterministic Specs

```
deterministic  
assume (params == params') {  
  // Parallel fractal render.  
  mandelbrot(params, img);  
} assert (img == img');
```



Background: Deterministic Specs

```
deterministic assume Pre( $s_0, s_0'$ ) {  
  P  
} assert Post( $s_1, s_1'$ )
```

- **Formally, specifies:**

$$\forall s_0 \xrightarrow{P} s_1, s_0' \xrightarrow{P} s_1' : \\ \text{Pre}(s_0, s_0') \Rightarrow \text{Post}(s_1, s_1')$$

Background: Deterministic Specs

```
deterministic  
assume (params == params') {  
  // Parallel fractal render.  
  mandelbrot(params, img);  
} assert (img == img');
```

“Bridge”
predicate

Background: Deterministic Specs

```
deterministic  
assume (params == params') {  
  // Parallel fractal render.  
  mandelbrot(params, img);  
} assert (img == img');
```

“Bridge”
assertion

Background: Deterministic Specs

```
// Parallel fractal render.  
mandelbrot(params, img);
```

- Much simpler than functional correctness:

$$\forall_{0 \leq x < width} \cdot \forall_{0 \leq y < height} \cdot$$

$$\left(\left| f_{iter}^{maxiter}(0) \right| < 2 \wedge img[x][y] = 0 \right)$$

$$\vee \exists_{1 \leq i < maxiter} \cdot \left| f_{iter}^i(0) \right| \geq 2 \wedge \forall_{1 \leq j < i} \cdot \left| f_{iter}^j(0) \right| < 2$$

$$\wedge img[x][y] = \text{HSB}\left((i/maxiter)^y, 1, 1\right)$$

where $f_{iter}(c) = c^2 + (xcenter + (xoff + x)/res)$
 $+ i(ycenter + (yoff - y)/res)$

Background: Deterministic Specs

```
set t = new RedBlackTreeSet();  
...  
deterministic  
assume (t.equals(t')) {  
    t.add(3) || t.add(5);  
} assert (t.equals(t'));
```

- Resulting sets are *semantically* equal.

Background: Deterministic Specs

```
double A[][] , b[] , x[];
...
deterministic
assume (A == A' and b == b') {
    // Solve A*x = b in parallel
    lufact_solve(A, b, x);
} assert (|x - x'| < ε);
```

Background: Deterministic Specs

```
deterministic
assume (data == data') {
  // Parallel branch-and-bound
  Tree t = min_phylo_tree(data);
} assert (t.cost == t'.cost());
```

- Determinism is **user-specified**.

Background: Deterministic Specs

- Can effectively test deterministic specs.
 - Added assertions to 13 benchmarks.
 - Ran CalFuzzer to test if concurrency issues (data races, atomicity violations, etc.) could lead to violations of deterministic spec.
- Specification inference would help automate the above testing.
- Also aid program understanding.

Outline

- Motivation and Overview
- Background: Deterministic Specs
- **Specification Inference Problem**
- Inferring Deterministic Specifications
- Experimental Evaluation
- Related Work
- Conclusions

Specification Inference Problem

```
// Parallel branch-and-bound  
(tree, cost) =  
    min_phylo_tree(N, data);
```

Specification Inference Problem

```
deterministic assume (???) {  
  // Parallel branch-and-bound  
  (tree, cost) =  
    min_phylo_tree(N, data);  
} assert (???)
```

Specification Inference Problem

```
deterministic assume (???) {  
  // Parallel branch-and-bound  
  (tree, cost) =  
    min_phylo_tree(N, data);  
} assert (???)
```

- **Observation:** Deterministic pre- and postcondition have simple structure.
 - Conjunction of **equality bridge predicates:**

$$N = N' \quad tree.equals(tree') \quad |cost - cost'| < \epsilon$$

Specification Inference Problem

```
deterministic assume (???) {  
  // Parallel branch-and-bound  
  (tree, cost) =  
    min_phylo_tree(N, data);  
} assert (???)
```

- Four possible deterministic preconditions:

true

data = data'

N = N'

data = data' \wedge N = N'

Specification Inference Problem

```
deterministic assume (???) {  
  // Parallel branch-and-bound  
  (tree, cost) =  
    min_phylo_tree(N, data);  
} assert (???)
```

- Six possible deterministic postconditions:

true $|cost - cost'| < \epsilon$ $cost = cost'$

$tree.equals(tree')$ $tree.equals(tree')$ $tree.equals(tree')$
 $\wedge |cost - cost'| < \epsilon$ $\wedge cost = cost'$

Specification Inference Problem

```
deterministic assume (???) {  
  // Parallel branch-and-bound  
  (tree, cost) =  
    min_phylo_tree(N, data);  
} assert (???)
```

- Six
true

Which specification
should we choose?

itions:

st'

tree.equals(tree')

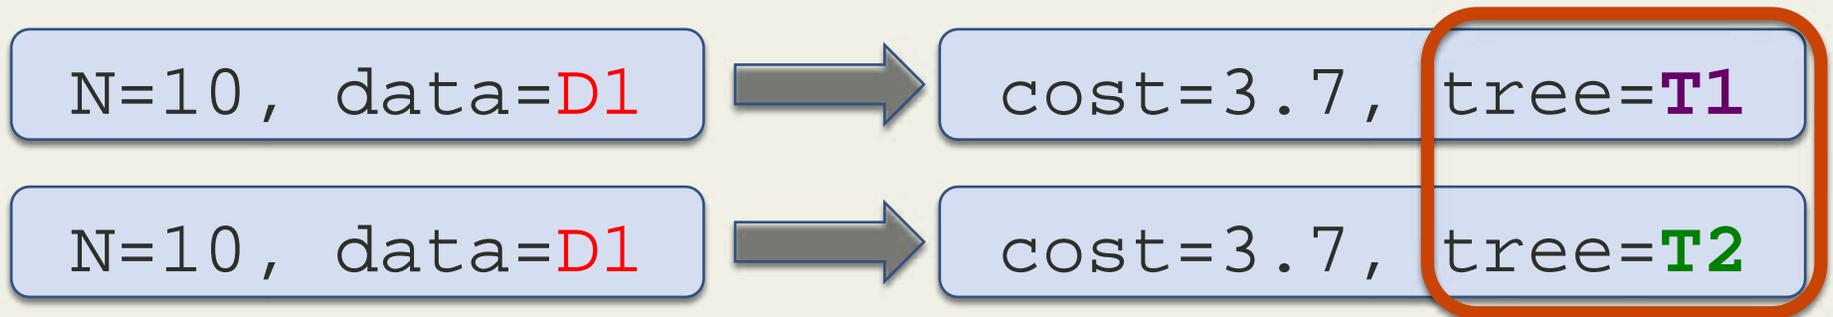
$\wedge |cost - cost'| < \epsilon$

$\wedge cost = cost'$

tree.equals(tree')

Specification Inference Problem

- Principles for specification inference:
 - Must be consistent with observed runs.**

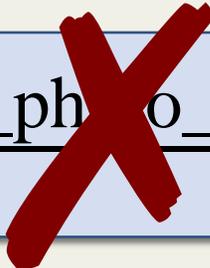


$data = data' \xrightarrow{\text{min_ph} \times \text{o_tree}} tree.equals(tree')$

Specification Inference Problem

- Principles for specification inference:
 - Must be consistent with observed runs.
 - **Postcondition as strong as possible.**

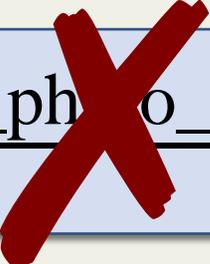
$$data = data' \xrightarrow{\text{min_phylo_tree}} |cost - cost'| < \varepsilon$$

$$data = data' \xrightarrow{\text{min_phylo_tree}} \text{true}$$


Specification Inference Problem

- Principles for specification inference:
 - Must be consistent with observed runs.
 - Postcondition as strong as possible.
 - **Precondition as weak as possible, *for post***

$$data = data' \xrightarrow{\text{min_phylo_tree}} |cost - cost'| < \epsilon$$

$$data = data' \wedge N = N' \xrightarrow{\text{min_phylo_tree}} |cost - cost'| < \epsilon$$


Specification Inference Problem

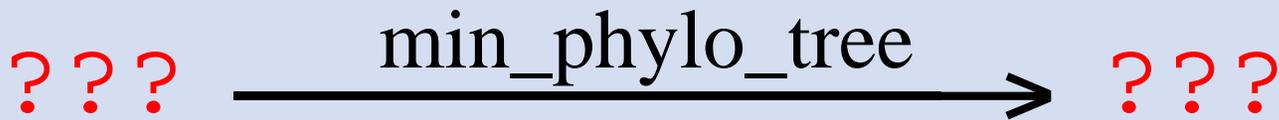
- Principles for specification inference:
 - Must be consistent with observed runs.
 - Postcondition as strong as possible.
 - Precondition as weak as possible, *for post*.

How do we compute such a deterministic specification?

Outline

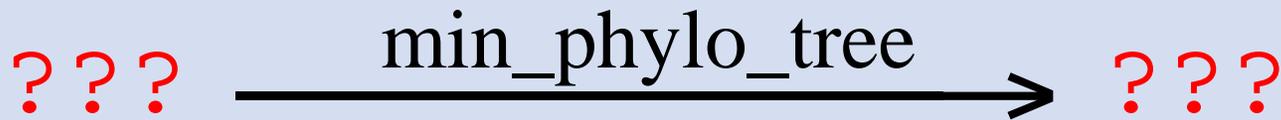
- Motivation and Overview
- Background: Deterministic Specs
- Specification Inference Problem
- **Inferring Deterministic Specifications**
- Experimental Evaluation
- Related Work
- Conclusions

Inferring Deterministic Specs

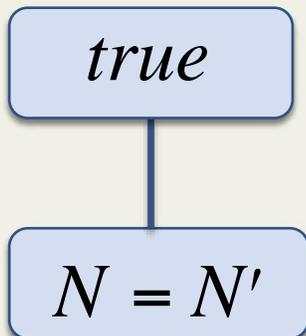


- M variables $\Rightarrow \Omega(2^M)$ specifications
 - Exhaustive search infeasible.
- Two-step algorithm:
 - Compute strongest consistent postcondition.
 - Compute weakest consistent precondition for the inferred postcondition.

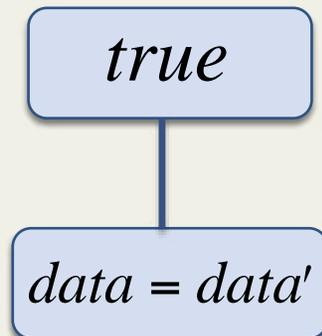
Inferring Deterministic Specs



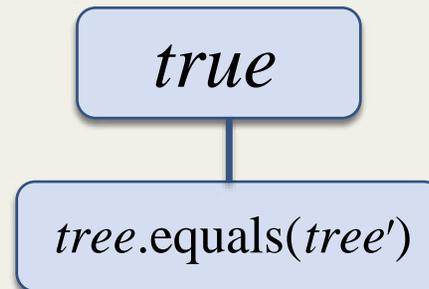
N:



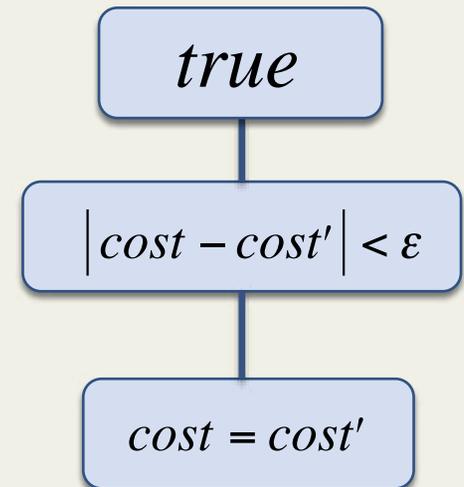
data:



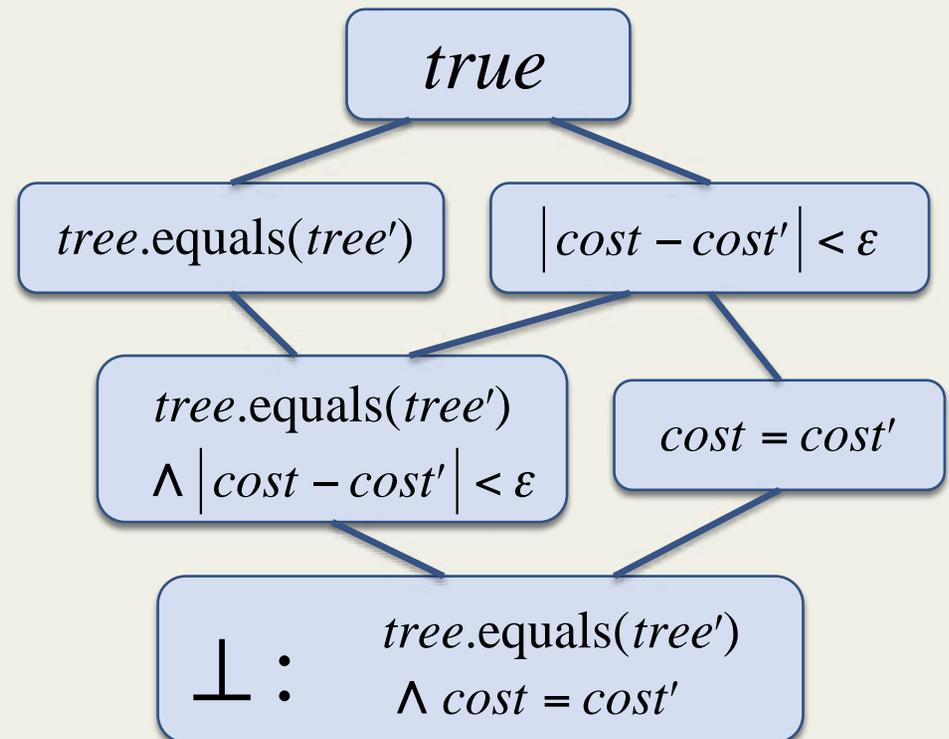
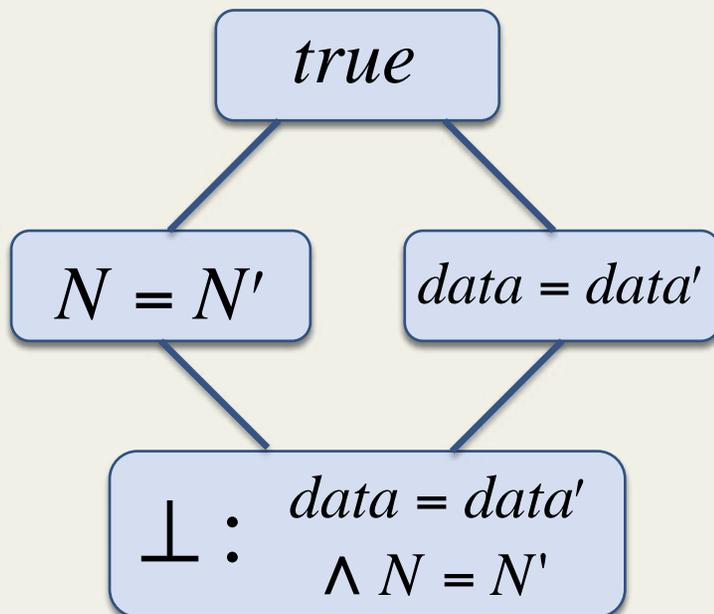
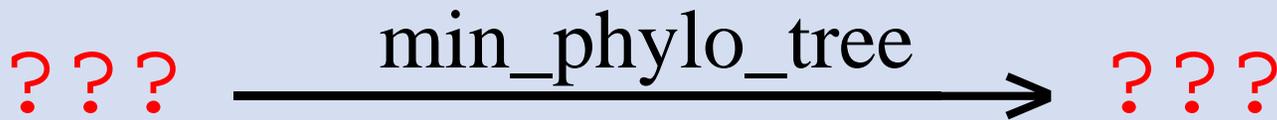
tree:



cost:

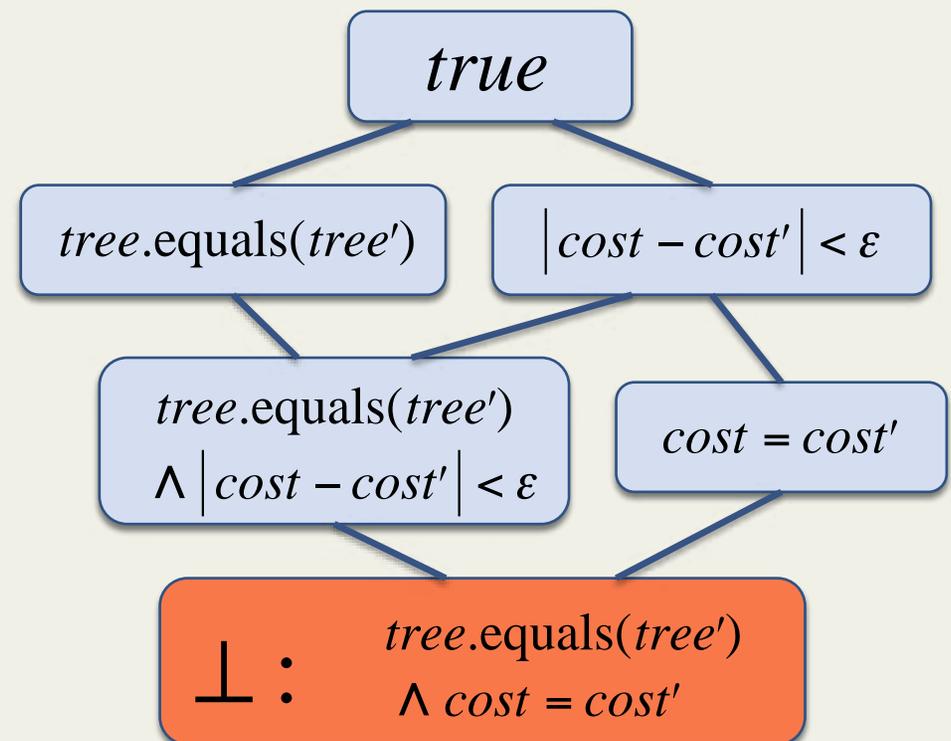
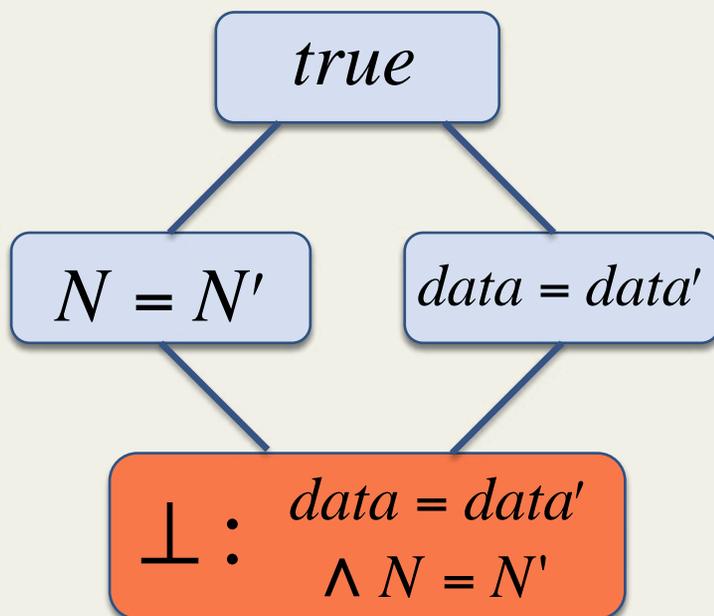


Inferring Deterministic Specs



Inferring Strongest Post

- For every pair of observed executions:
 - If they satisfy precondition \perp , ensure that the postcondition holds.



Inferring Strongest Post I

$N=10$, $data=D1$

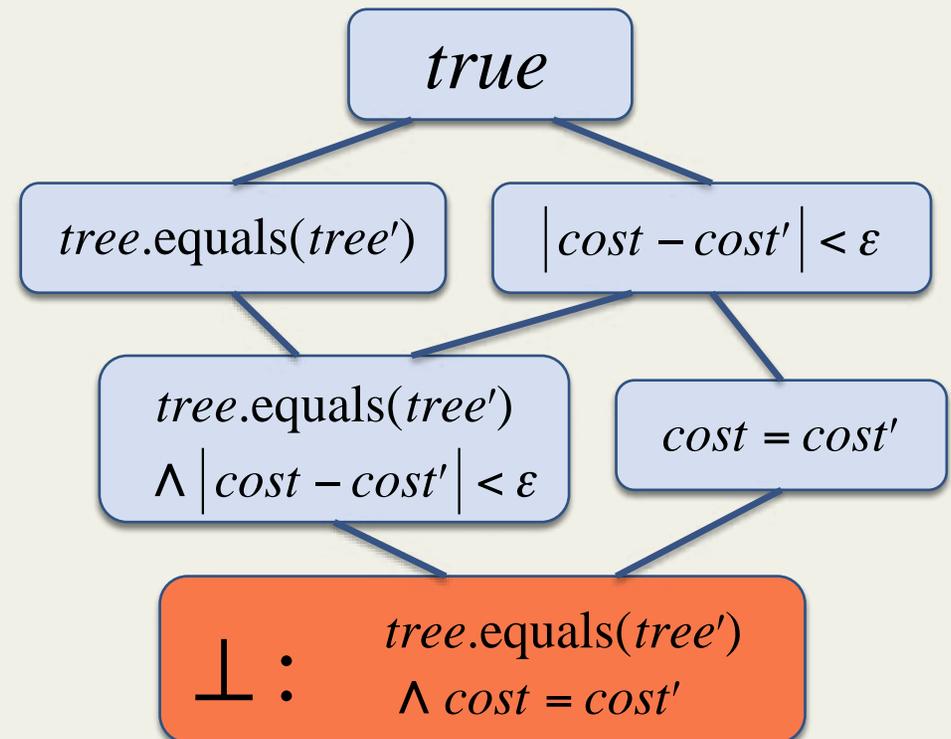
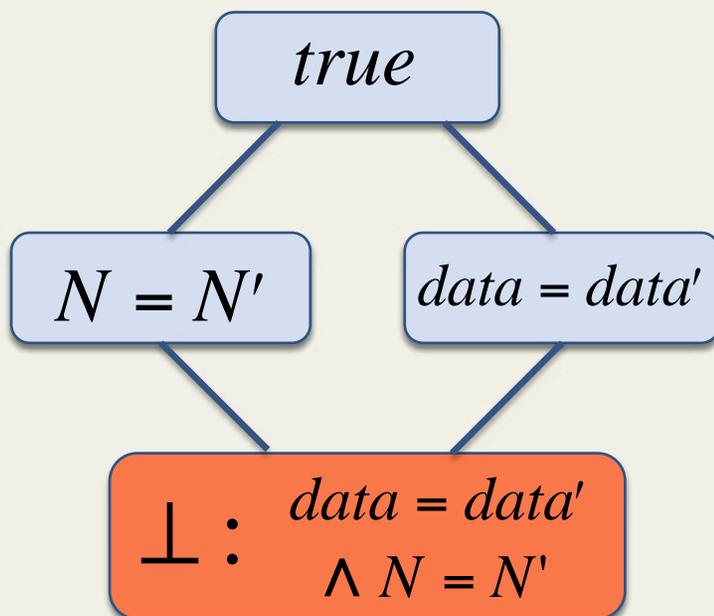


$cost=3.7$, $tree=T1$

$N=10$, $data=D1$



$cost=3.7$, $tree=T2$



Inferring Strongest Post I

$N=10$, $data=D1$

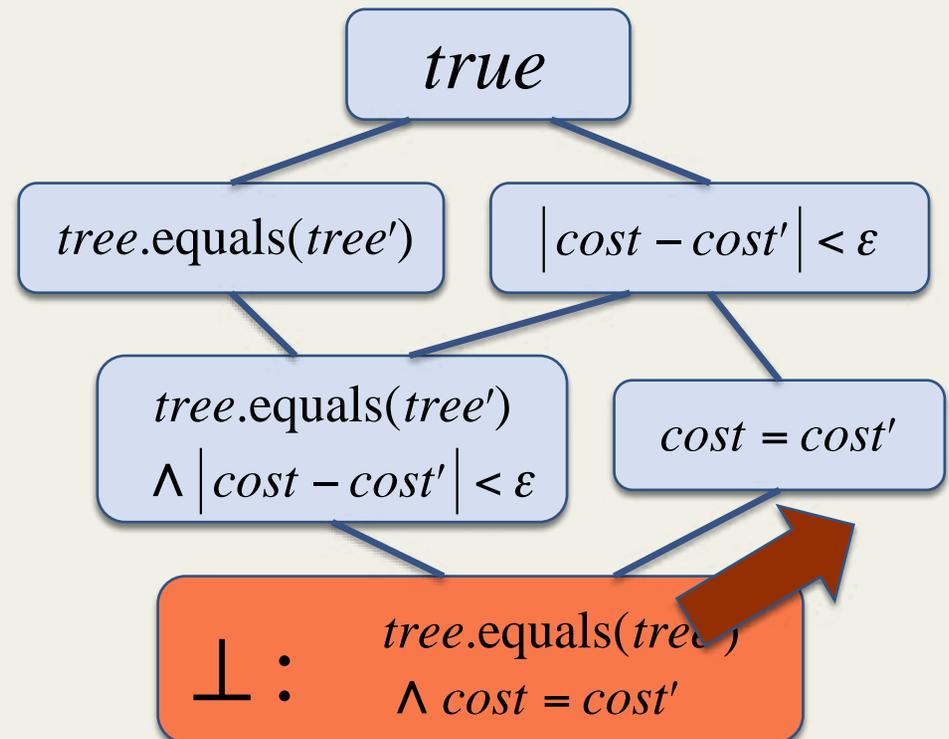
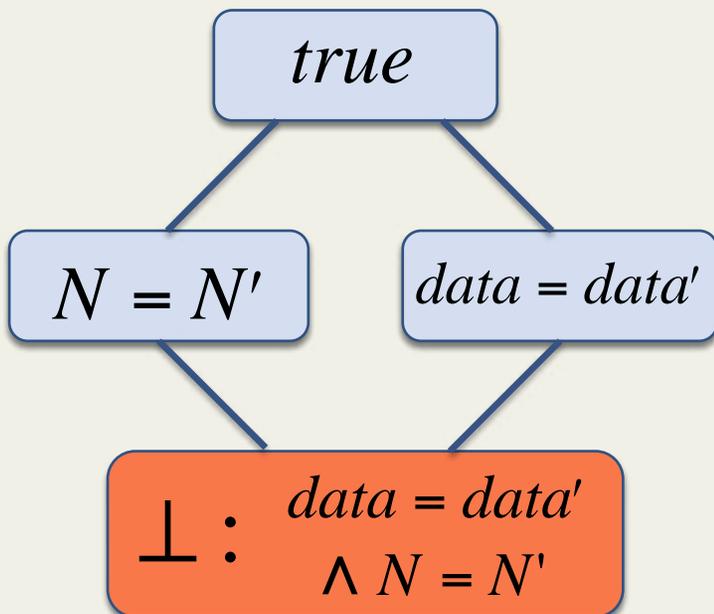


$cost=3.7$, $tree=T1$

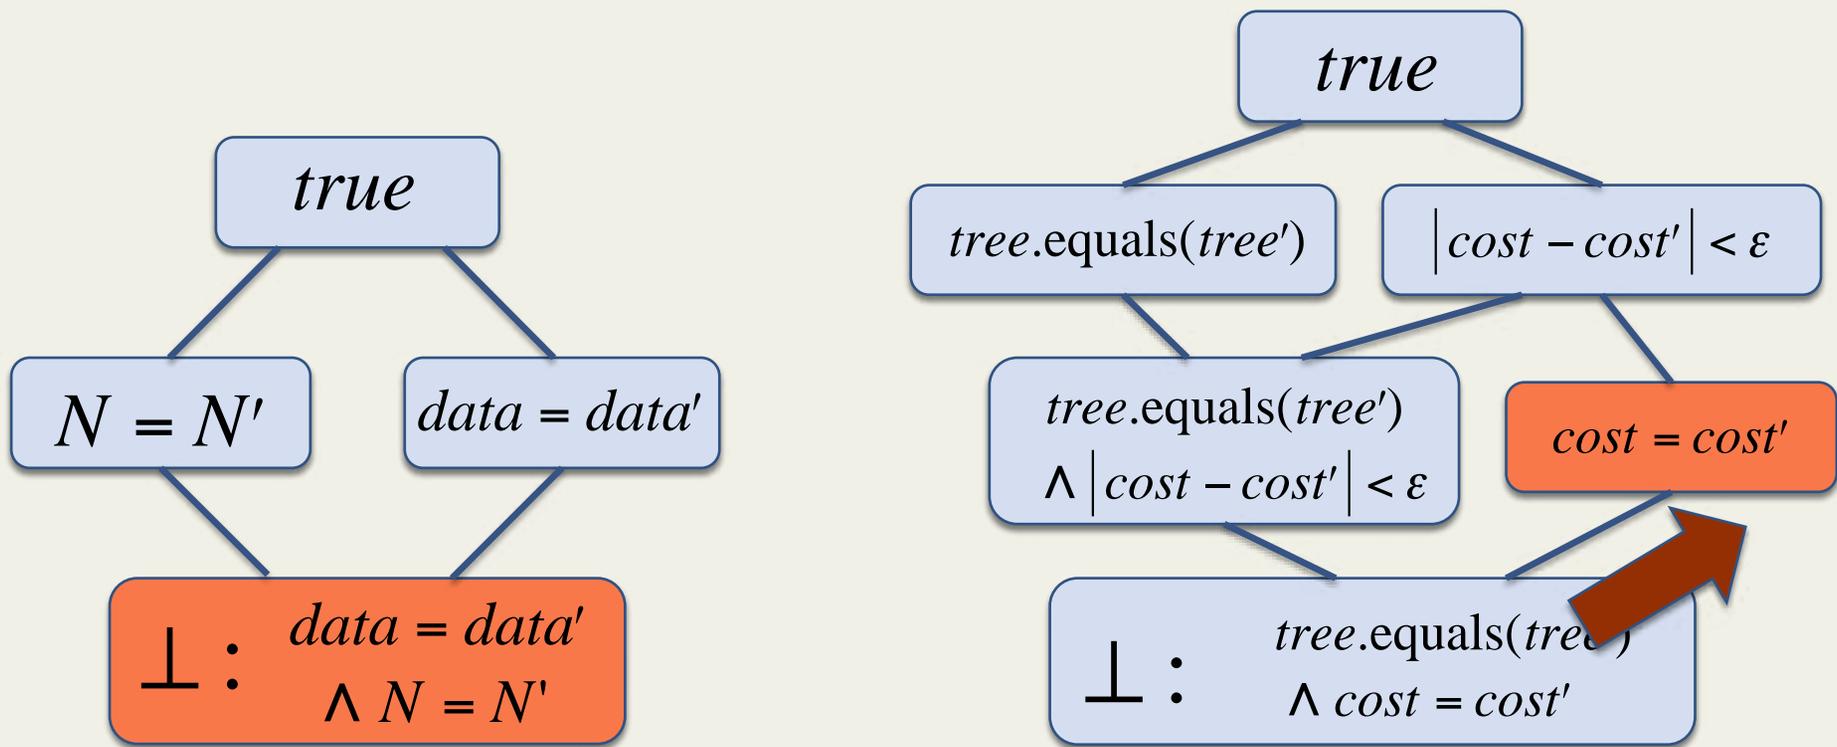
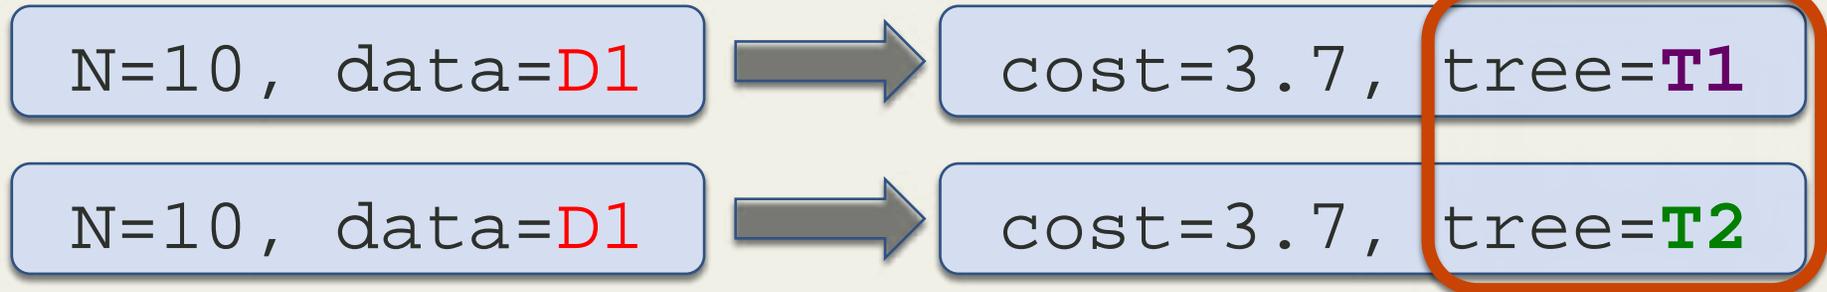
$N=10$, $data=D1$



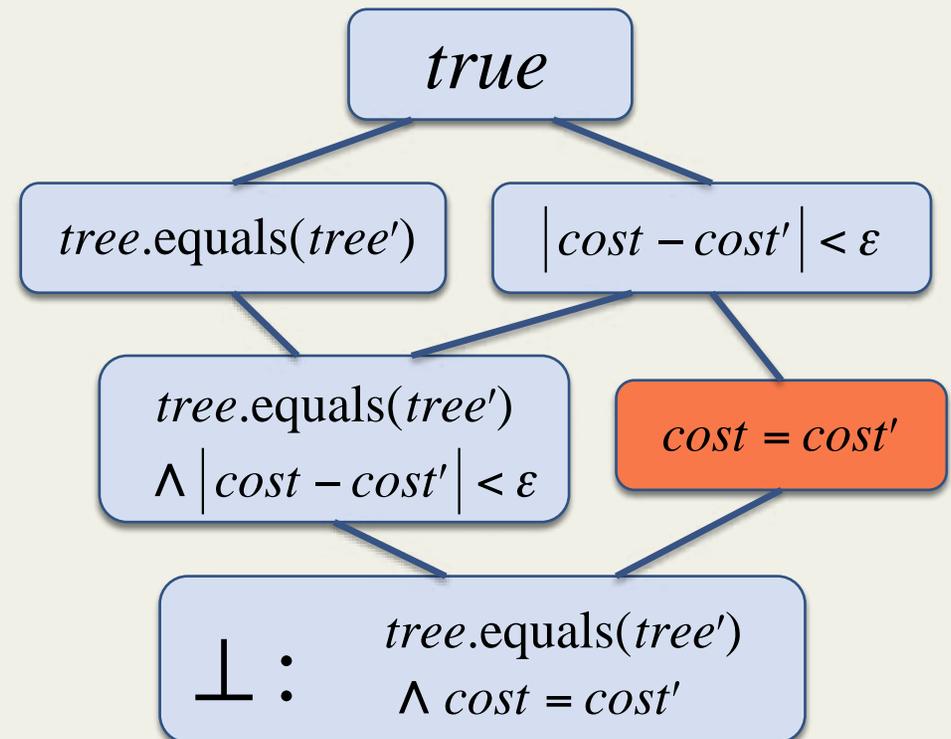
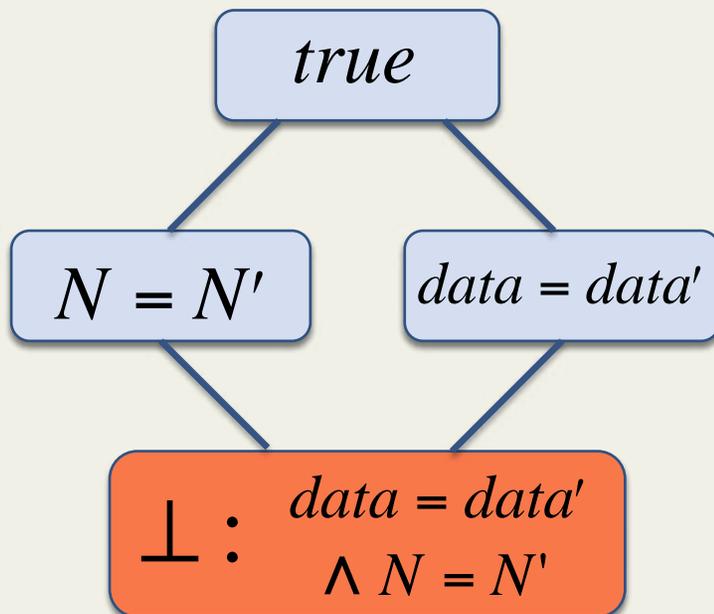
$cost=3.7$, $tree=T2$



Inferring Strongest Post I

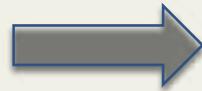


Inferring Strongest Post I



Inferring Strongest Post II

$N=10$, data=**D1**

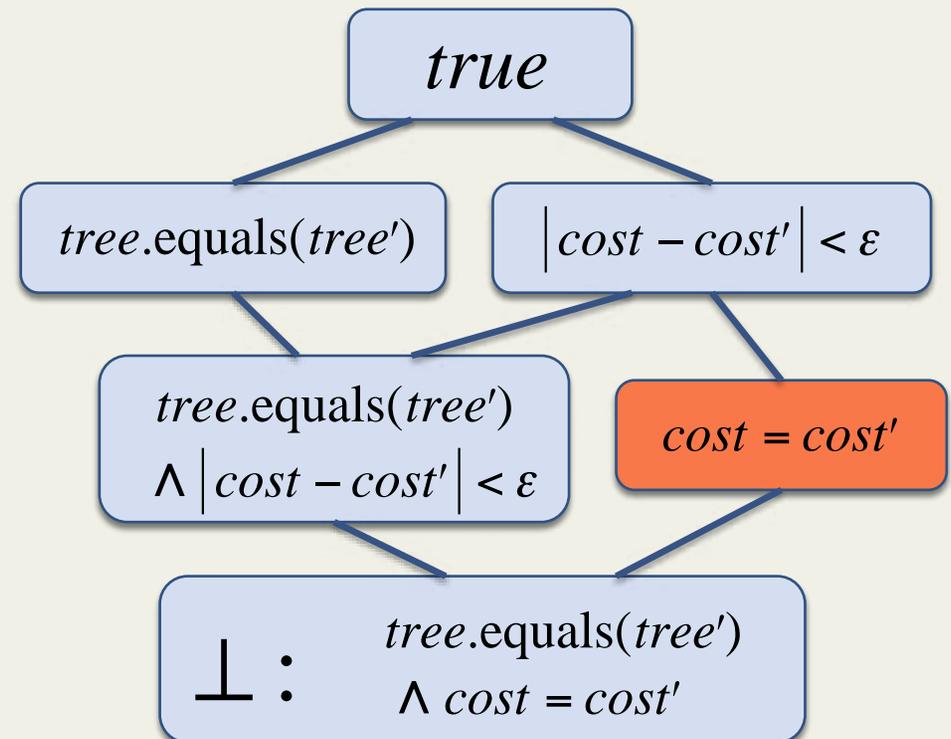
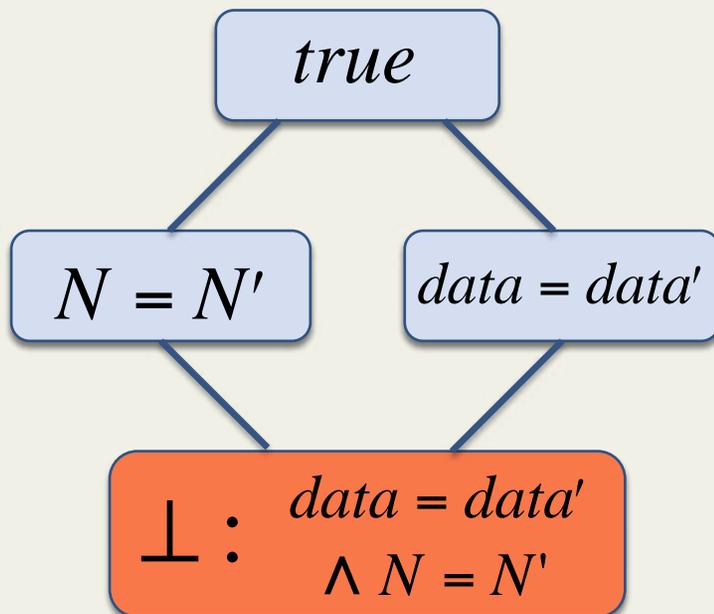


cost=3.7, tree=**T1**

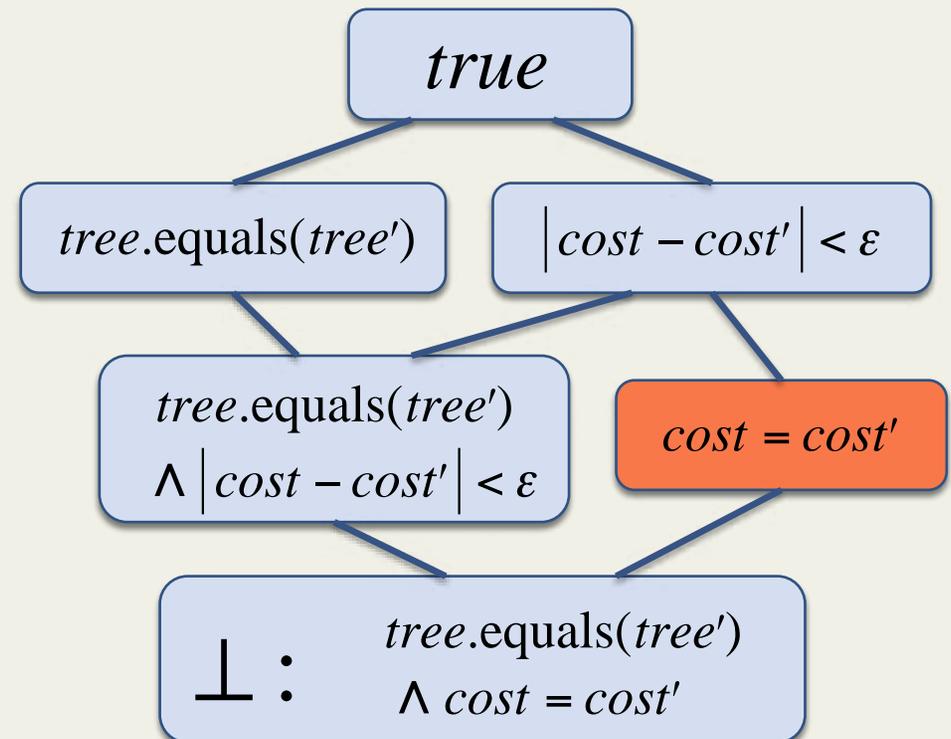
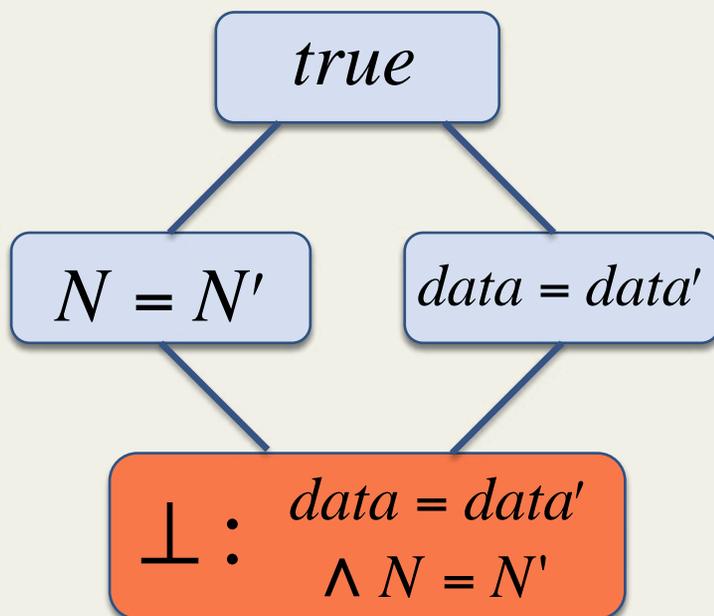
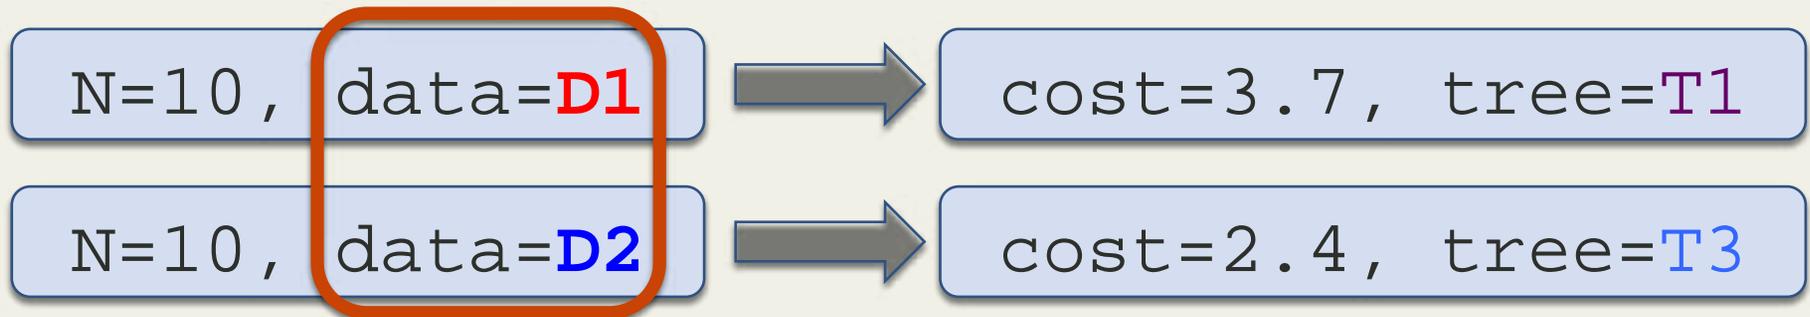
$N=10$, data=**D2**



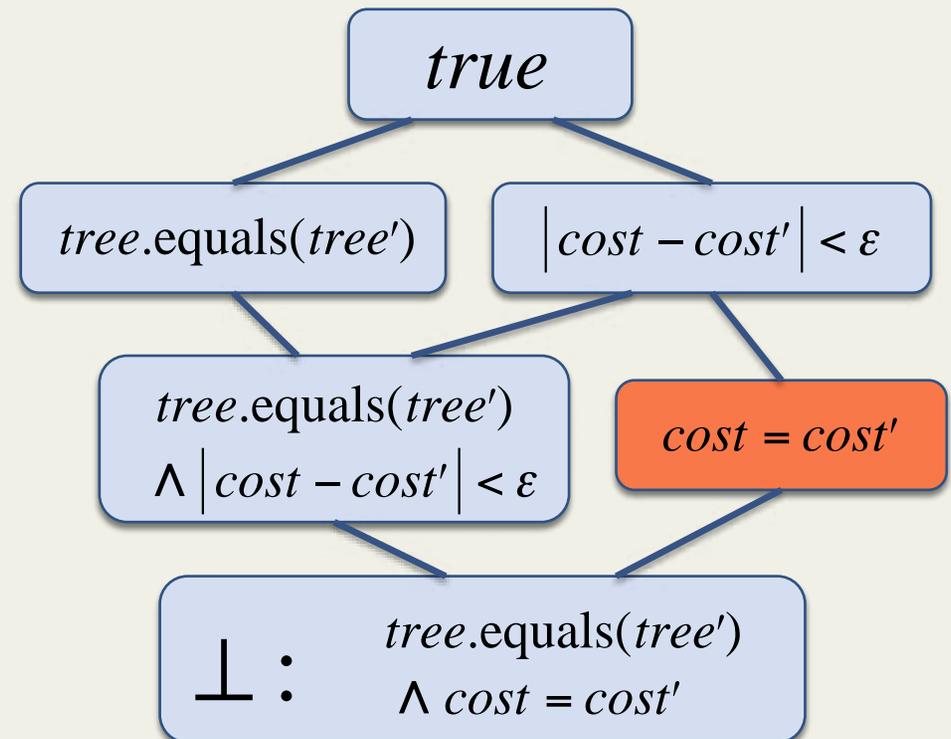
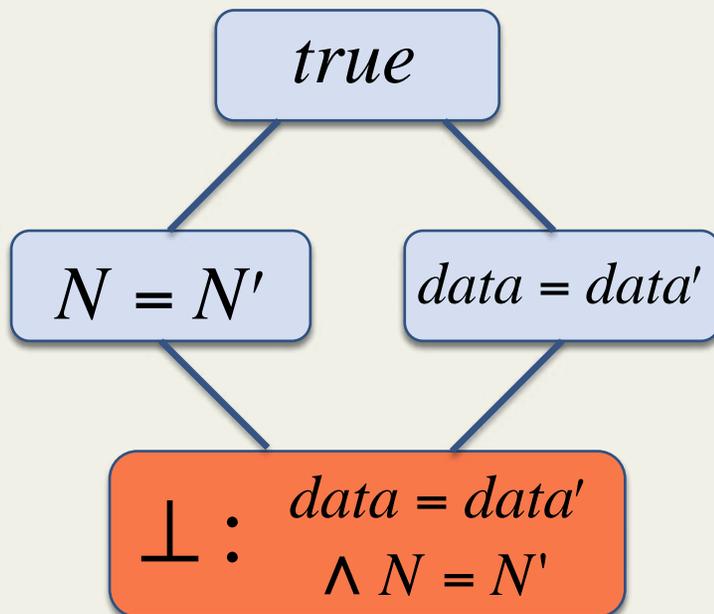
cost=2.4, tree=**T3**



Inferring Strongest Post II



Inferring Strongest Post II



Inferring Strongest Post III

$N=10$, $data=D2$

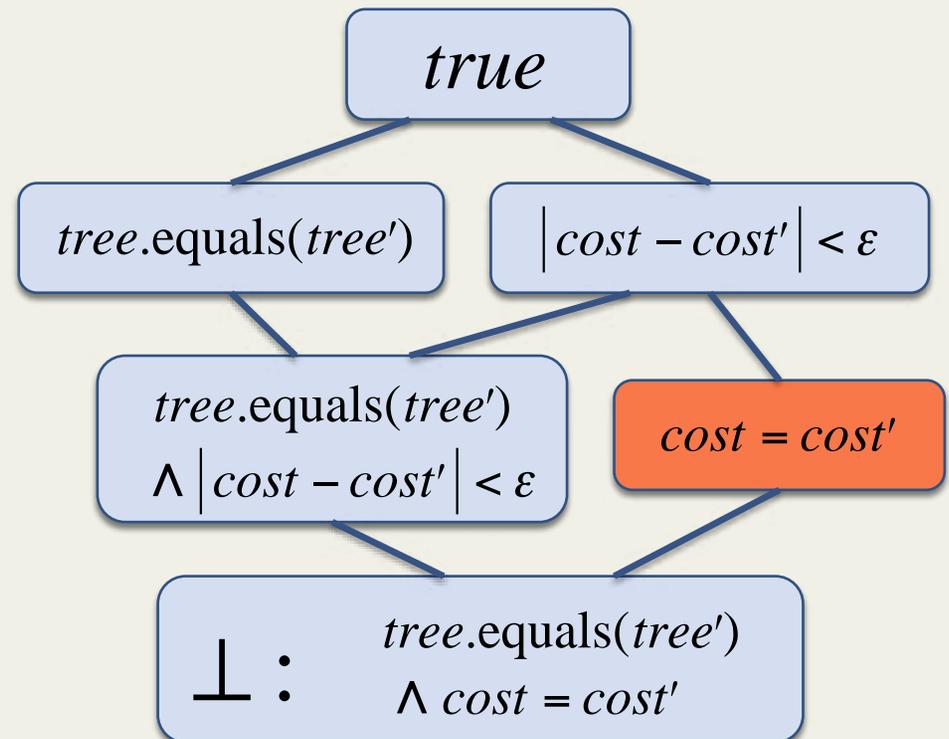
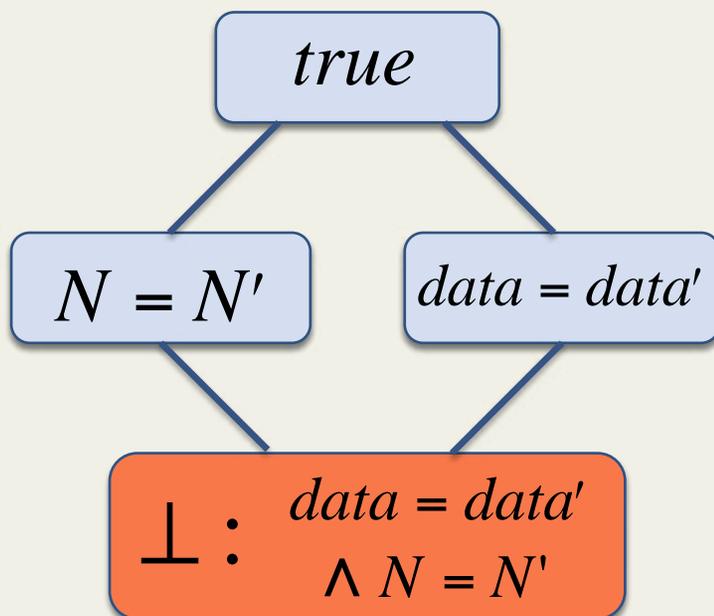


$cost=2.4$, $tree=T3$

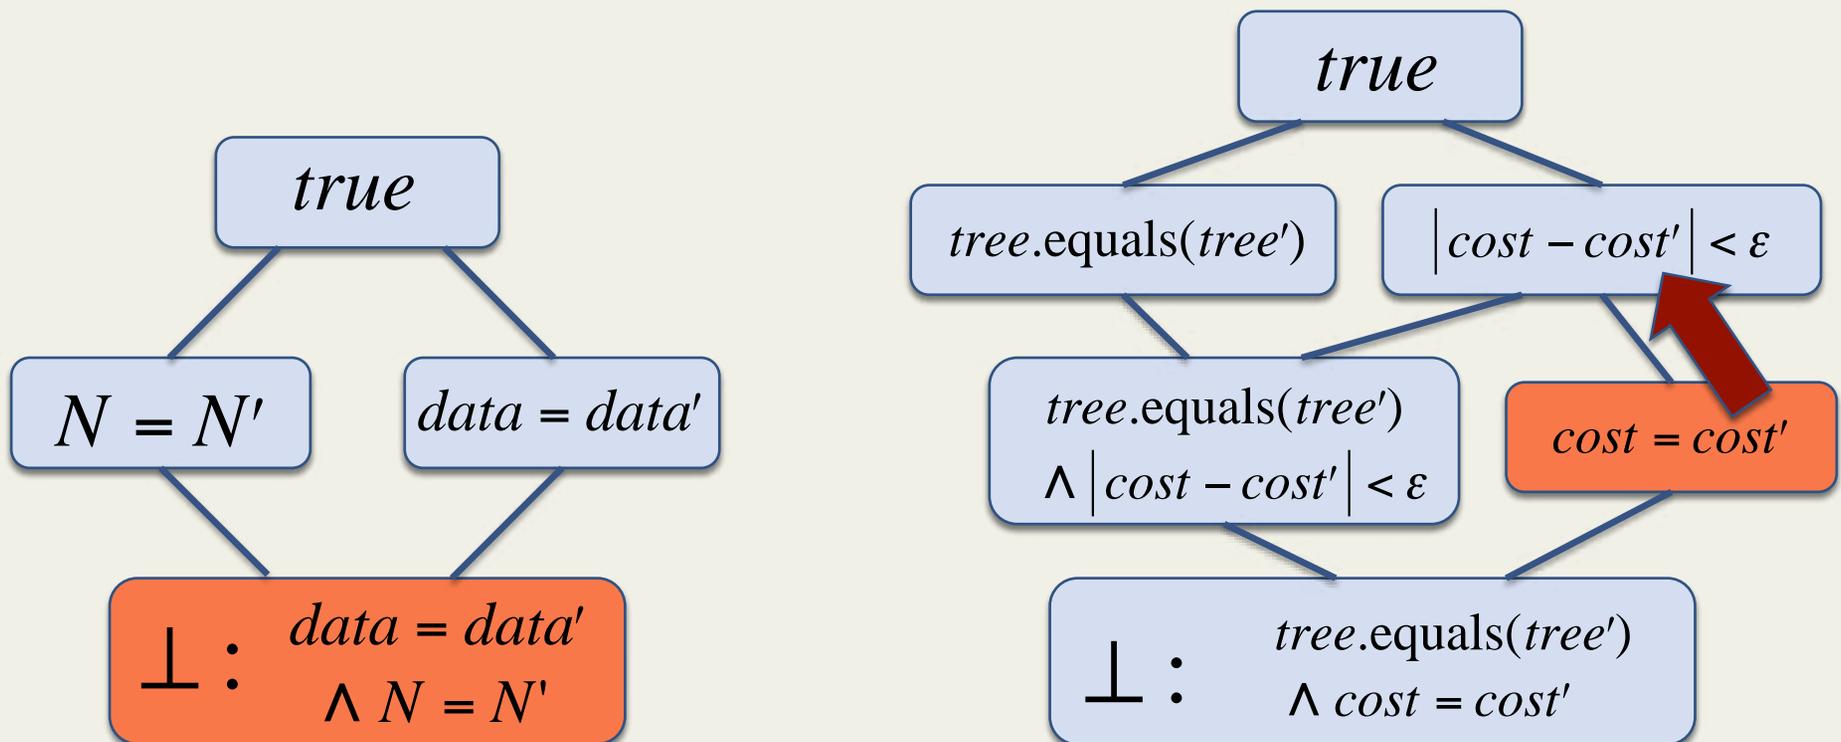
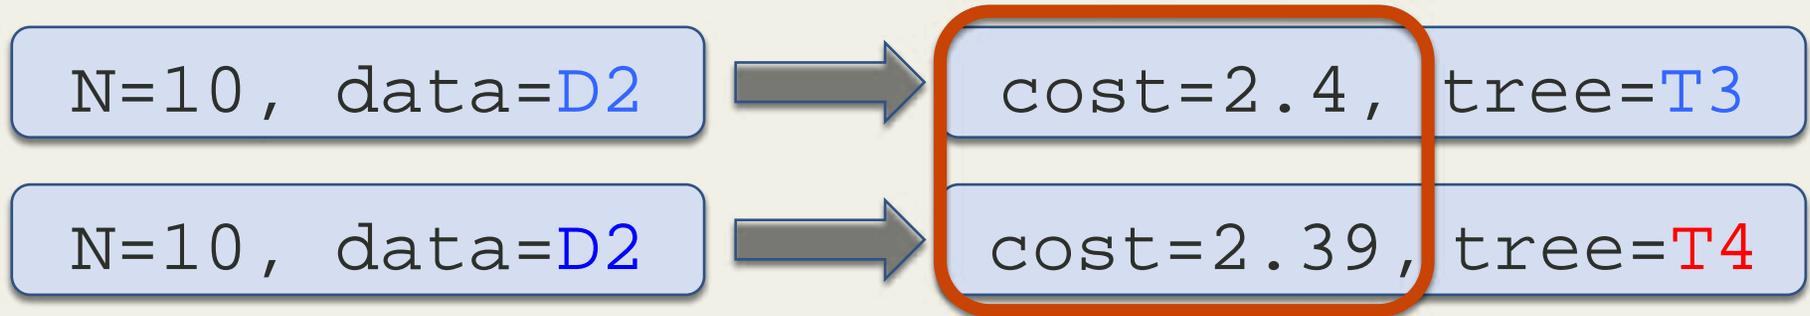
$N=10$, $data=D2$



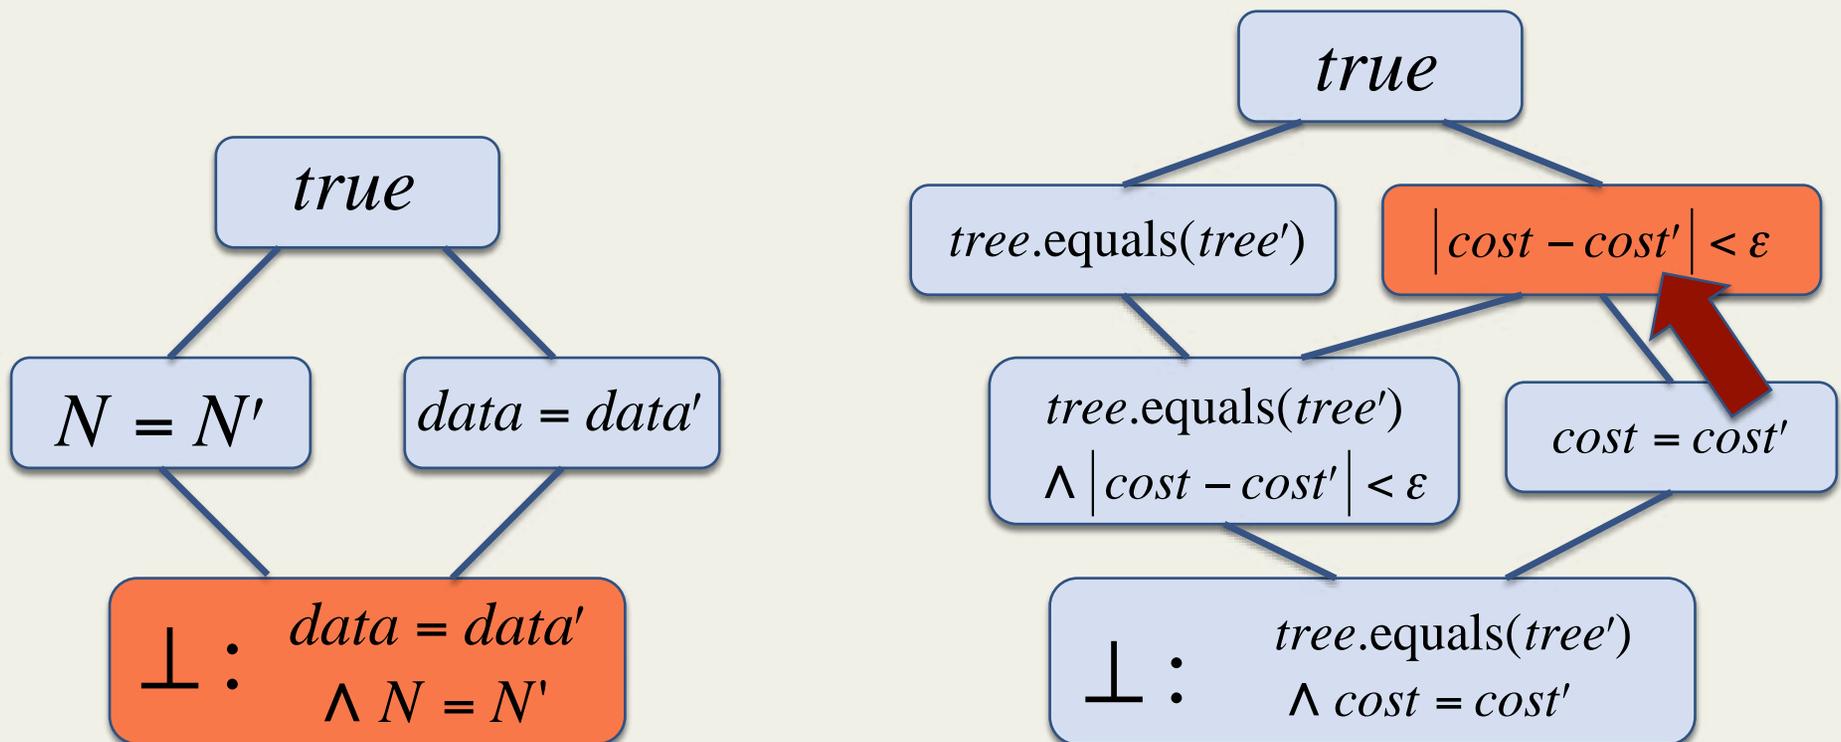
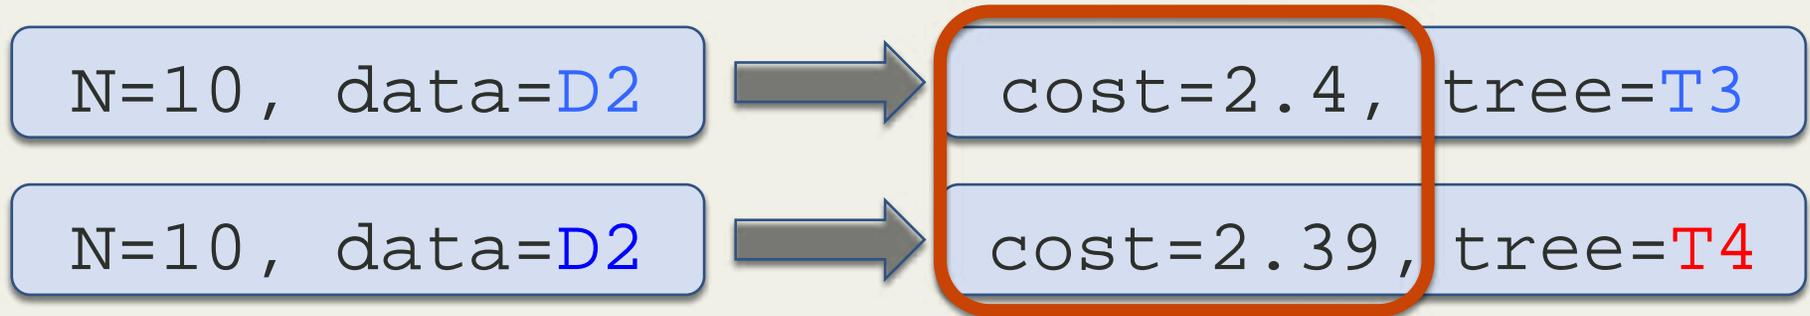
$cost=2.39$, $tree=T4$



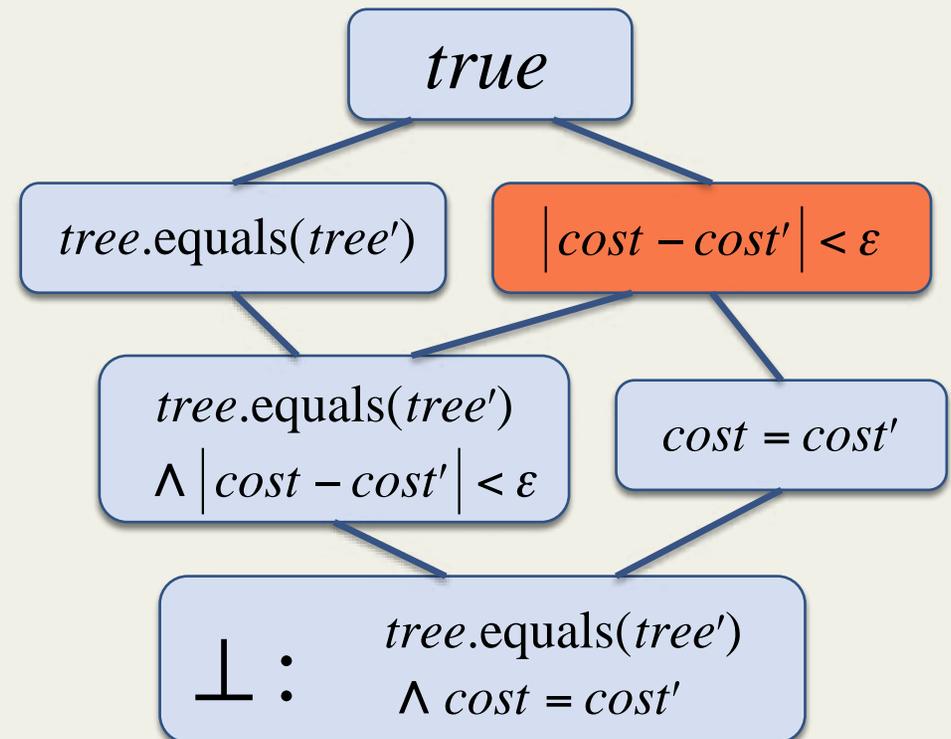
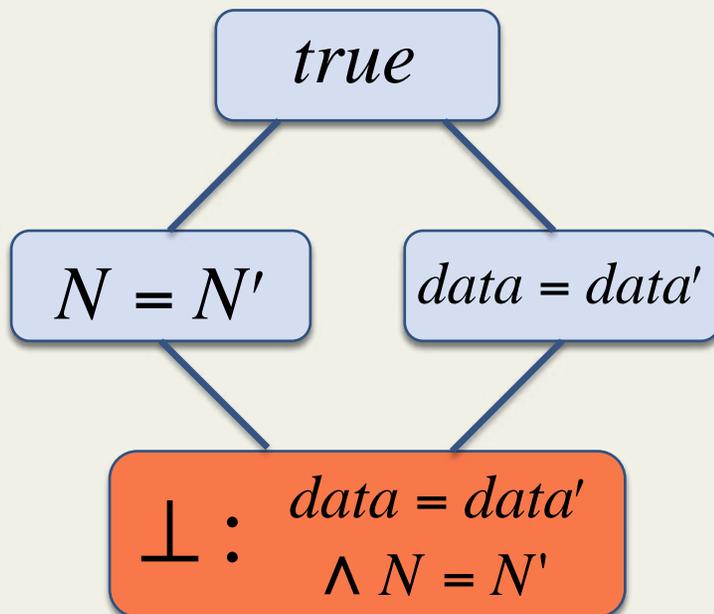
Inferring Strongest Post III



Inferring Strongest Post III



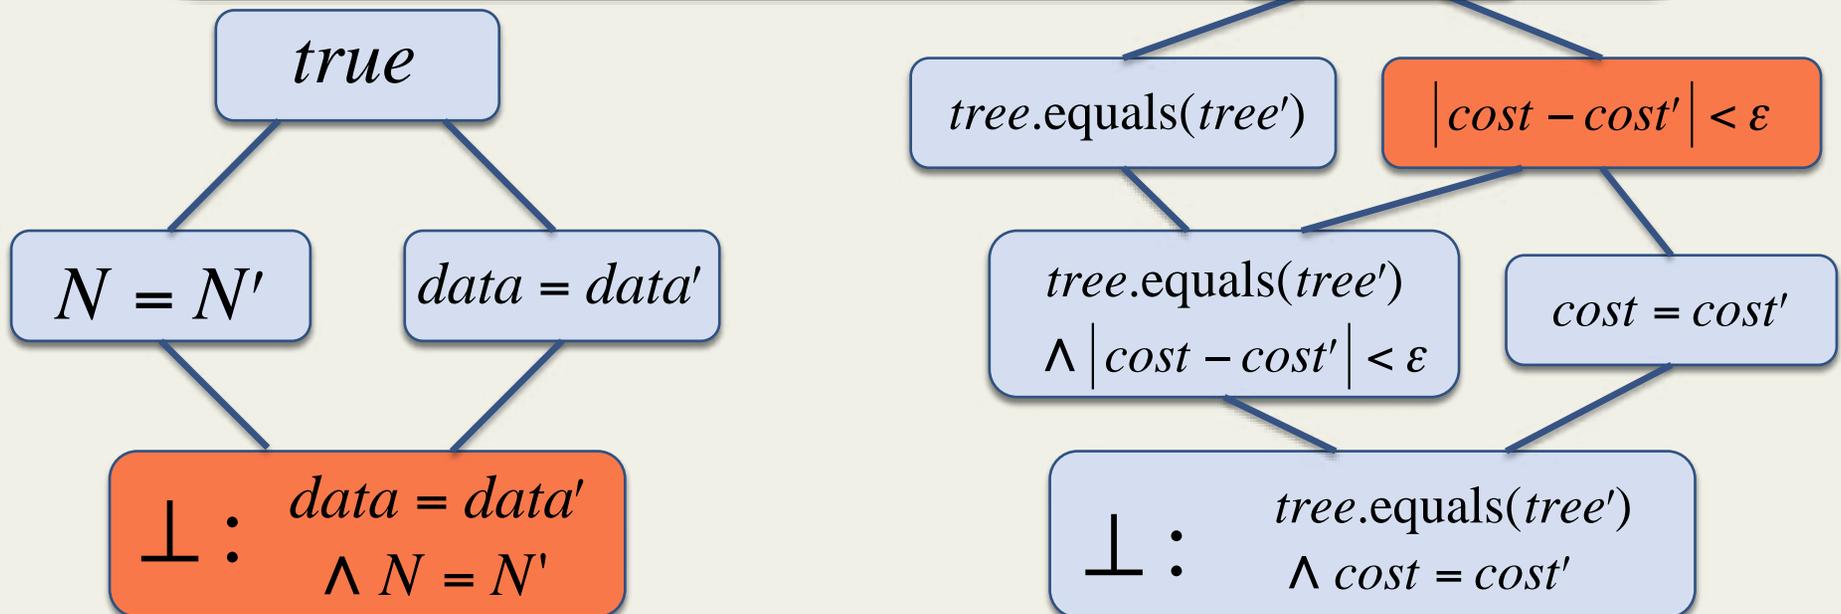
Inferred Strongest Post



Inferred Strongest Post

Proposition 1. The inferred postcondition $Post_R$ is the unique strongest postcondition given the observed runs.

I.e., if $\varphi_{pre} \xrightarrow{P,R} \varphi_{post}$ then $Post_R \Rightarrow \varphi_{post}$.



Inferred Strongest Post

Proposition 1. The inferred postcondition $Post_R$ is the unique strongest postcondition given the observed runs.

I.e., if $\varphi_{pre} \xrightarrow{P,R} \varphi_{post}$ then $Post_R \Rightarrow \varphi_{post}$.

true

Corollary 2. Let the inferred specification be $Pre_R \rightarrow Post_R$. Then, $Post_R$ is the strongest postcondition of Pre_R .

I.e., if $Pre_R \xrightarrow{P,R} \varphi_{post}$ then $Post_R \Rightarrow \varphi_{post}$.

$\leq \varepsilon$

post'

$\wedge N = N'$

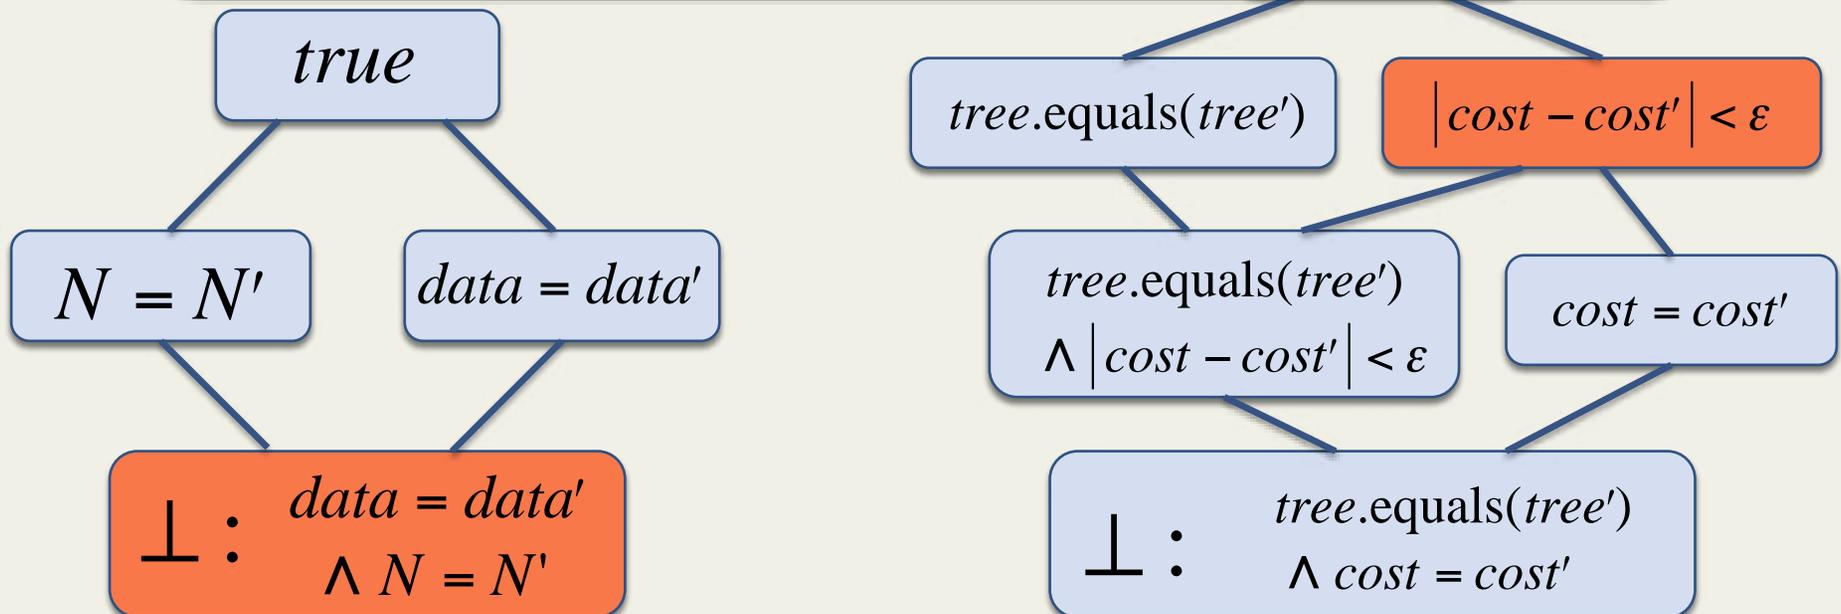
$\wedge cost = cost'$

N

Inferred Strongest Post

Proposition 3. The inferred postcondition $Post_R$ is at least as strong as the true unique strongest postcondition $SP_P(\perp)$.

I.e., if $\varphi_{pre} \xrightarrow{P} \varphi_{post}$ then $Post_R \Rightarrow \varphi_{post}$.



Inferred Strongest Post

Proposition 3. The inferred postcondition $Post_R$ is at least as strong as the true unique strongest postcondition $SP_P(\perp)$.

I.e., if $\varphi_{pre} \xrightarrow{P} \varphi_{post}$ then $Post_R \Rightarrow \varphi_{post}$.

true

Proposition 4. More observed executions lead to a weaker inferred postcondition.

That is, if $R_1 \subseteq R_2 \subseteq R_3 \subseteq \dots$ then

$Post_{R_1} \Rightarrow Post_{R_2} \Rightarrow \dots \Rightarrow SP_P(\perp)$.

$\wedge N = N'$

$\wedge cost = cost'$

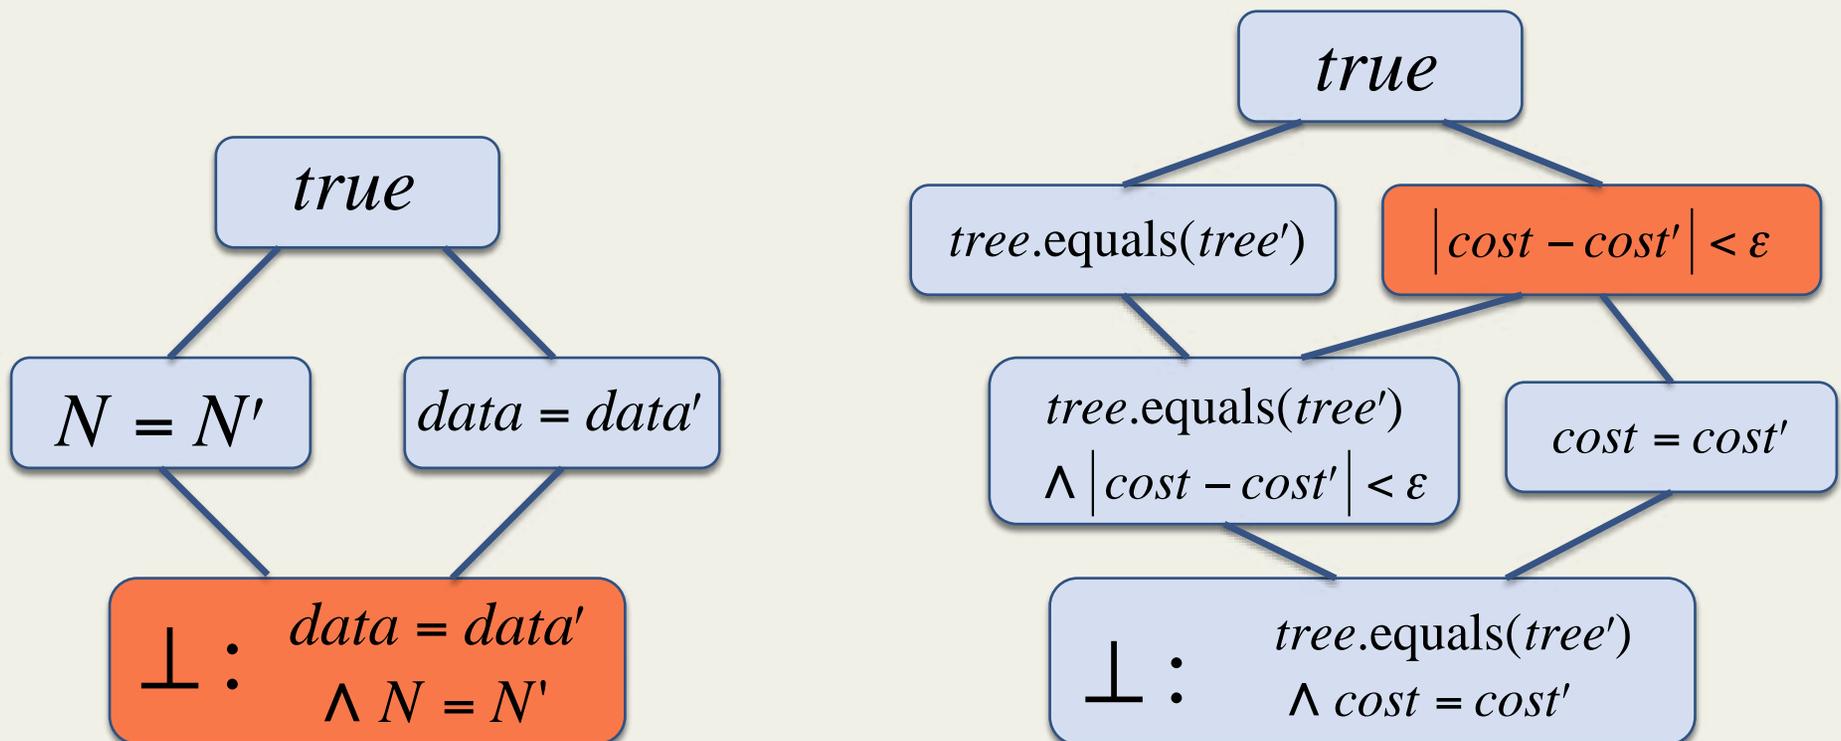
$< \epsilon$

post'

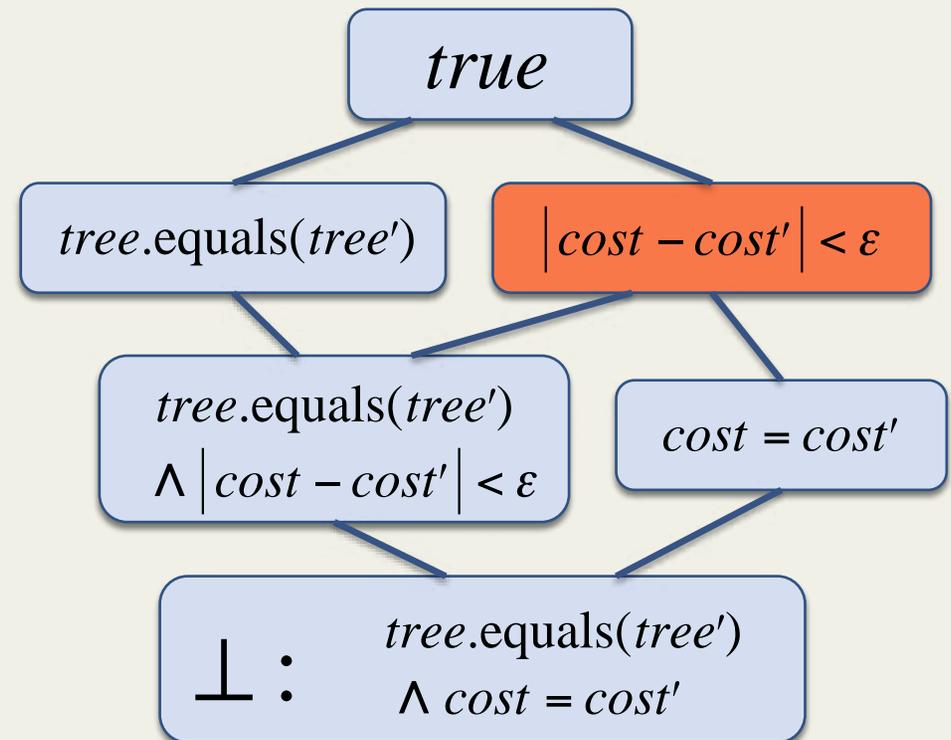
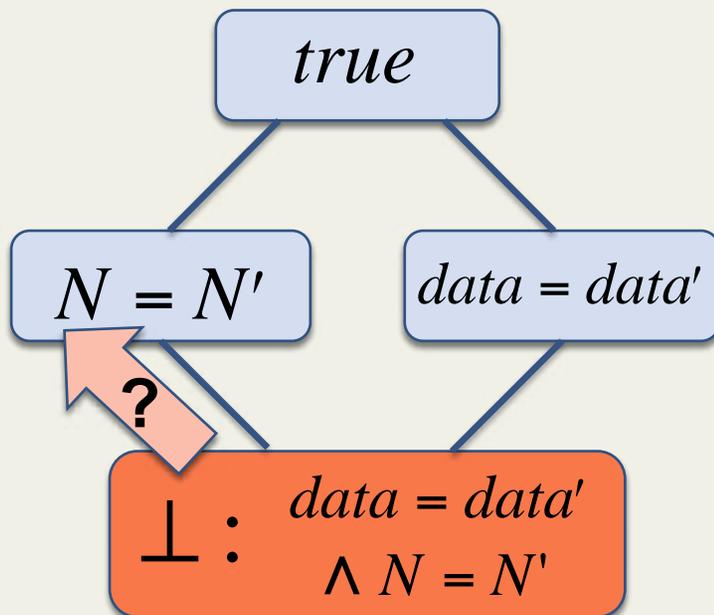
N

Inferring Weakest Precondition

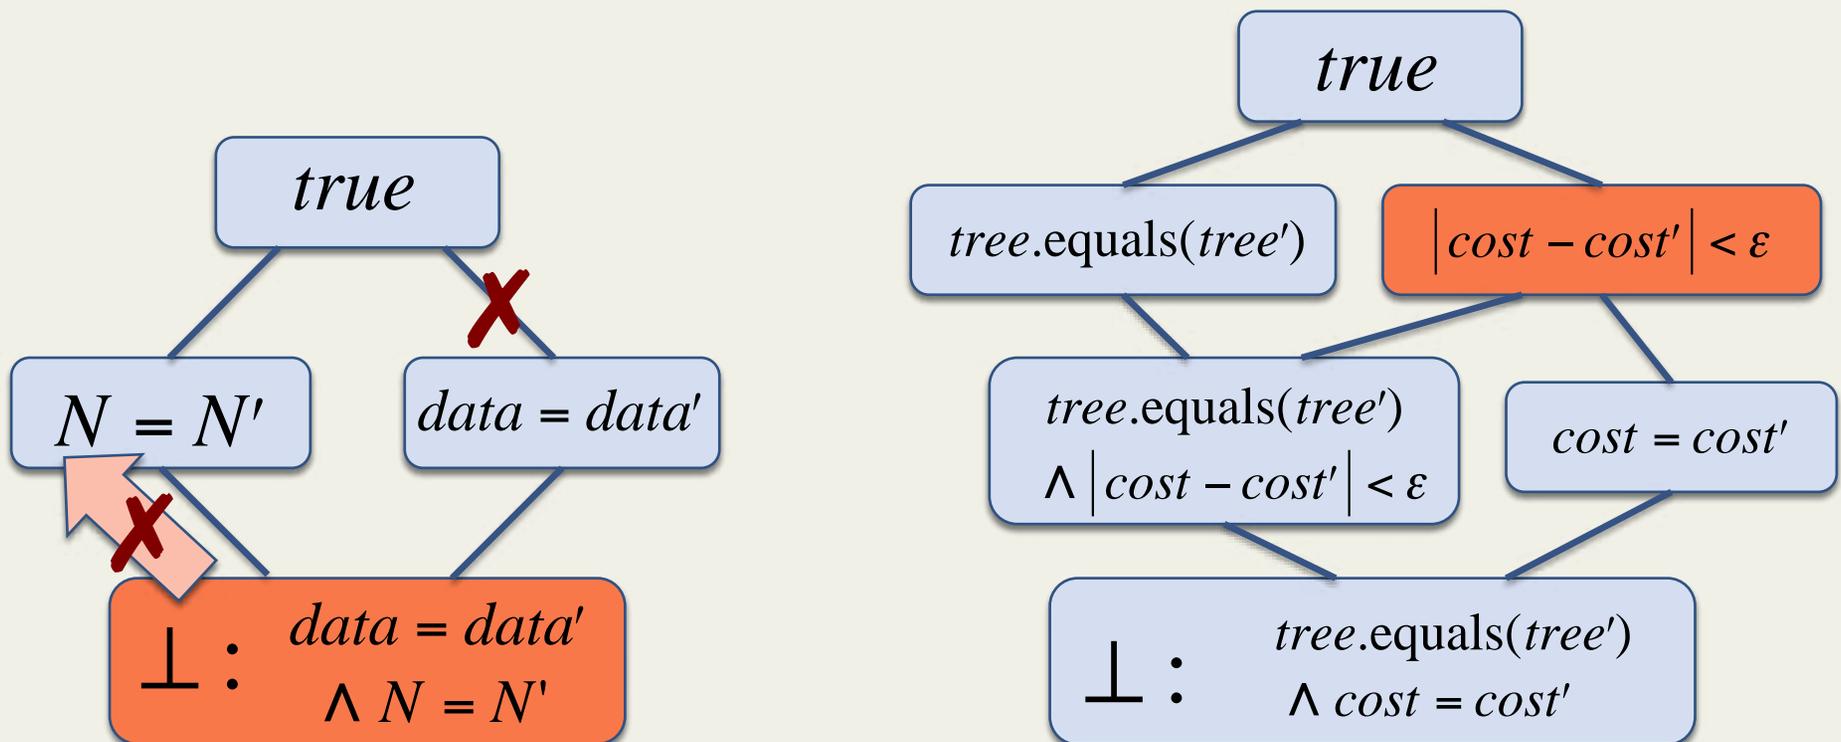
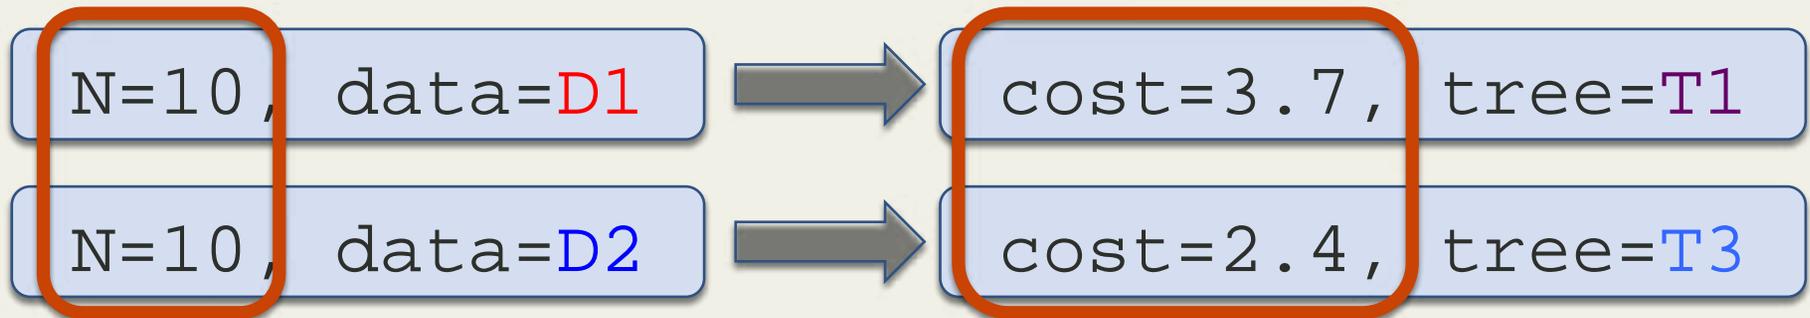
- Repeatedly weaken the precondition
 - As long as it still ensures the postcondition on every pair of observed executions.



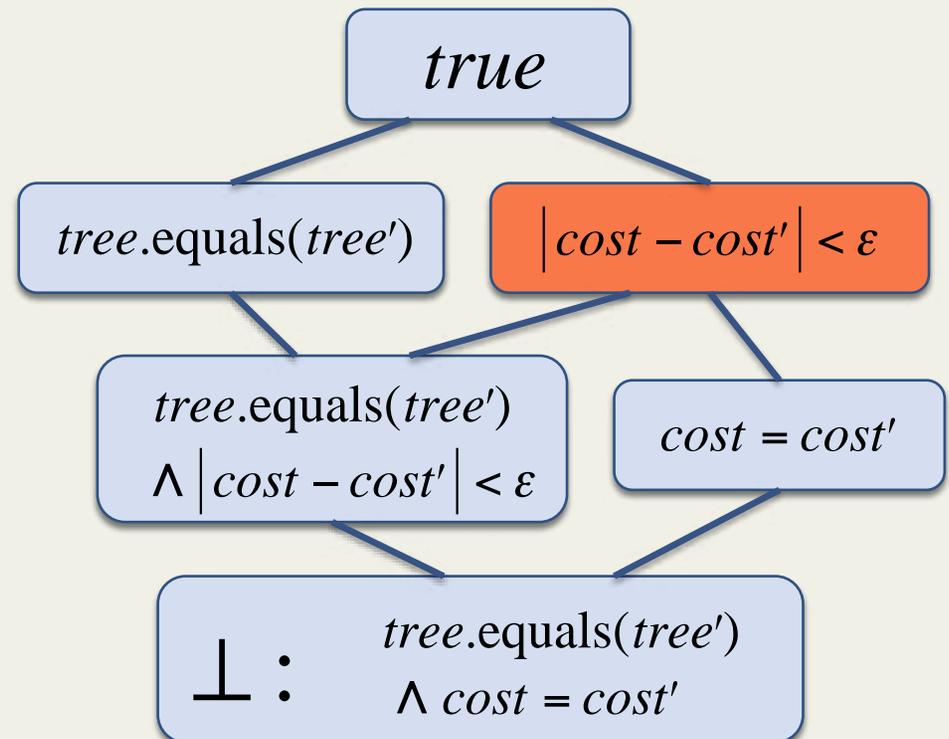
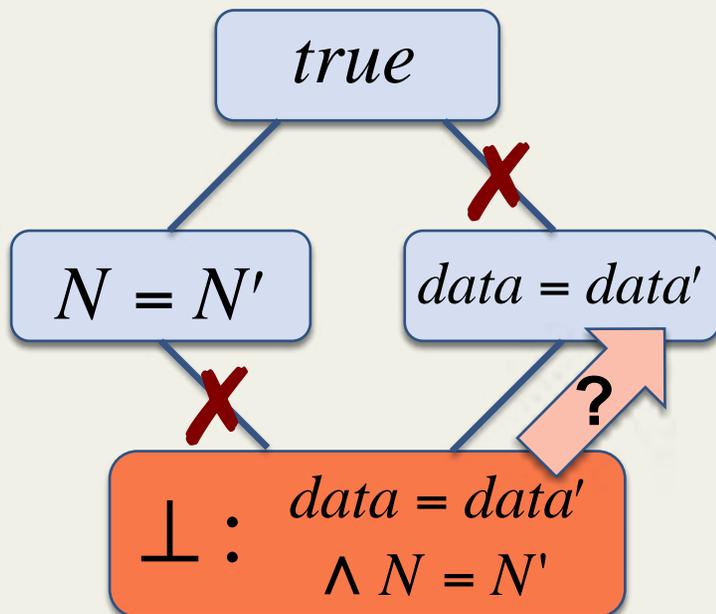
Inferring Weakest Pre I



Inferring Weakest Pre I



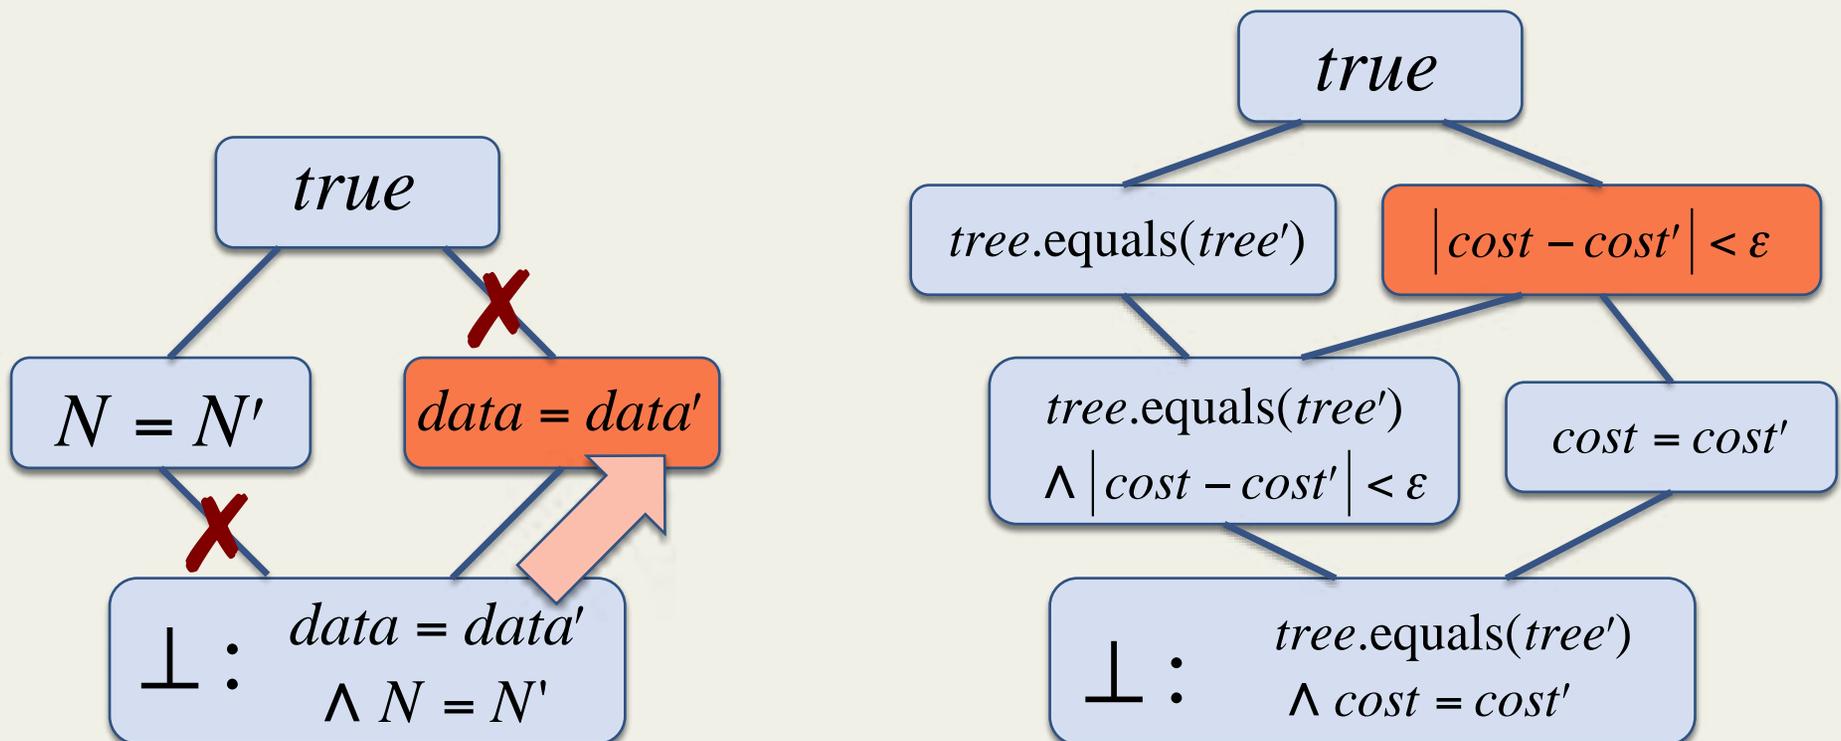
Inferring Weakest Pre II



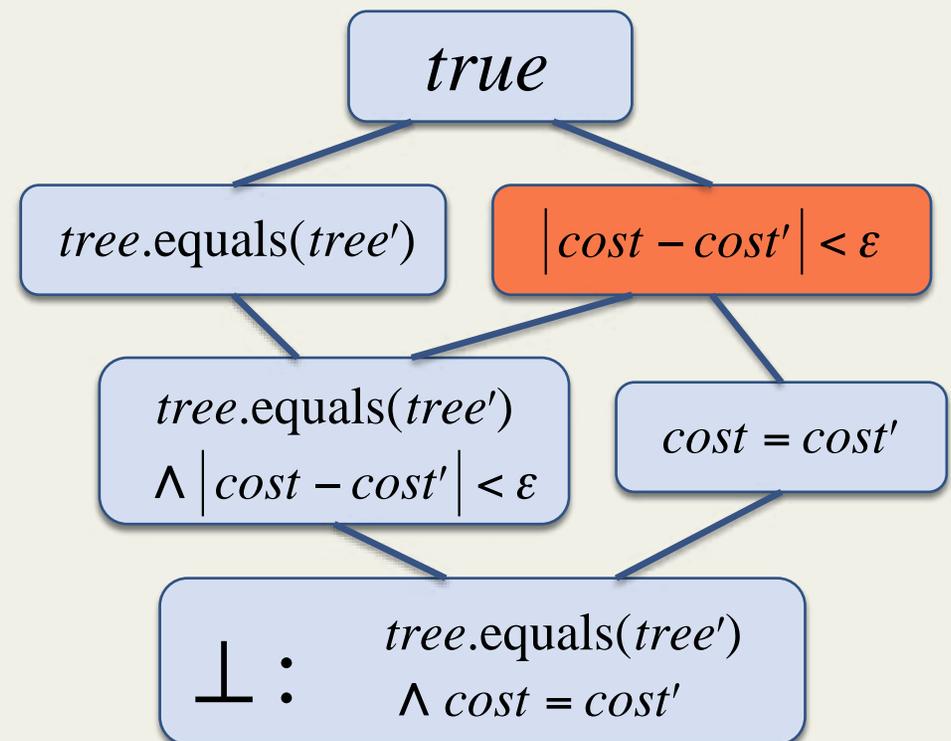
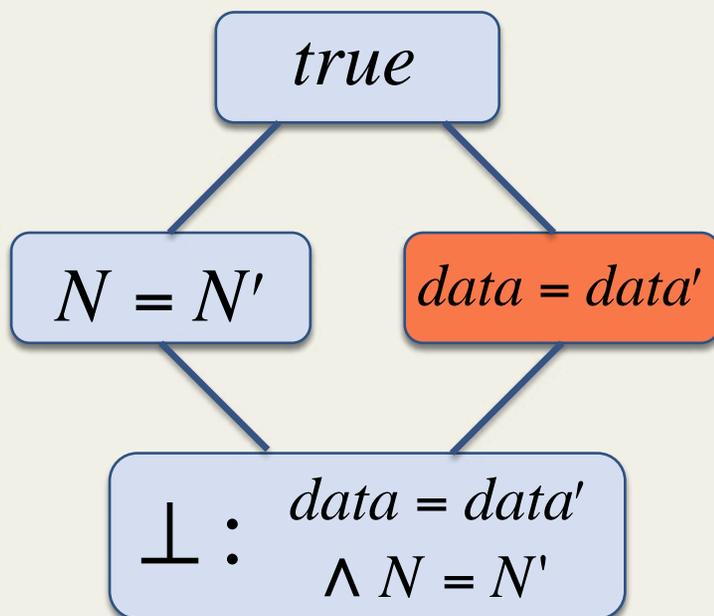
Inferring Weakest Pre II

For all observed pairs of runs:

$$data = data' \xrightarrow{\text{min_phylo_tree}} |cost - cost'| < \epsilon$$



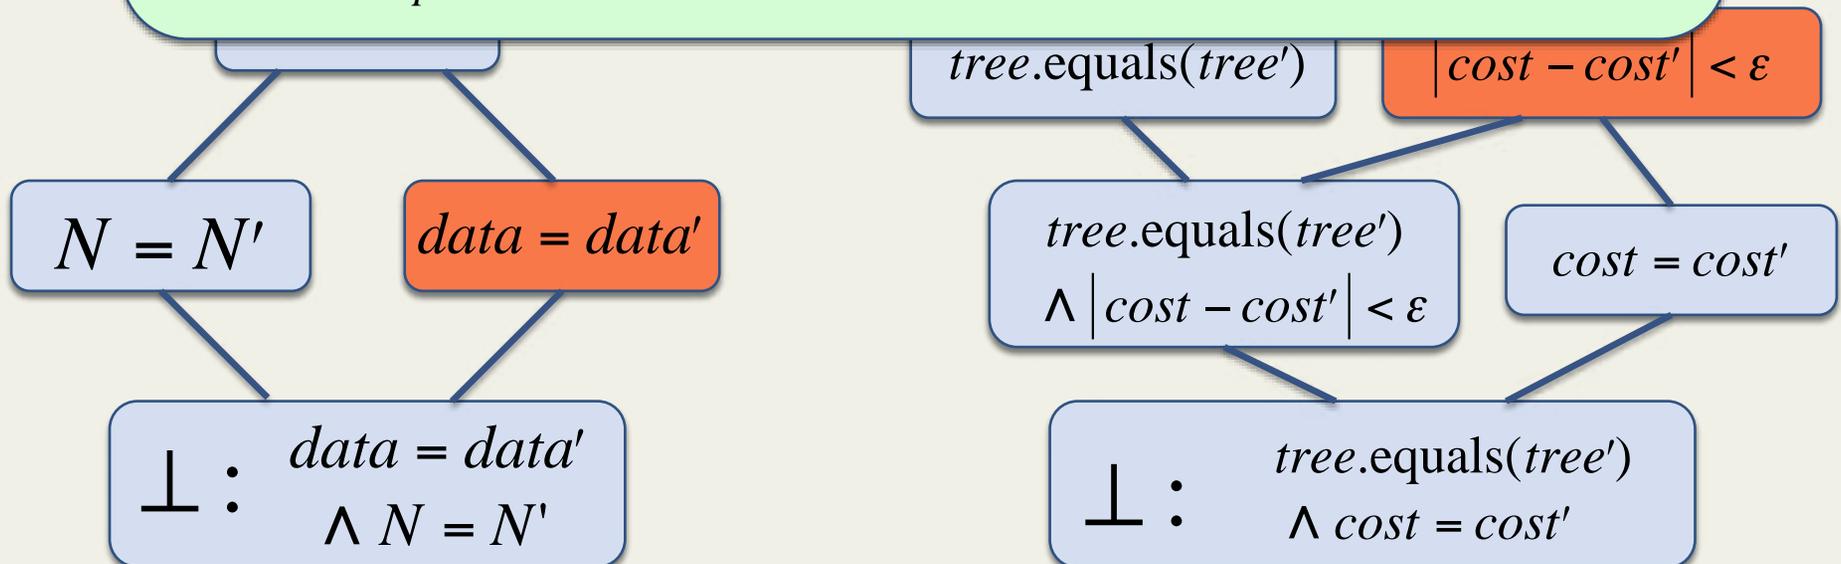
Inferred Weakest Precondition



Inferred Weakest Precondition

Proposition 5. Let the inferred specification be $Pre_R \rightarrow Post_R$. Then, Pre_R is a weakest precondition of $Post_R$ for the observed runs.

I.e., if $\varphi_{pre} \xrightarrow{P,R} Post_R$ and $\varphi_{pre} \Rightarrow Pre_R$, then $\varphi_{pre} = Pre_R$.



Outline

- Motivation and Overview
- Background: Deterministic Specs
- Specification Inference Problem
- Inferring Deterministic Specifications
- **Experimental Evaluation**
- Related Work
- Conclusions

Determinism Inference Experiments

- For previous benchmarks, infer specs for manually-identified deterministic blocks.
 - Benchmarks: 8 from Java Grande Forum (JGF), 4 from Parallel Java (PJ) Library.

Implementation (for Java)

- Record memory graph at open and close of deterministic block.
- Three equality predicates considered:
 - equals(), approximate equality, set equality
 - Compare any chain of fields (up to length 8):
e.g., `this.tree.cost`, `Harness.matrix`
- Heuristics to reduce specification size.
 - By removing “uninteresting” conjuncts.

Implementation (for Java)

- Heuristics needed to shrink specifications:
 - **Remove inputs from postcondition:**
If no run changes v , don't include $v=v'$.

$data=data' \Rightarrow cost=cost' \wedge data=data'$

Implementation (for Java)

- Heuristics needed to shrink specifications:
 - Remove inputs from postcondition.
 - **Remove constants from pre- and post-:**
If v is equal in every run, don't include $v=v'$.

$data=data' \wedge MAX_MEM=MAX_MEM'$
 $\Rightarrow cost=cost' \wedge done=done'$

Implementation (for Java)

- Heuristics needed to shrink specifications:
 - Remove inputs from postcondition.
 - Remove constants from pre- and post-.
 - **Remove redundant conditions:**

params=params'

=> point.equals(point')

^ point.x=point.x'

^ point.y=point.y'

Determinism Inference Experiments

- For previous benchmarks, infer specs for manually-identified deterministic blocks.
 - Benchmarks: 8 from Java Grande Forum (JGF), 4 from Parallel Java (PJ) Library.
- Compare to manual specifications.
 - Are inferred specs correct?
 - Capture intended deterministic behavior?
 - Small enough to be useful?

Experimental Results

- Inferred specification vs. manual one:
 - Equivalent for 7/13 benchmarks.

Manual:

`params=params' → matrix=matrix'`

Inferred:

`params=params' → matrix=matrix'
' ∧ img.equals(img')`

Experimental Results

- Inferred specification vs. manual one:
 - Equivalent for 7/13 benchmarks.
 - Inferred correct while manual wrong for 2/13!

Experimental Results

- Inferred specification vs. manual one:
 - Equivalent for 7/13 benchmarks.
 - Inferred correct while manual wrong for 2/13!
 - 1/13 is correct but stronger than desired.

Manual:

$$\text{params}=\text{params}' \longrightarrow |ek[0]-ek[0]'| < \varepsilon$$

Inferred:

$$\text{params}=\text{params}' \wedge \textit{nthreads}=\textit{nthreads}' \longrightarrow |ek-ek'| < \varepsilon$$

Experimental Results: JGF

Bench	LoC	Size of Precondition		Size of Postcondition	
		Manual	Inferred	Manual	Inferred
series	800	1	3	1	1
crypt	1.1k	1	3	2	2
moldyn	1.3k	2	14	3	7
raytracer	1.9k	2	3	1	1
monte	3.6k	1	2	1	1

Experimental Results: PJ and tsp

Bench	LoC	Size of Precondition		Size of Postcondition	
		Manual	Inferred	Manual	Inferred
pi3	150	2	3	1	1
keysearch3	200	3	5	1	3
mandelbrot	250	7	11	1	5
phylogeny	4.4k	3	5	2	11
tsp*	700	1	3	1	2

Experimental Results

- Limitations:
 - For 3/13 benchmarks, inferred spec is incorrect because of insufficient test suite.

Manual:

```
graph=graph' → tour.cost=tour.cost'
```

Inferred:

```
graph=graph' → tour.equals(tour')
```

Outline

- Motivation and Overview
- Background: Deterministic Specs
- Specification Inference Problem
- Inferring Deterministic Specs
- Experimental Evaluation
- **Related Work**
- Conclusions

Related Work: Determinism

- Deterministic languages and runtimes.
 - Deterministic Parallel Java (UIUC)
 - Kendo (Olszewski, et al, ASPLOS 09)
 - DMP (Devietti, et al, ASPLOS 09)
- Determinism Checking.
 - SingleTrack (Sadowski, et al, ESOP 09)
 - *Race detection* can be thought of as determinism checking.

Outline

- Motivation and Overview
- Background: Deterministic Specs
- Specification Inference Problem
- Inferring Deterministic Specs
- Experimental Evaluation
- Related Work
- **Conclusions**

Conclusions

- Deterministic specifications – lightweight spec of parallel correctness.
 - Much simpler structure than functional correctness specifications.
- Can infer high-quality deterministic specs.
 - From small number of observed runs.
 - Mostly recovered previous manual specs.
 - Found two errors in previous manual specs.

Any Questions?