

Formal Abstractions for Packet Scheduling

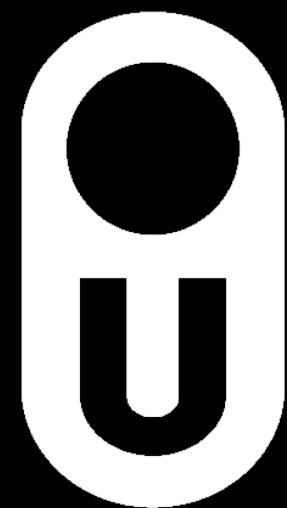
A. Mohan, Y. Liu, N. Foster, T. Kappé, D. Kozen



Formal Abstractions for Packet Scheduling

A. Mohan, Y. Liu, N. Foster, T. Kappé, D. Kozen

Has let me use
his slides!



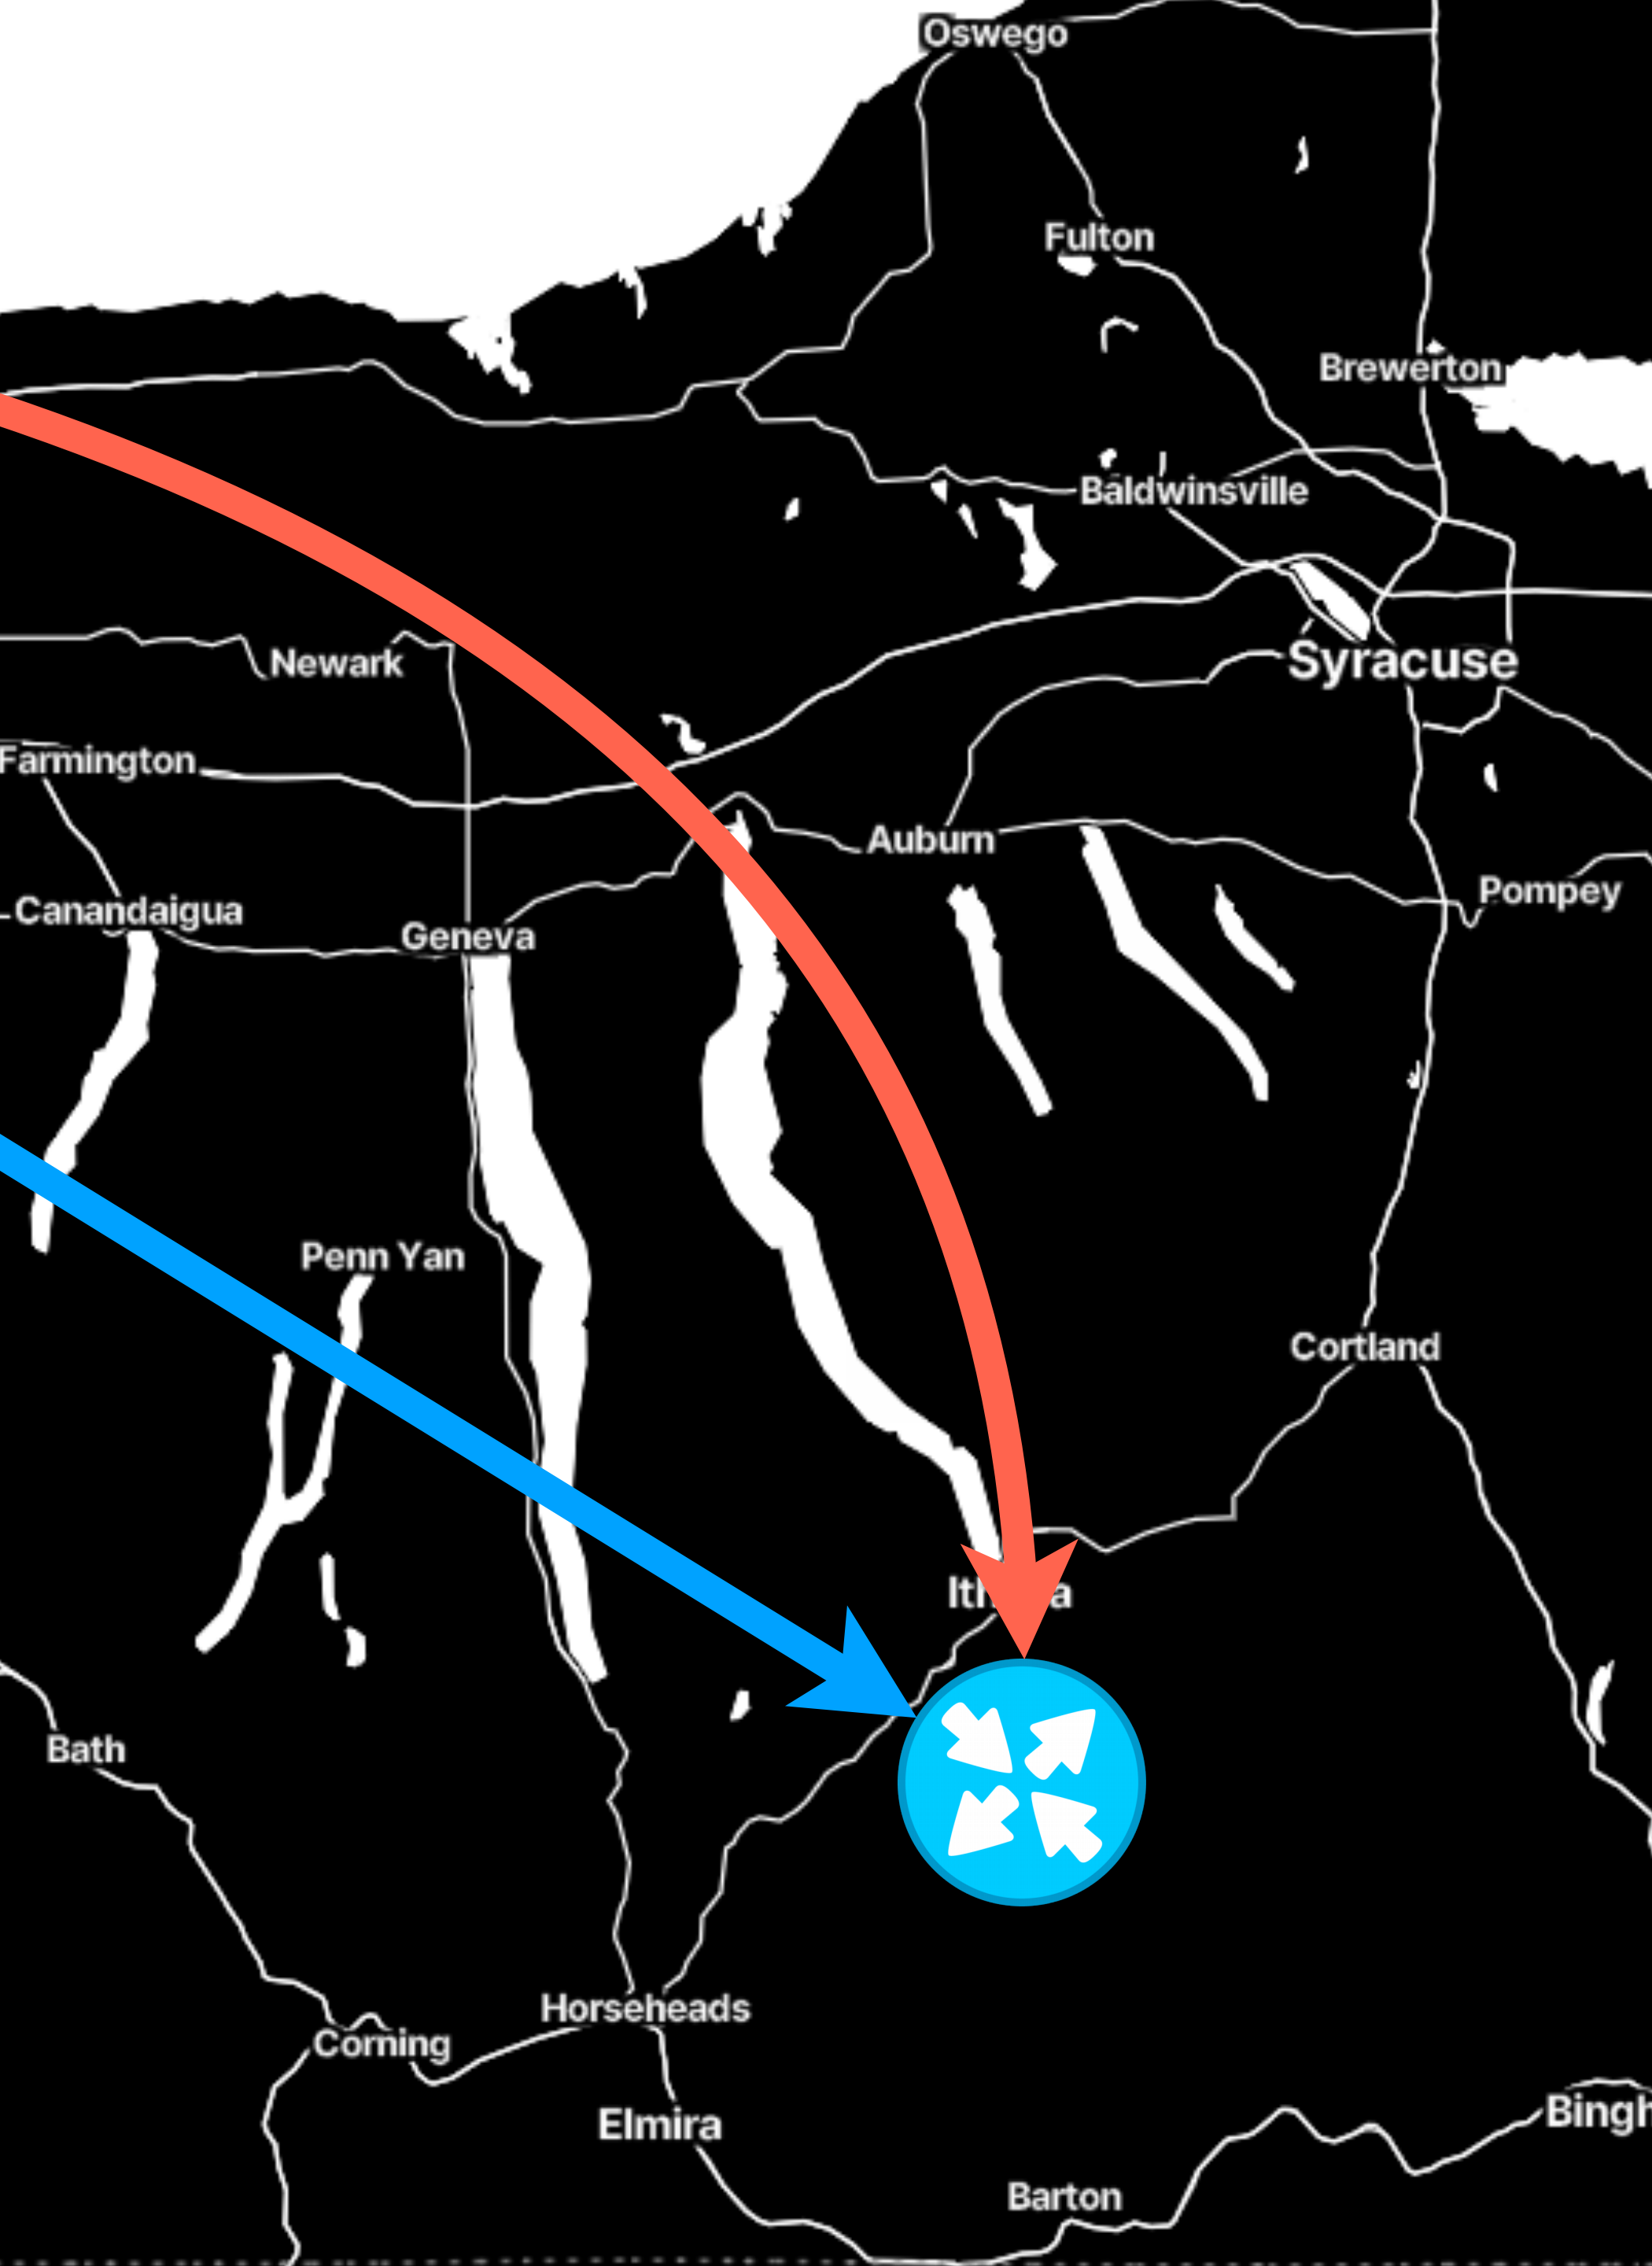


Software-defined networking
made networks programmable.



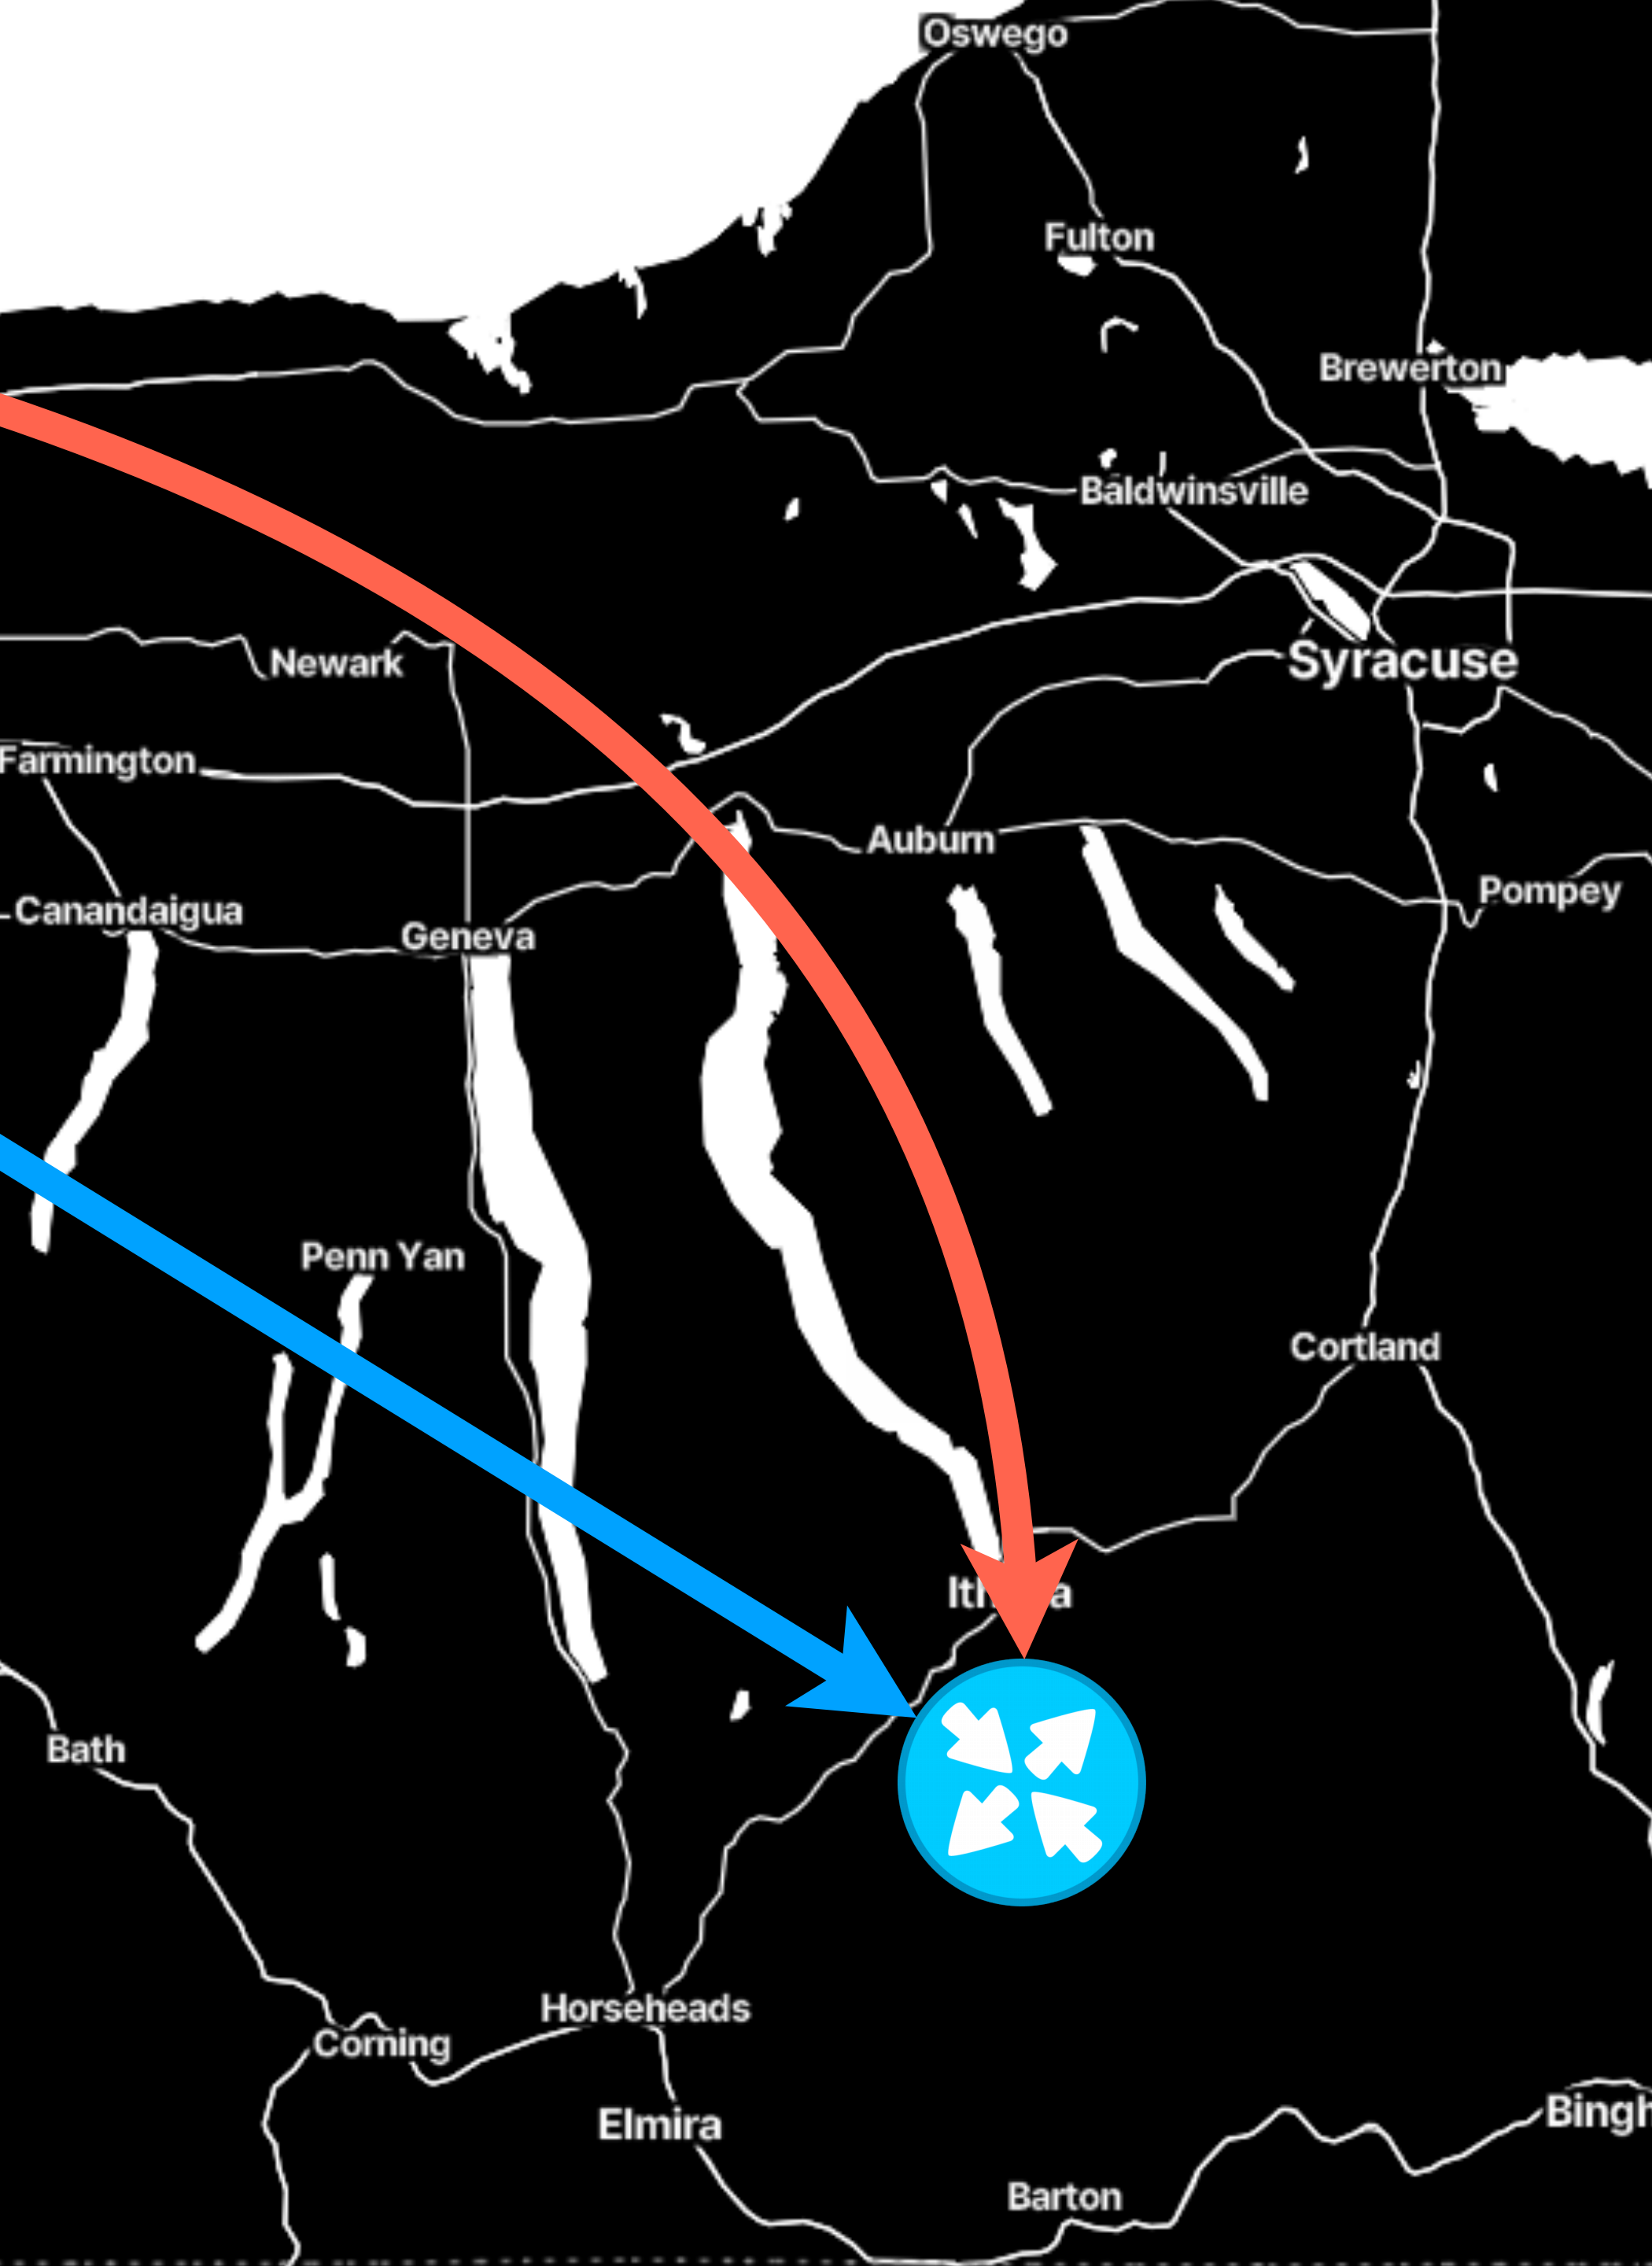
Software-defined networking
made networks programmable.

We want control over
packet scheduling.



Software-defined networking
made networks programmable.

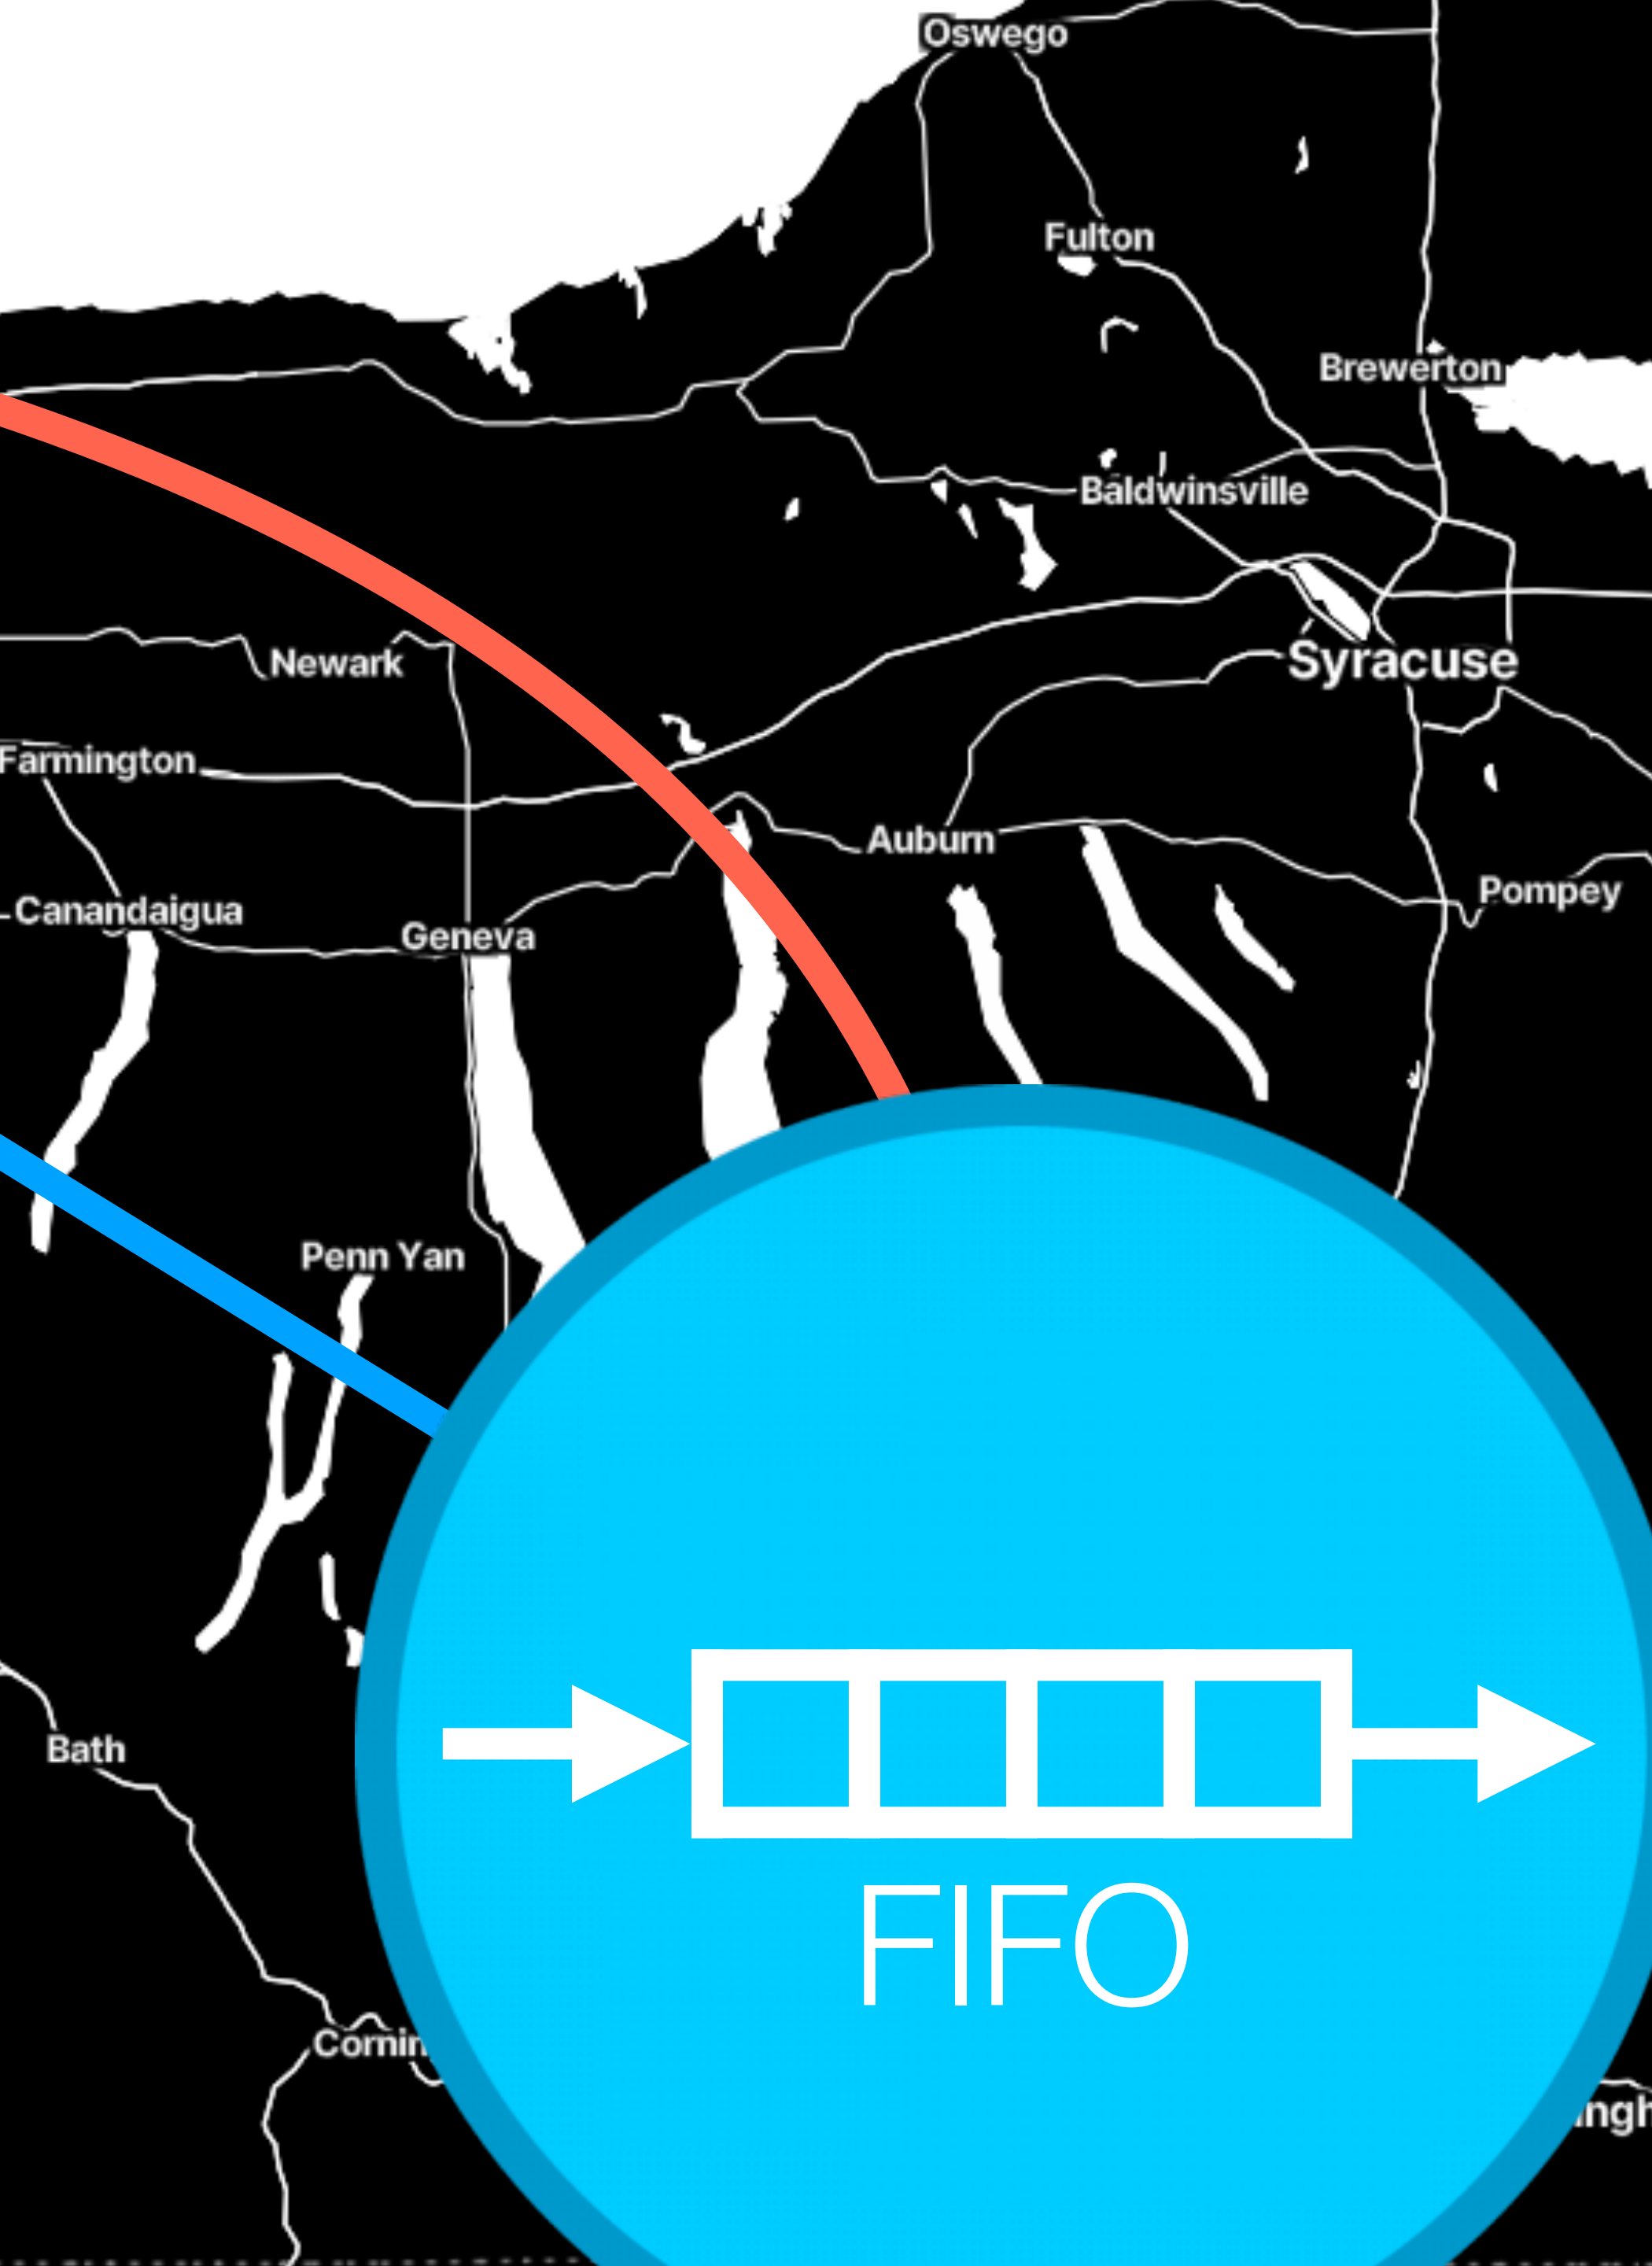
We want control over
packet scheduling.



Software-defined networking
made networks programmable.

We want control over
packet scheduling.

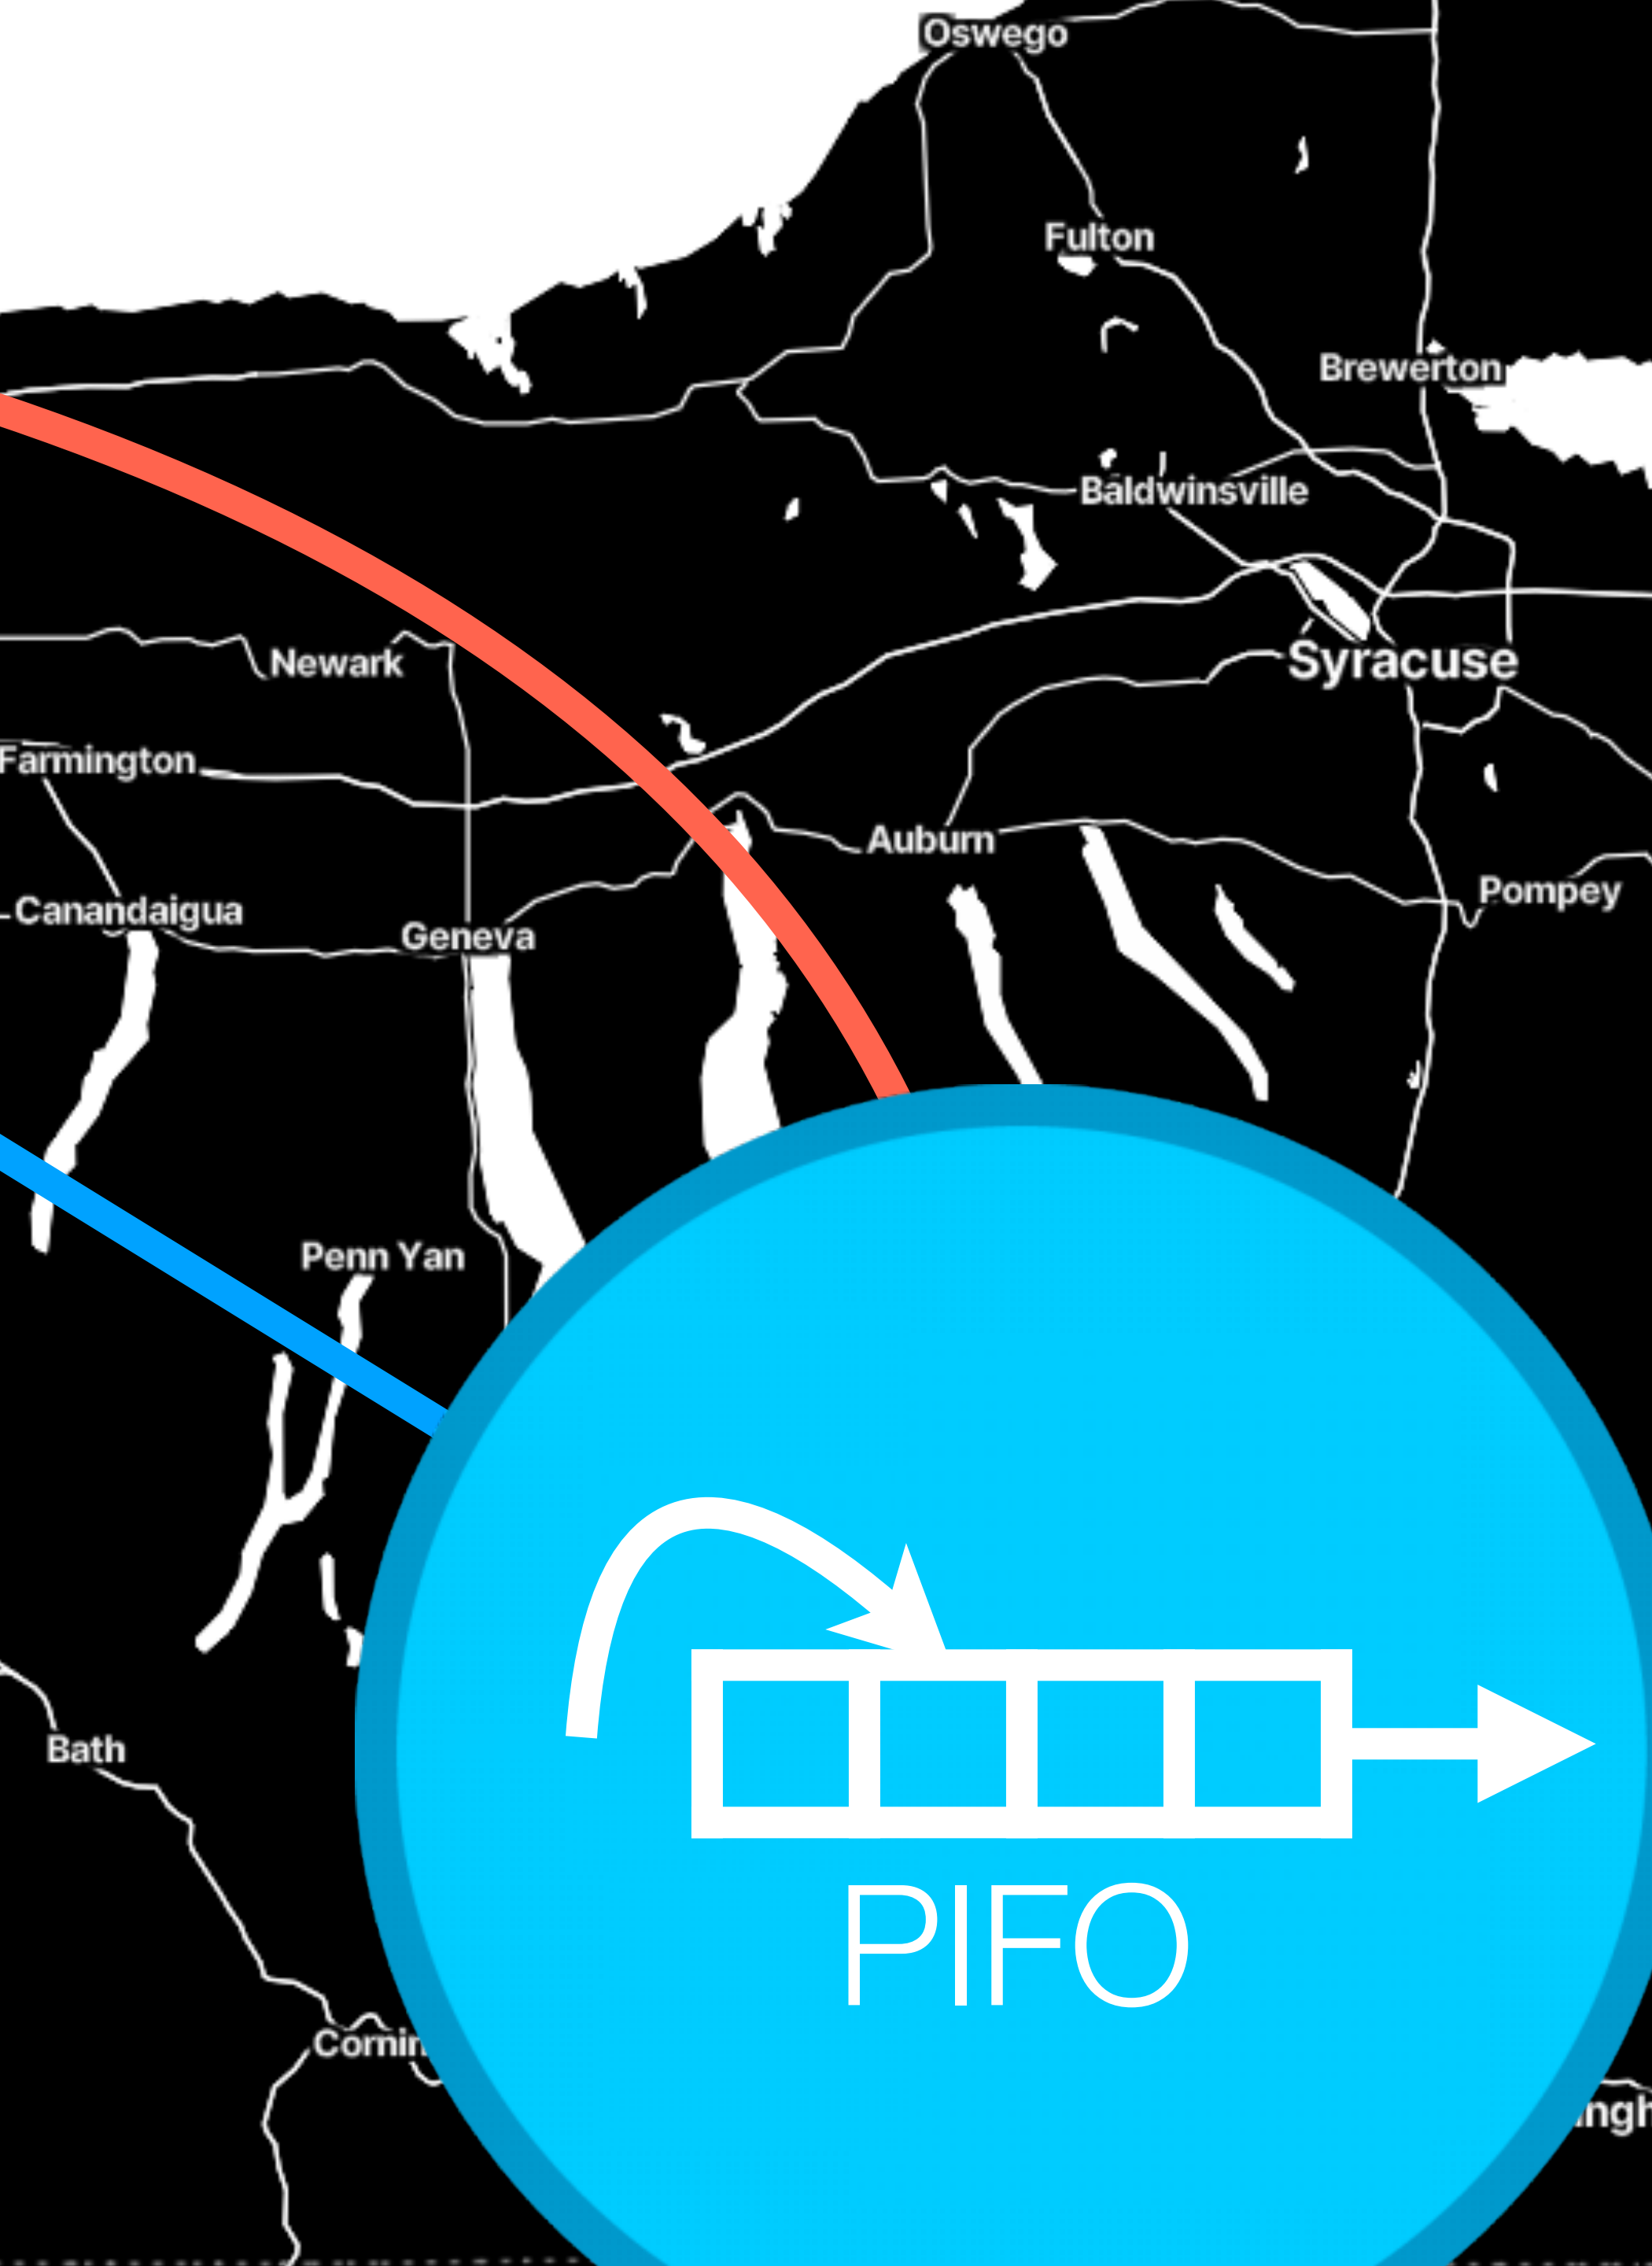
Basic tools work fine...



Software-defined networking
made networks programmable.

We want control over
packet scheduling.

Basic tools work fine...

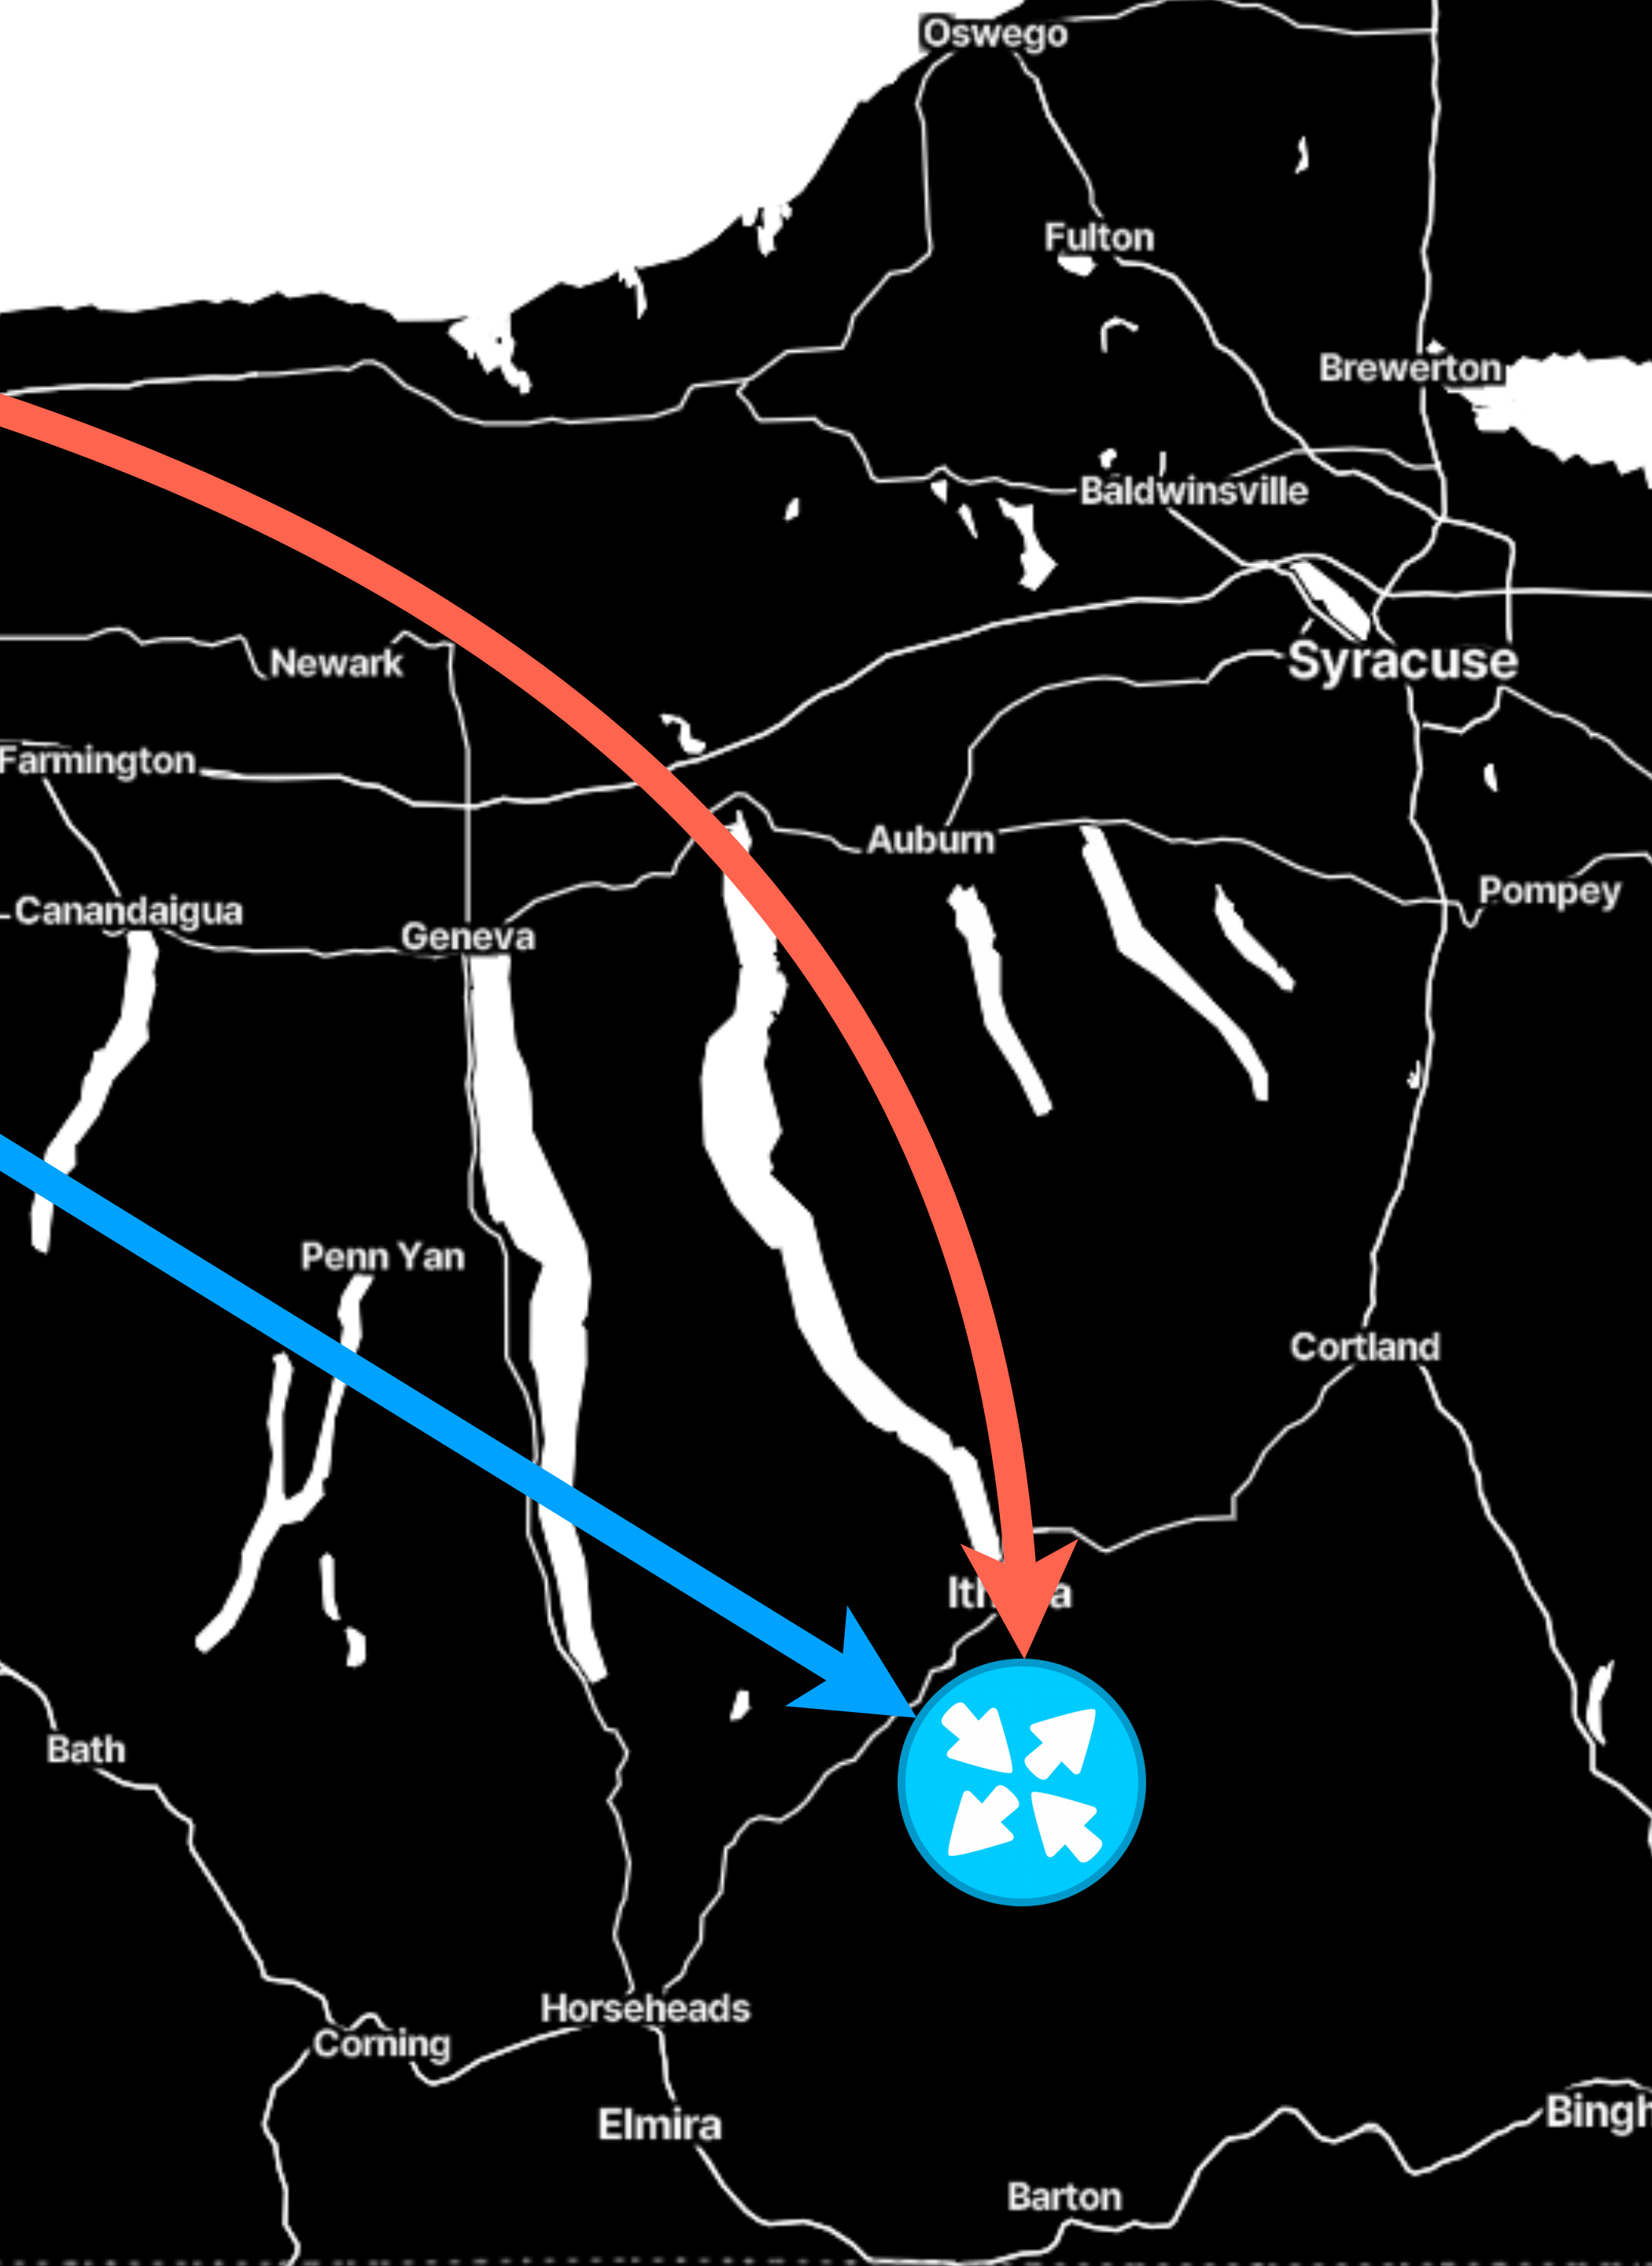


Software-defined networking
made networks programmable.

We want control over
packet scheduling.

Basic tools work fine...

...but modern scheduling
requires more.





...but modern scheduling
requires more.

R traffic goes to either
Pittsburgh or Toronto.



...but modern scheduling
requires more.

R traffic goes to either
Pittsburgh or **Toronto**.

Goal:

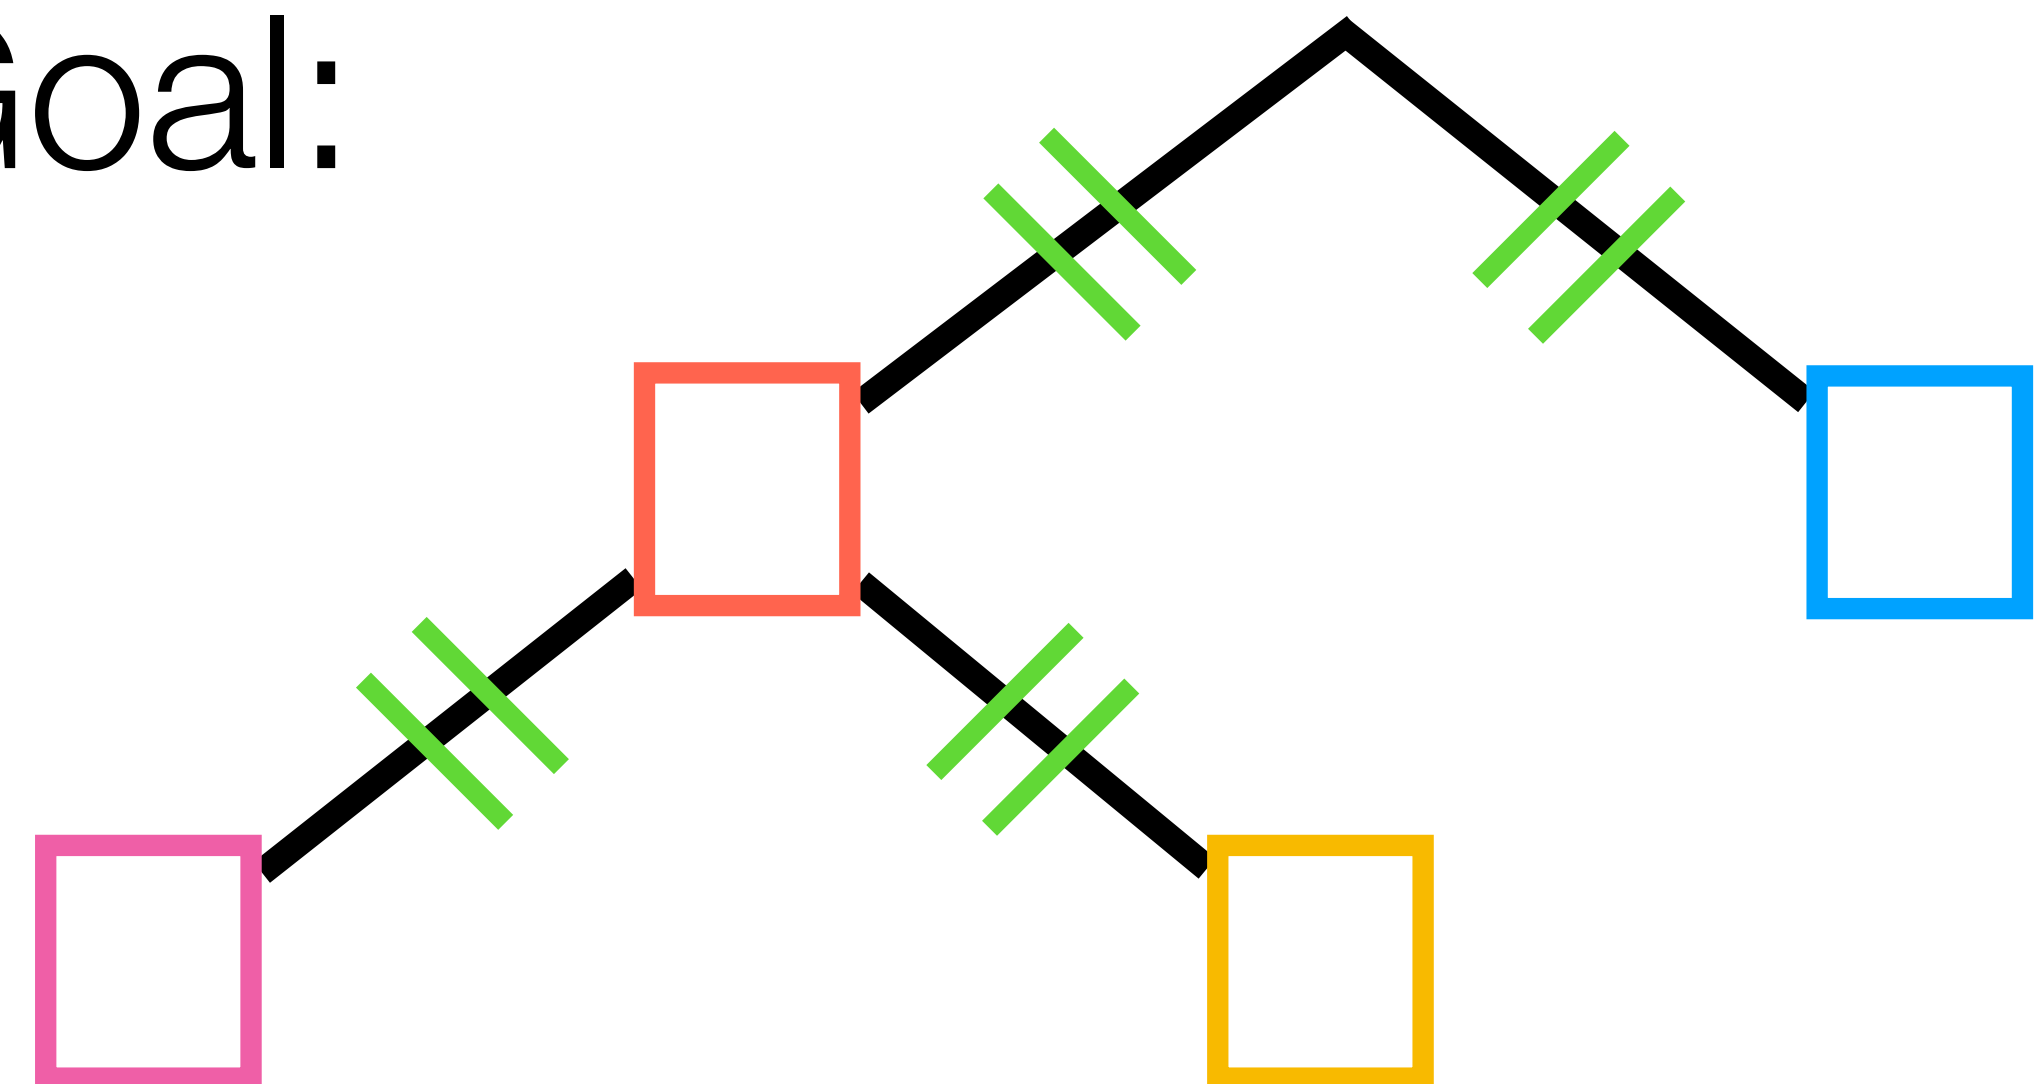
Interleave **R** and **B**;
interleave **P** and **T**.



...but modern scheduling
requires more.

R traffic goes to either
Pittsburgh or Toronto.

Goal:

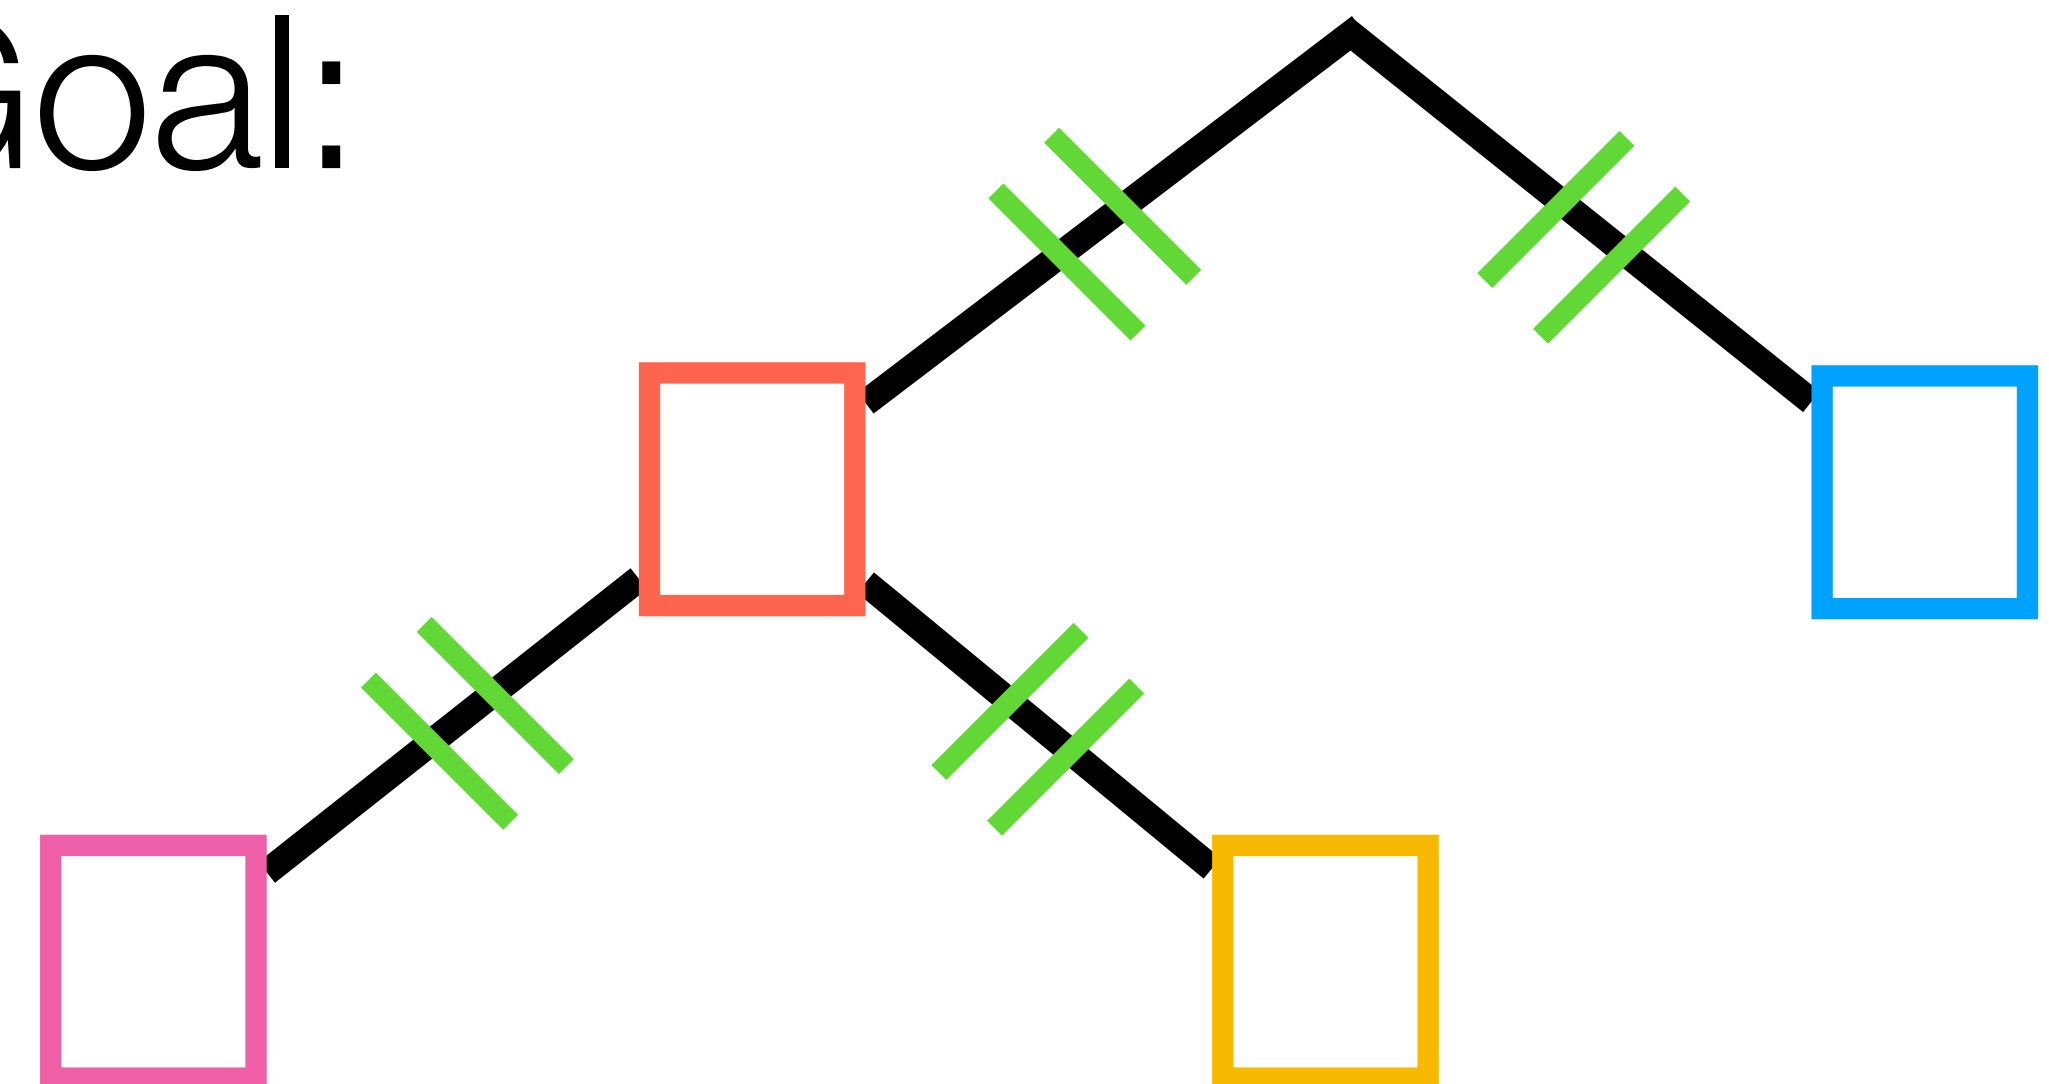




...but modern scheduling
requires more.

R traffic goes to either
Pittsburgh or Toronto.

Goal:





New plan!

PIFO Tree



New plan!

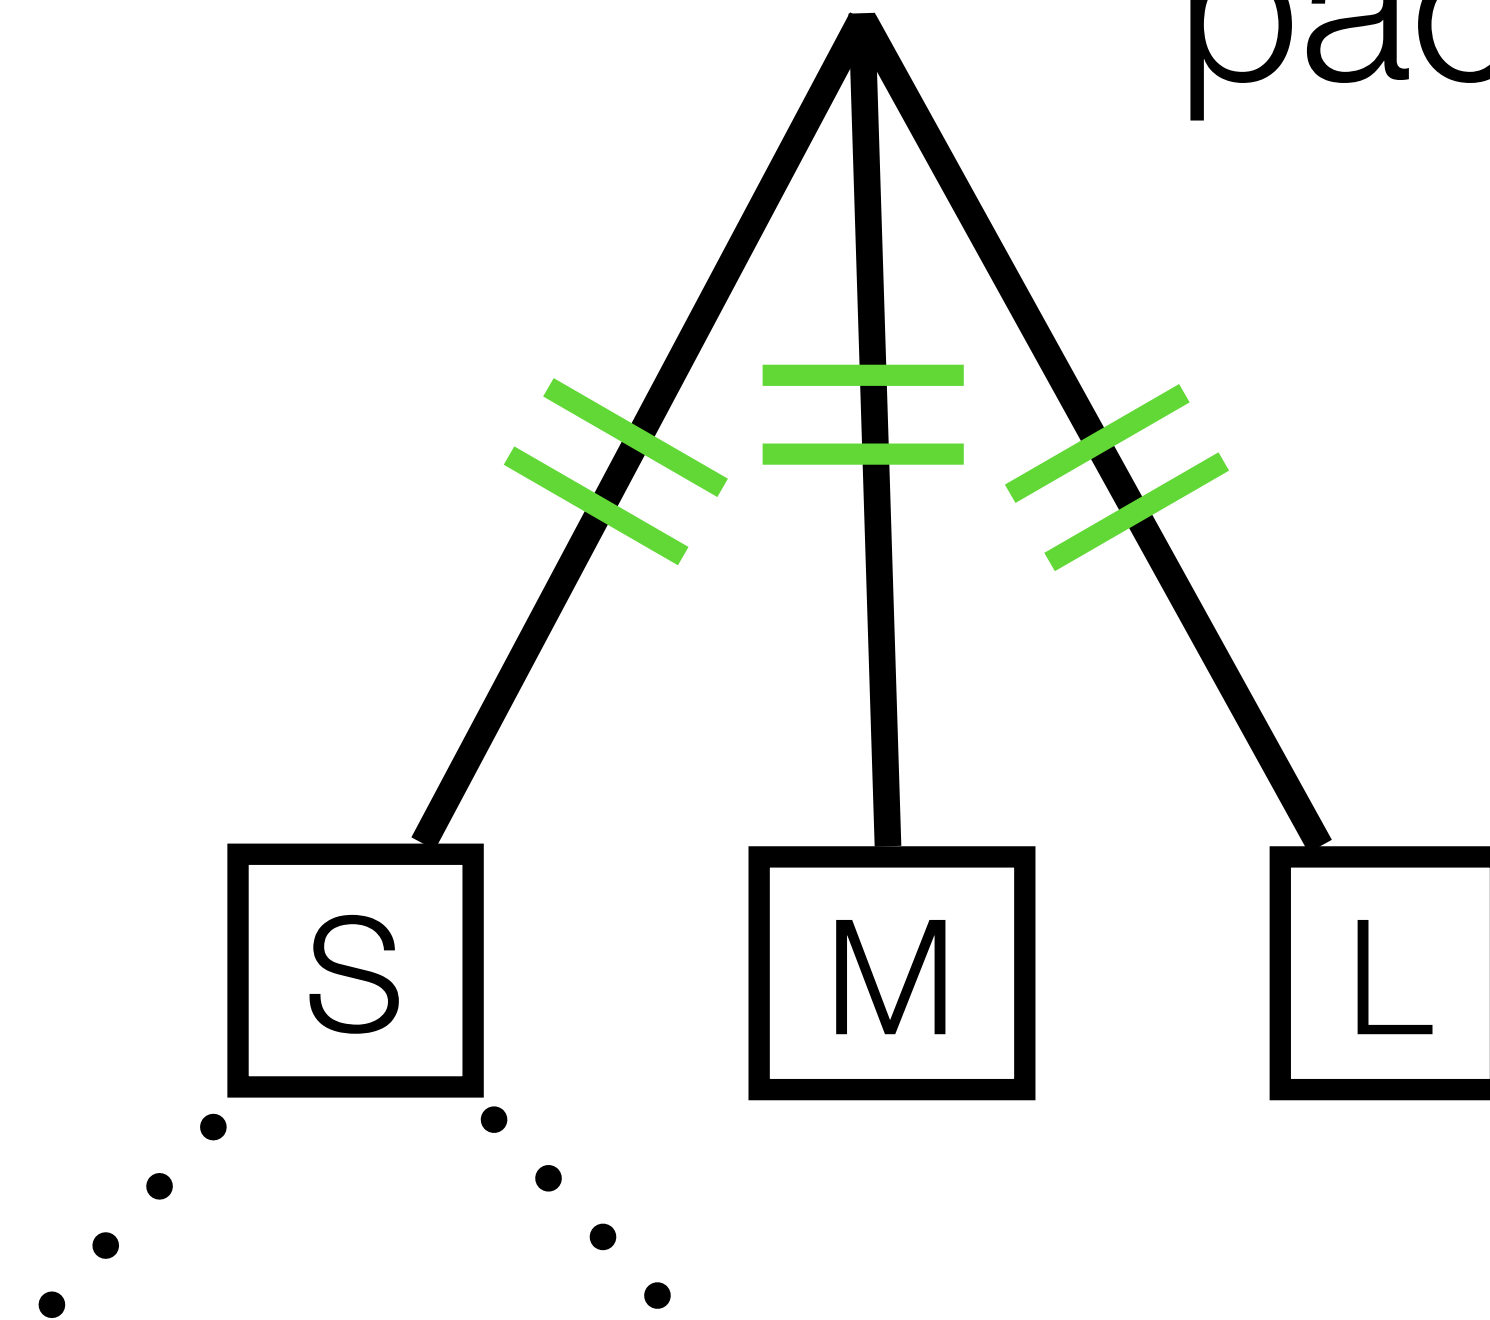
Interleave
small, medium, and large
packets.

PIFO Tree



New plan!

Interleave
small, medium, and large
packets.



No general way to deploy our gadget.

No general way to deploy our gadget.



A human needs a
range of trees.

No general way to deploy our gadget.



A human needs a
range of trees.

The hardware wants
to support *one* tree.

No general way to deploy our gadget.



A human needs a
range of trees.



The hardware wants
to support *one* tree.

No general way to deploy our gadget.



A human needs a
range of trees.

The hardware wants
to support *one* tree.

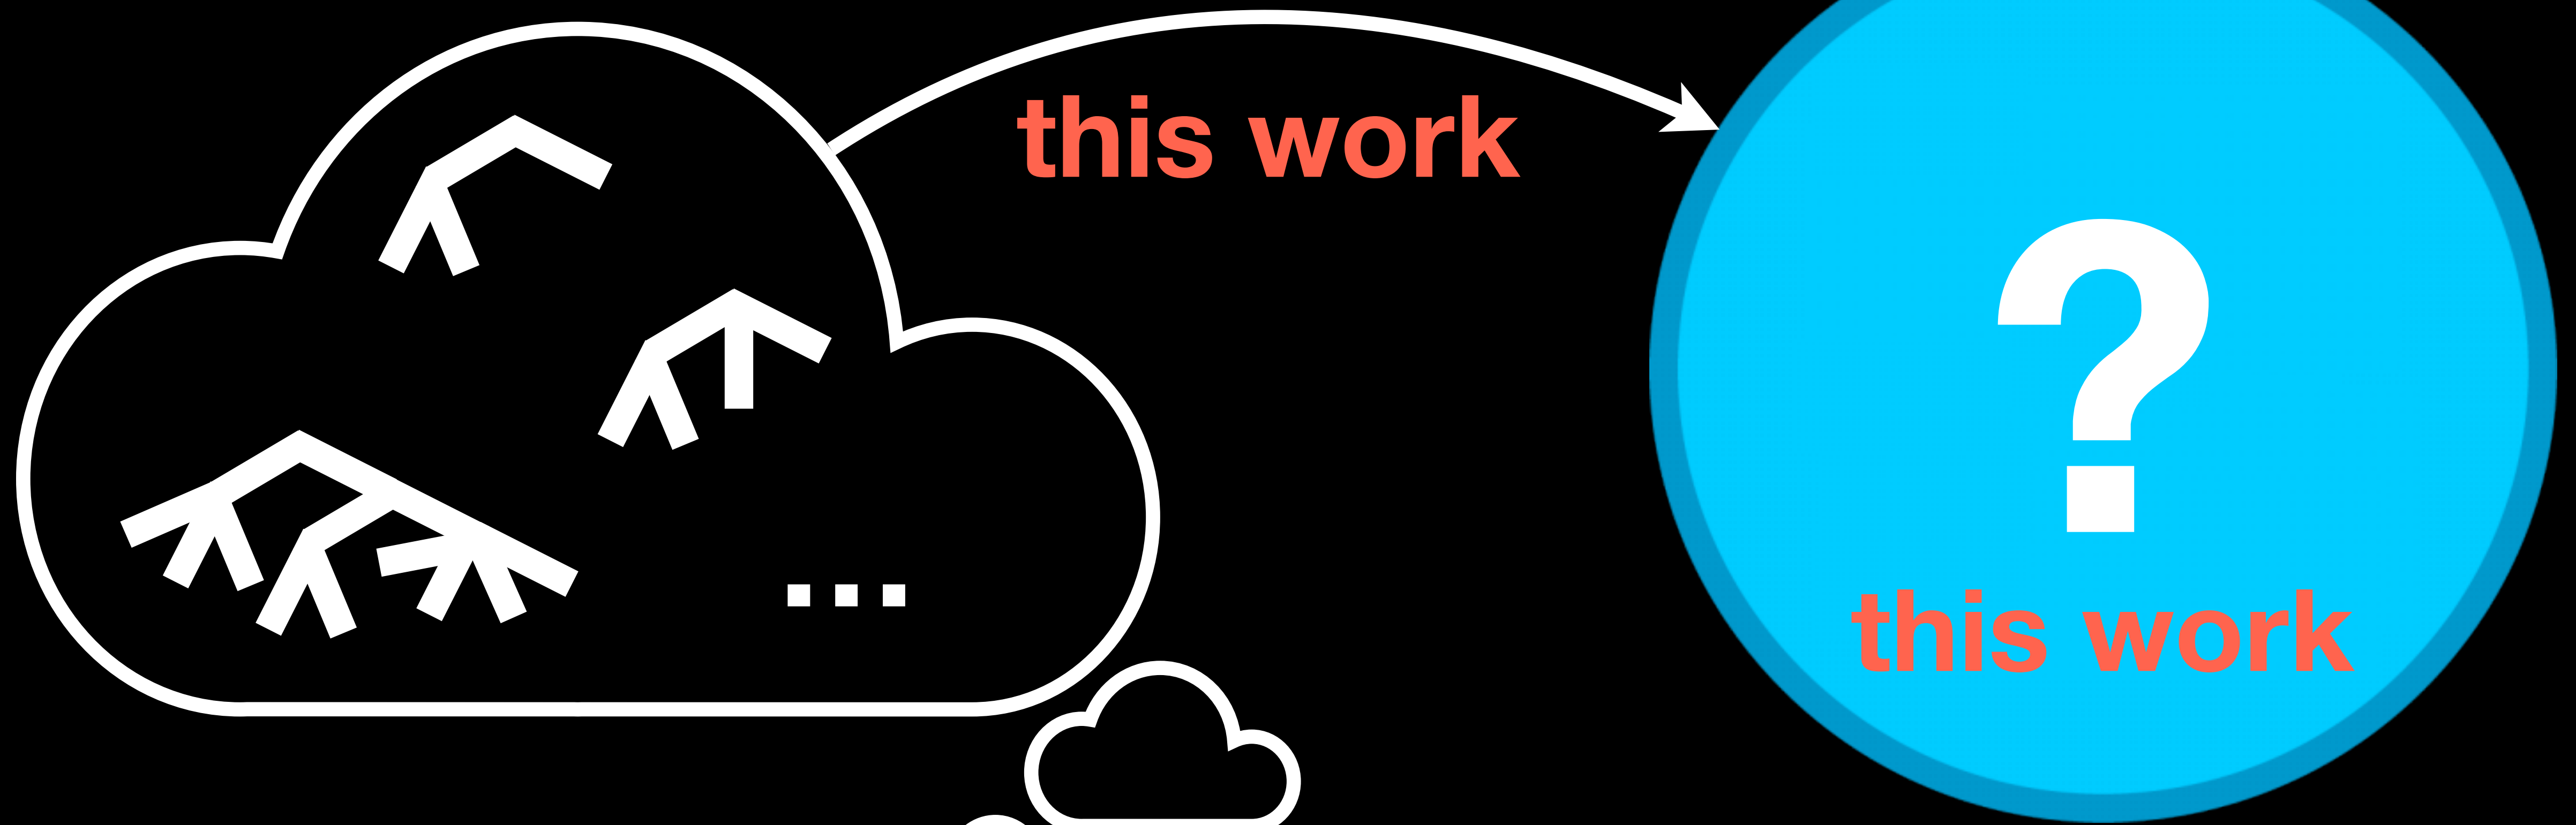
No general way to deploy our gadget.



A human needs a
range of trees.

The hardware wants
to support *one* tree.

No general way to deploy our gadget.



A human needs a
range of trees.

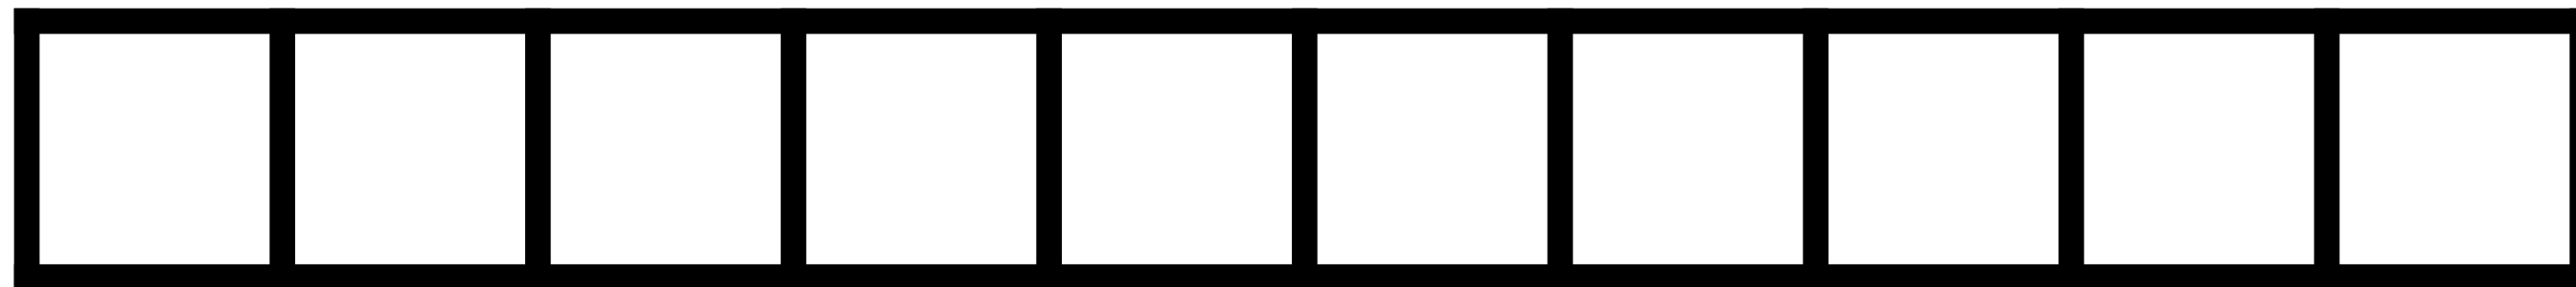
The hardware wants
to support *one* tree.

Aside: PIFO Trees

Sivaraman et al. at SIGCOMM '16

Review: FIFO

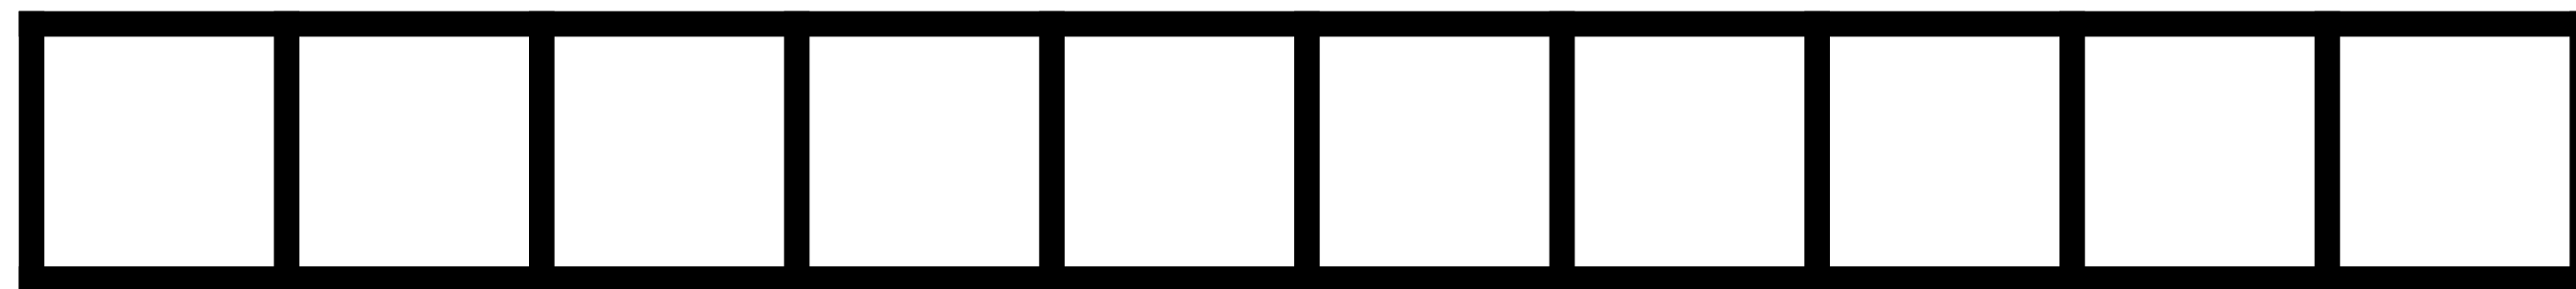
Just an ordered collection.



Review: FIFO

Just an ordered collection.

Two ways of interacting with the collection:

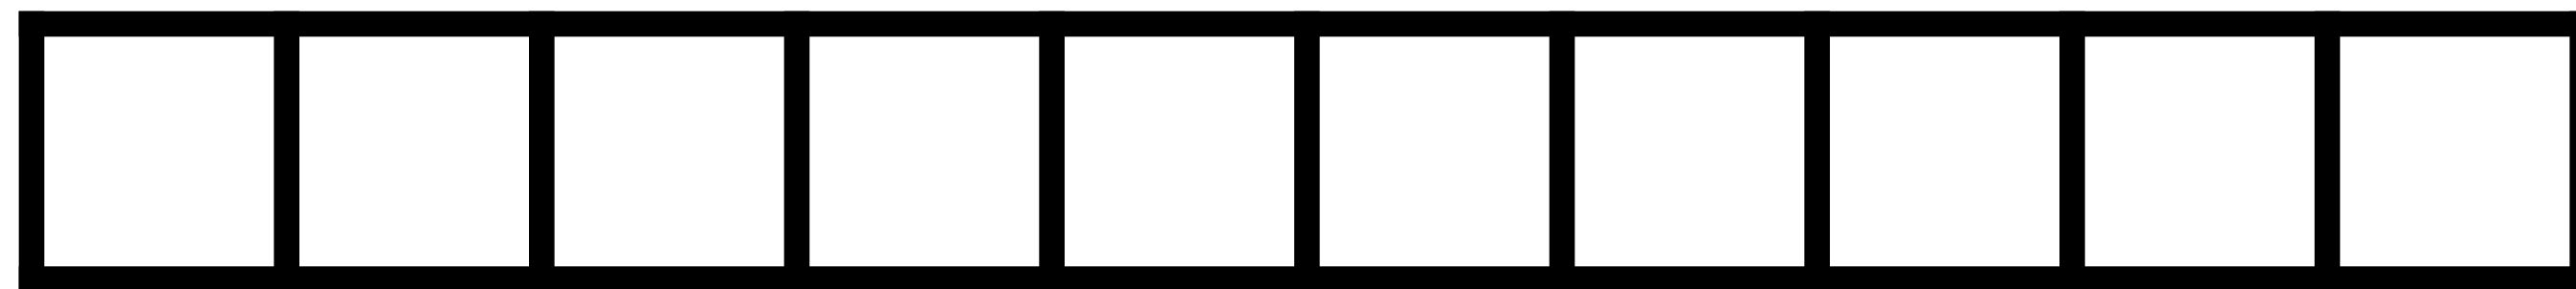


Review: FIFO

Just an ordered collection.

Two ways of interacting with the collection:

push

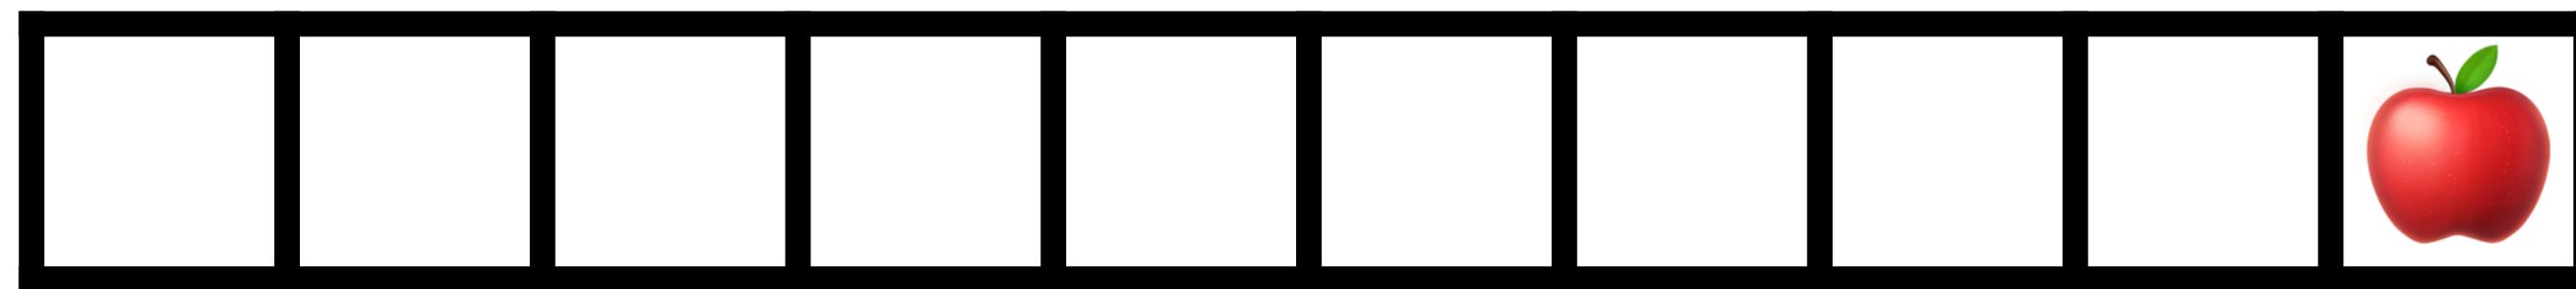


Review: FIFO

Just an ordered collection.

Two ways of interacting with the collection:

push

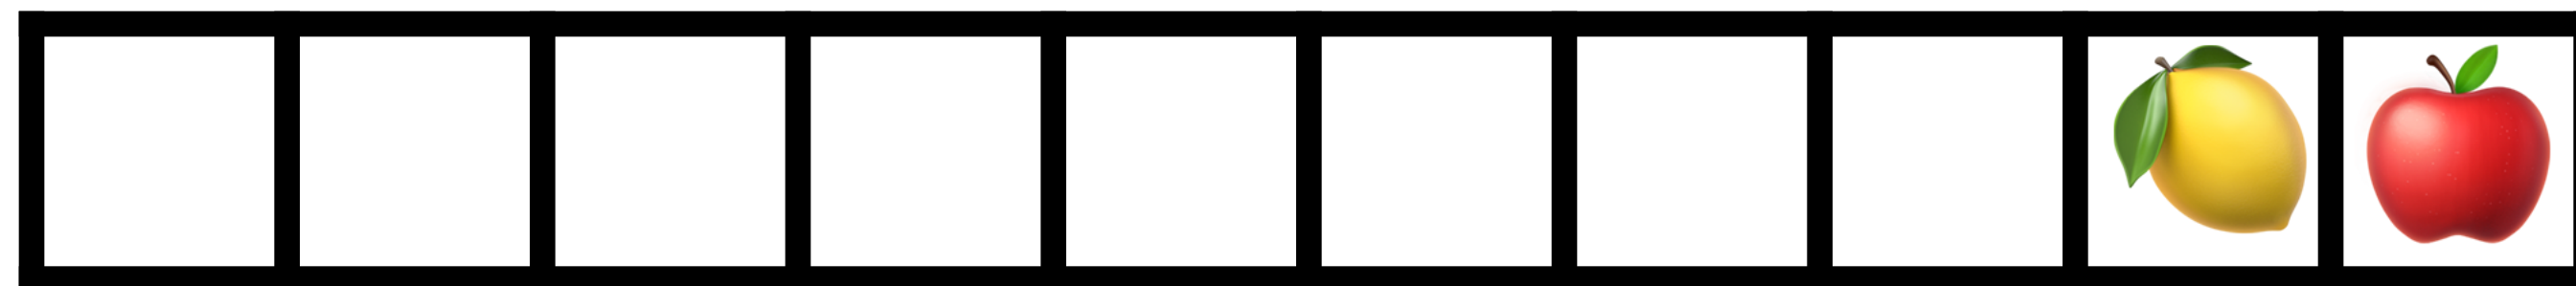


Review: FIFO

Just an ordered collection.

Two ways of interacting with the collection:

push



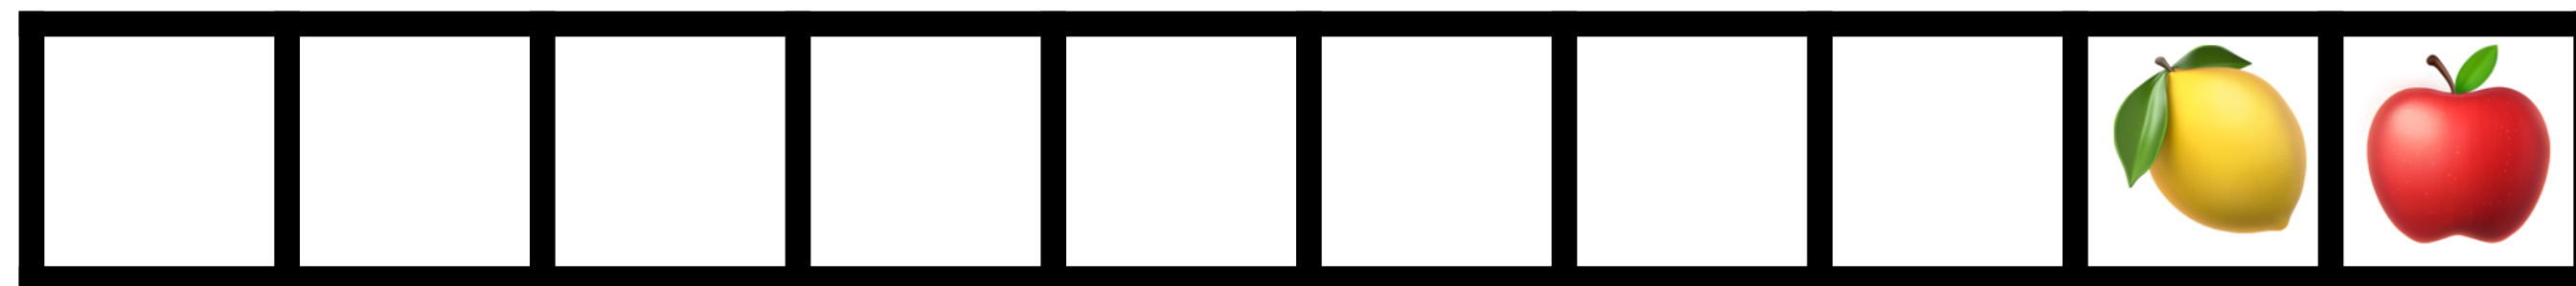
Review: FIFO

Just an ordered collection.

Two ways of interacting with the collection:

push

pop



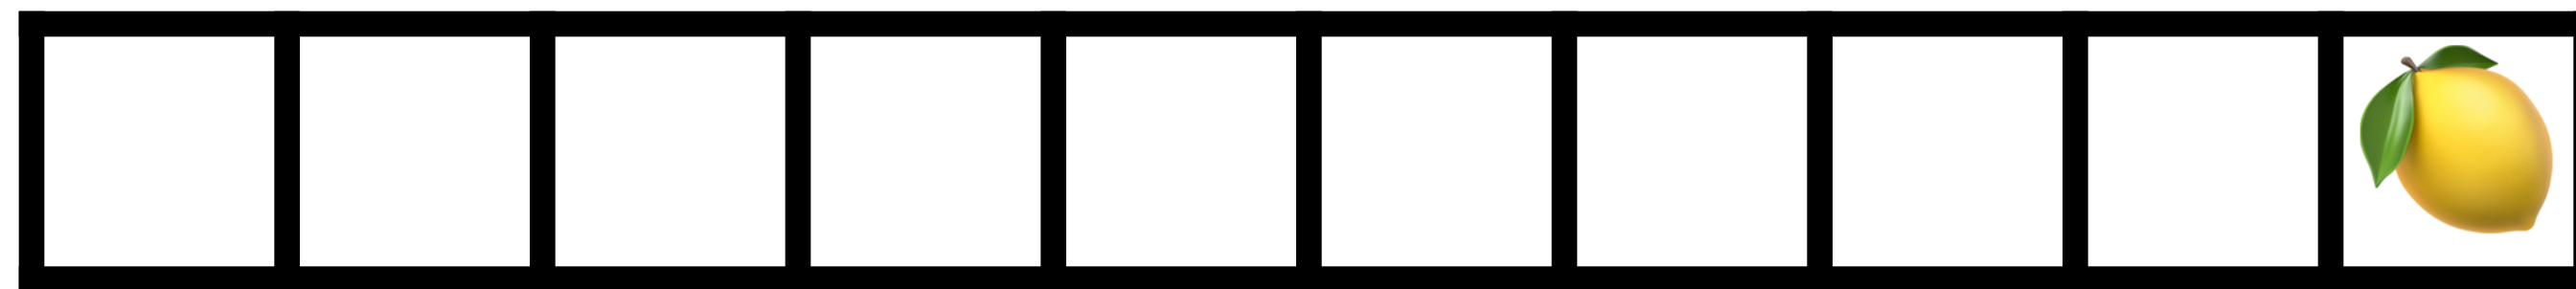
Review: FIFO

Just an ordered collection.

Two ways of interacting with the collection:

push

pop



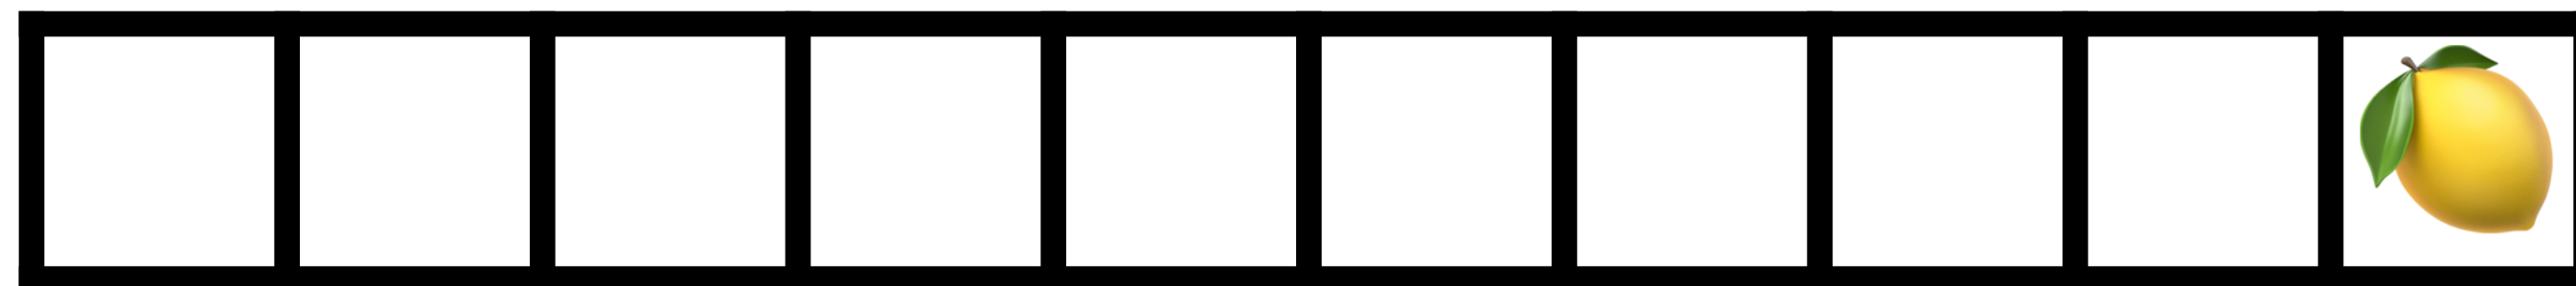
Review: FIFO

Just an ordered collection.

Two ways of interacting with the collection:

push

pop



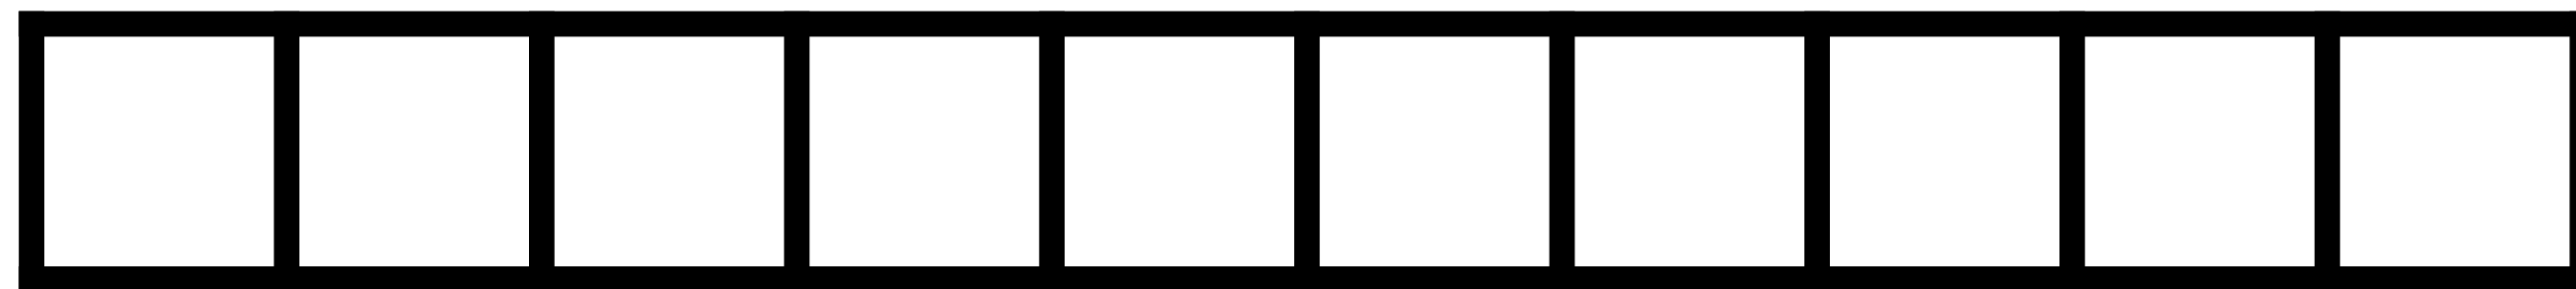
Review: FIFO

Just an ordered collection.

Two ways of interacting with the collection:

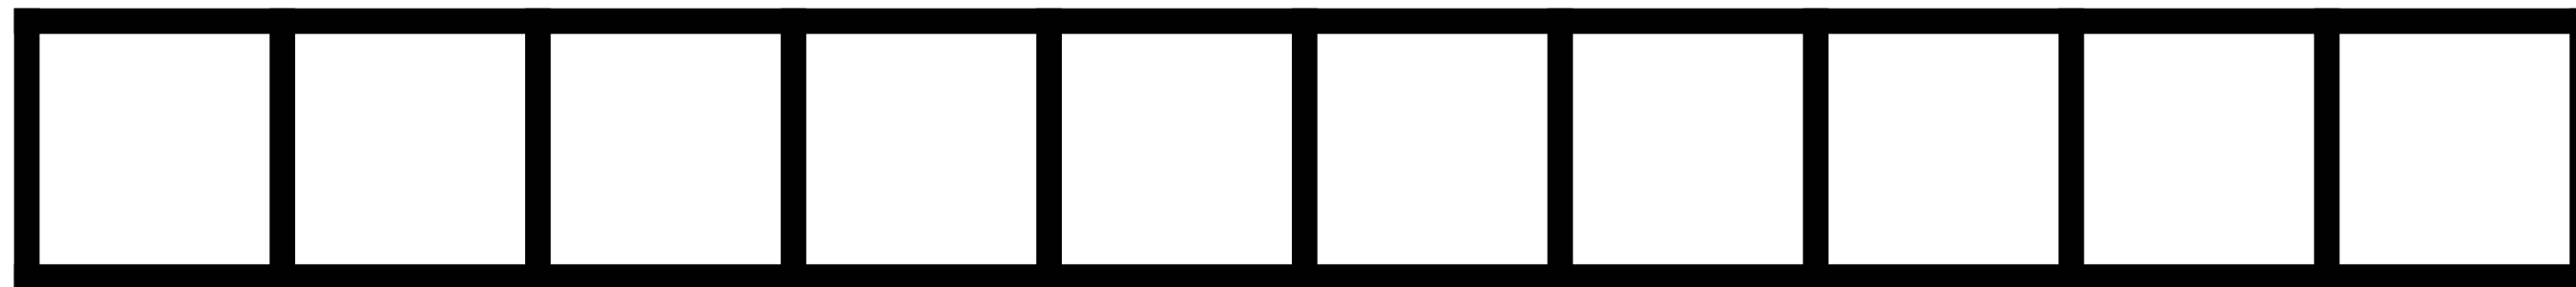
push

pop



Review: priority queue

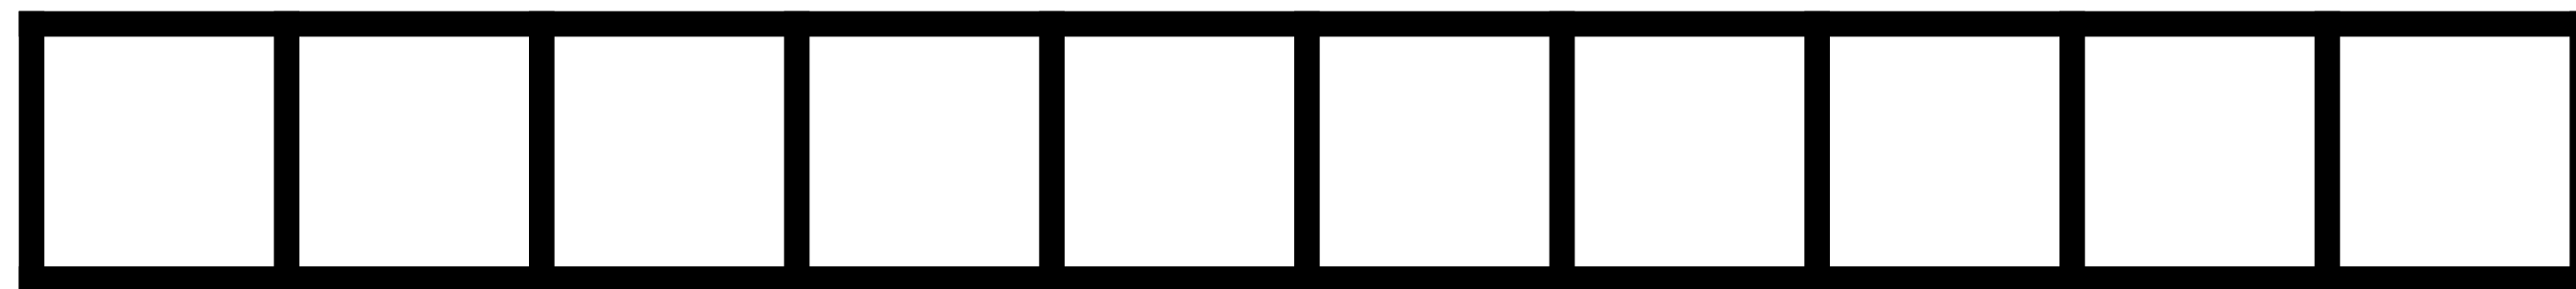
Everything from before holds,
but we have a little more control.



Review: priority queue

Everything from before holds,
but we have a little more control.

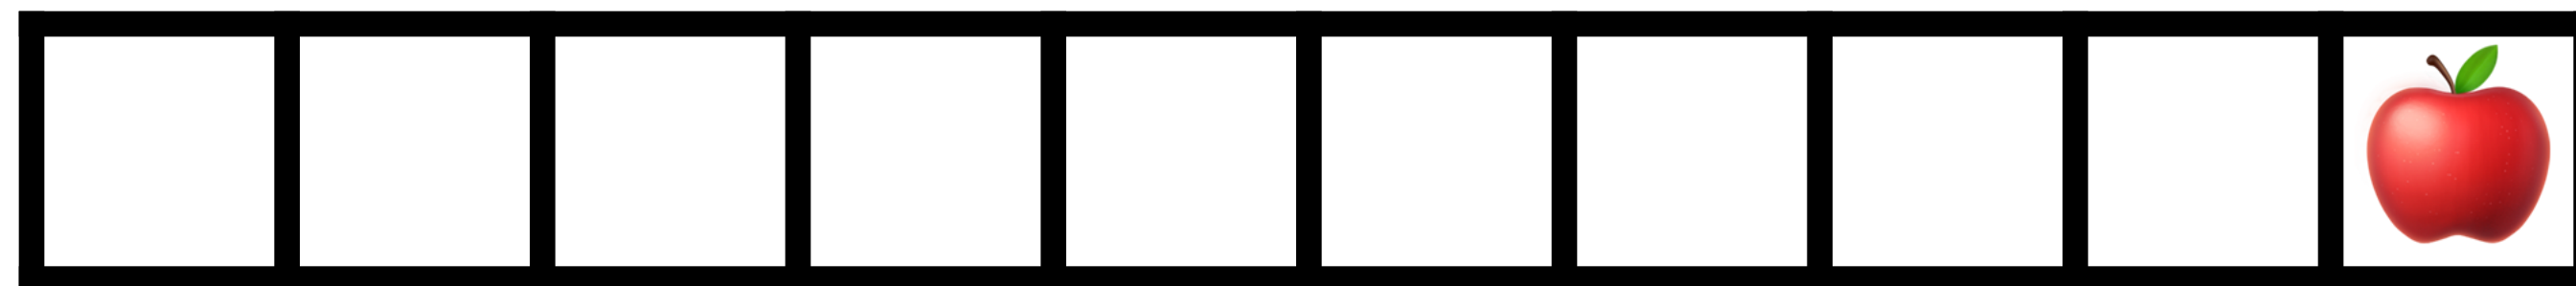
Say we have a queue *prioritized by pH*.



Review: priority queue

Everything from before holds,
but we have a little more control.

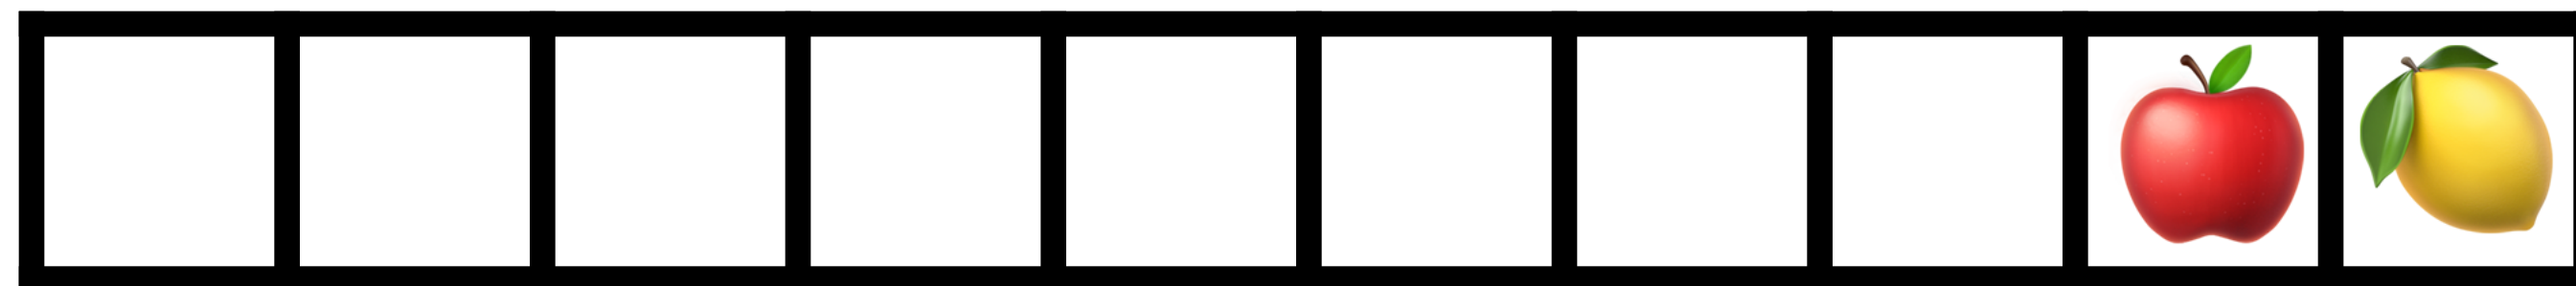
Say we have a queue *prioritized by pH*.



Review: priority queue

Everything from before holds,
but we have a little more control.

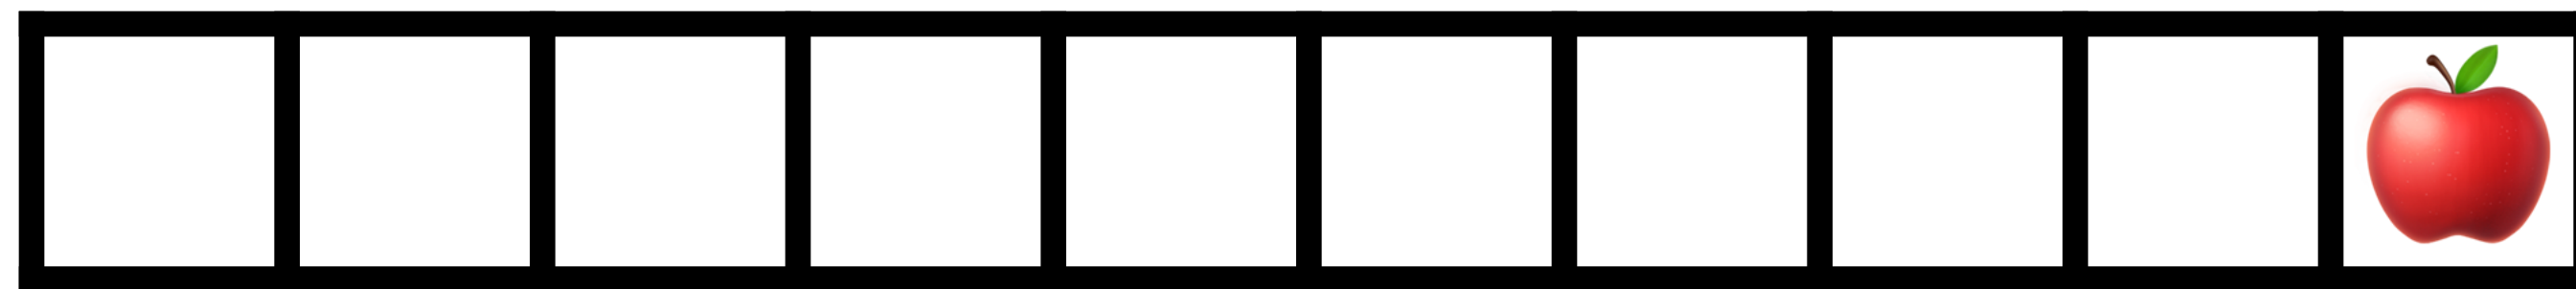
Say we have a queue *prioritized by pH*.



Review: priority queue

Everything from before holds,
but we have a little more control.

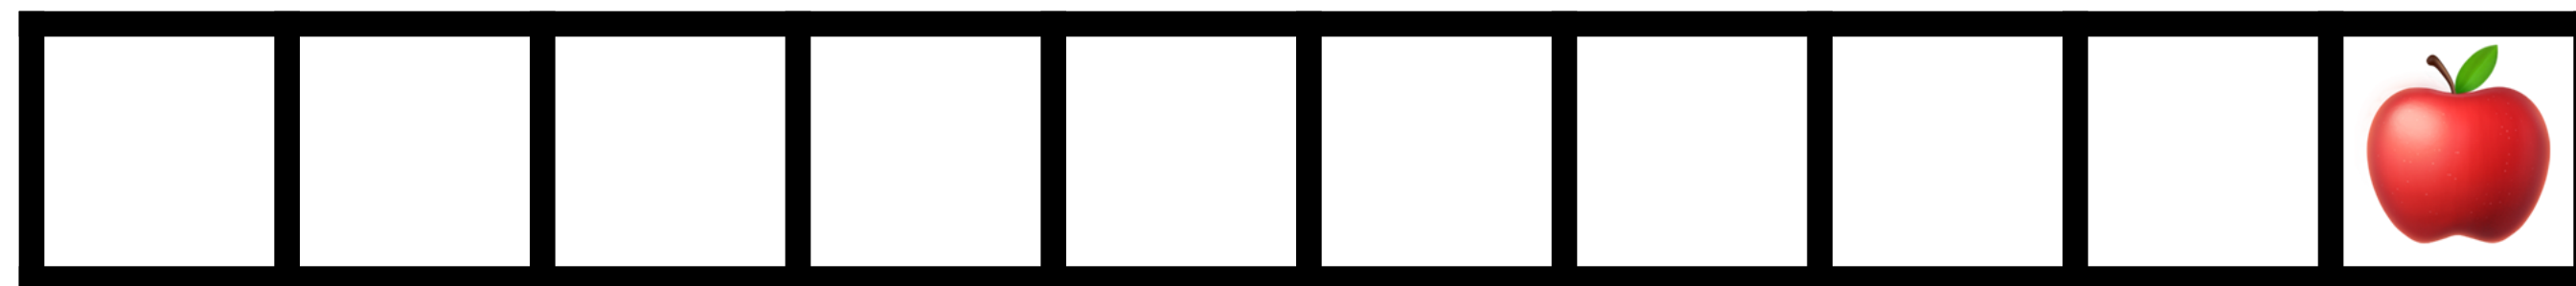
Say we have a queue *prioritized by pH*.



Review: priority queue

Everything from before holds,
but we have a little more control.

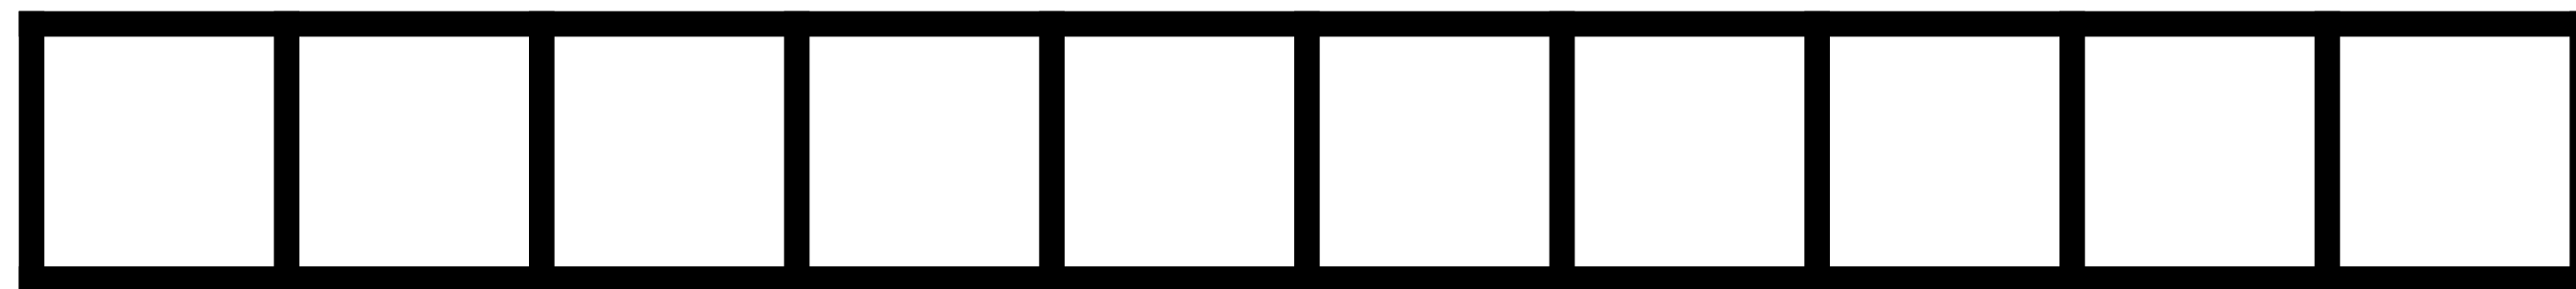
Say we have a queue *prioritized by pH*.



Review: priority queue

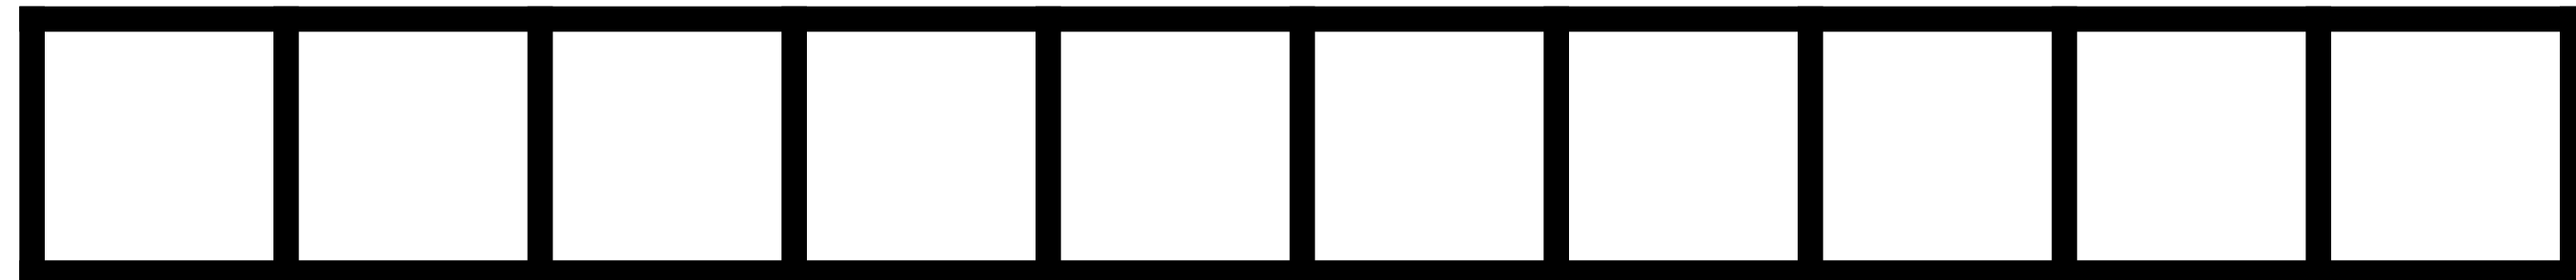
Everything from before holds,
but we have a little more control.

Say we have a queue *prioritized by pH*.



Review: priority queue

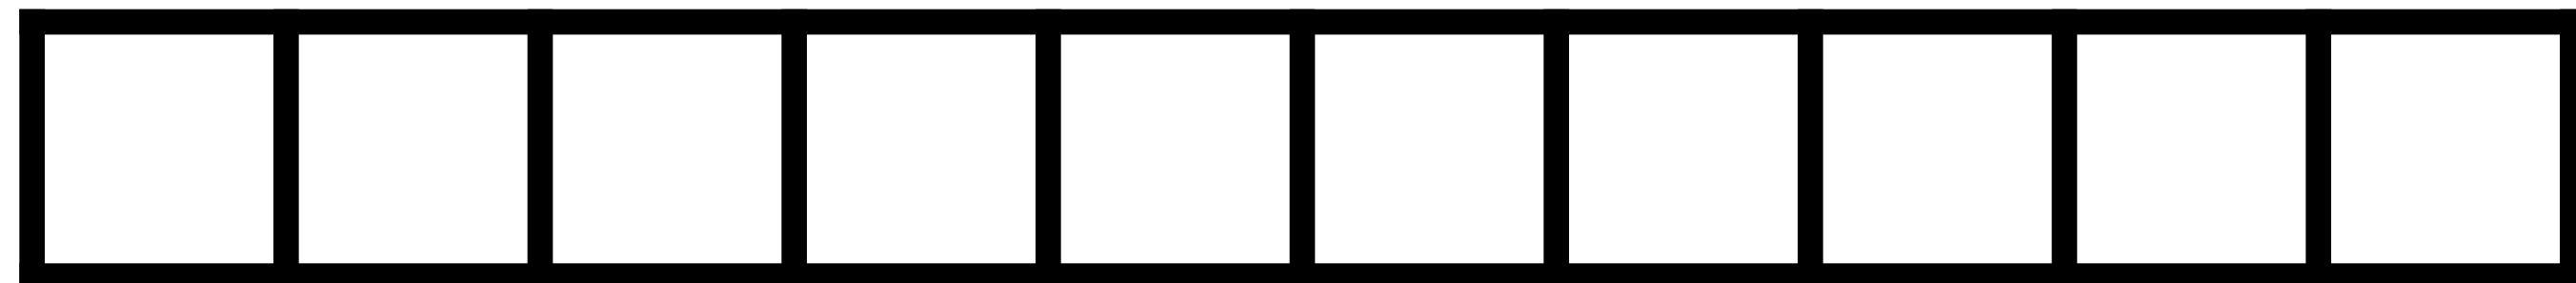
The priority need not be inherent to the item!



Review: priority queue

The priority need not be inherent to the item!

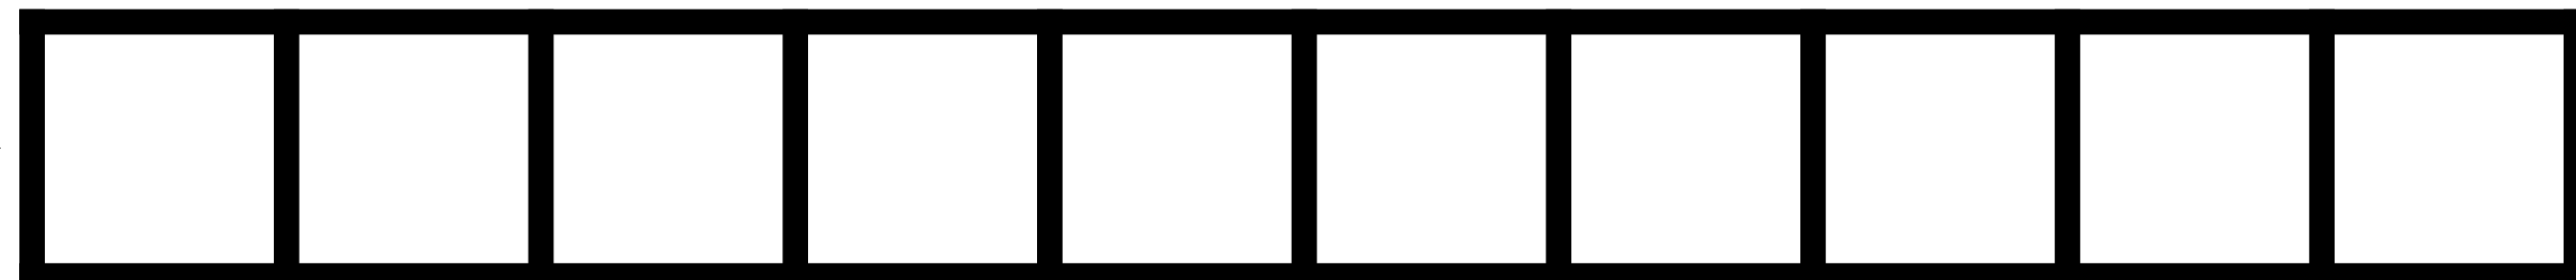
We can have a *ranking function*:



Review: priority queue

The priority need not be inherent to the item!

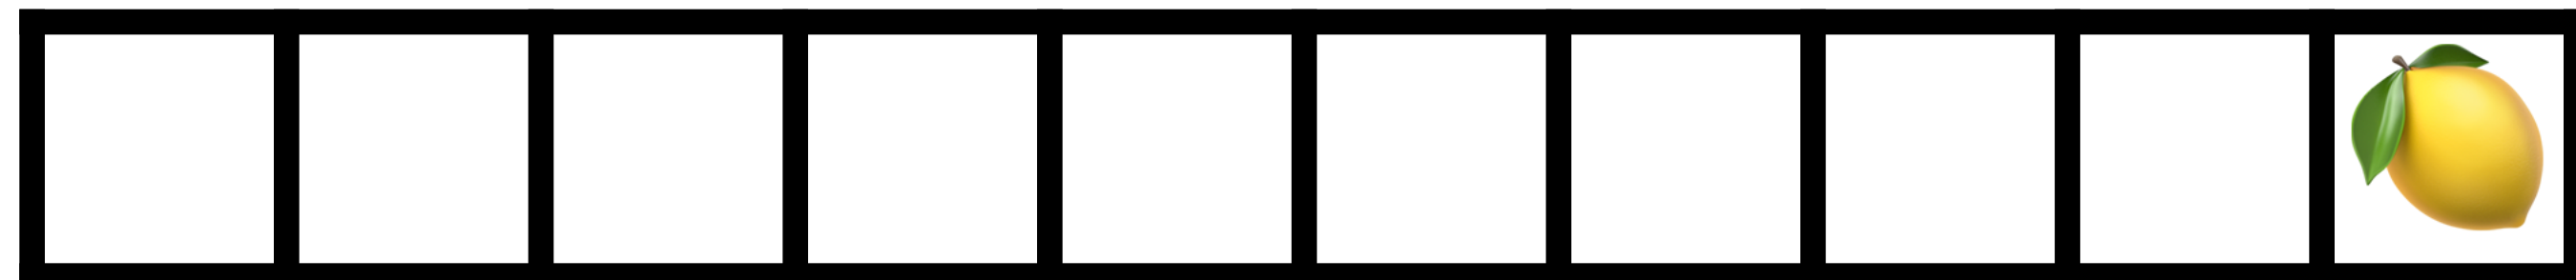
We can have a *ranking function*:



Review: priority queue

The priority need not be inherent to the item!

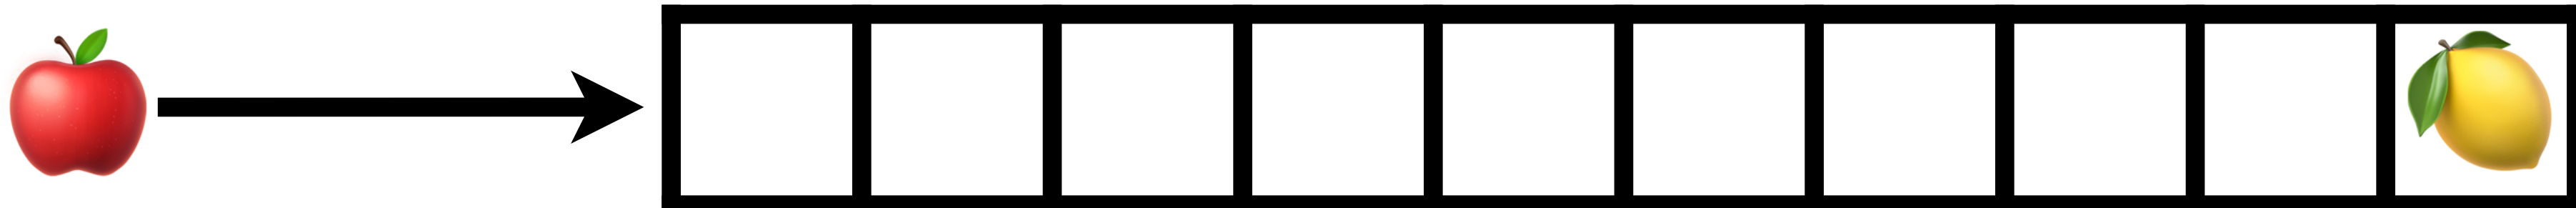
We can have a *ranking function*:



Review: priority queue

The priority need not be inherent to the item!

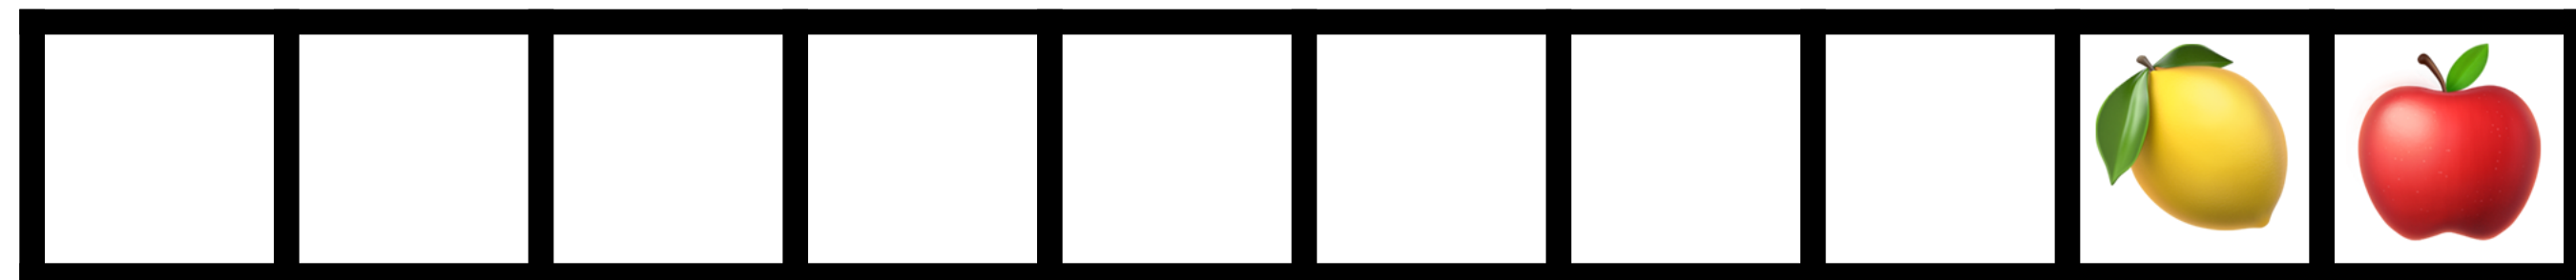
We can have a *ranking function*:



Review: priority queue

The priority need not be inherent to the item!

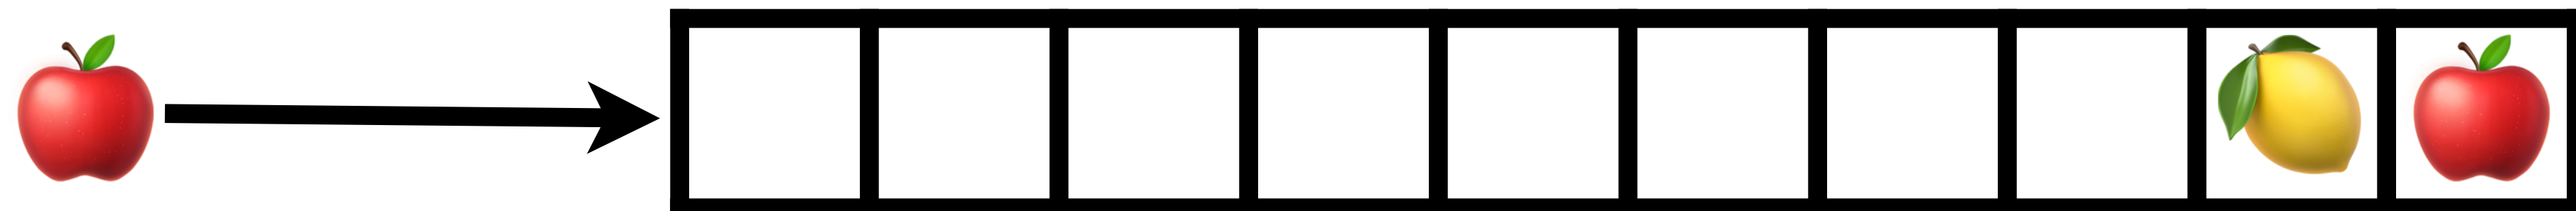
We can have a *ranking function*:



Review: priority queue

The priority need not be inherent to the item!

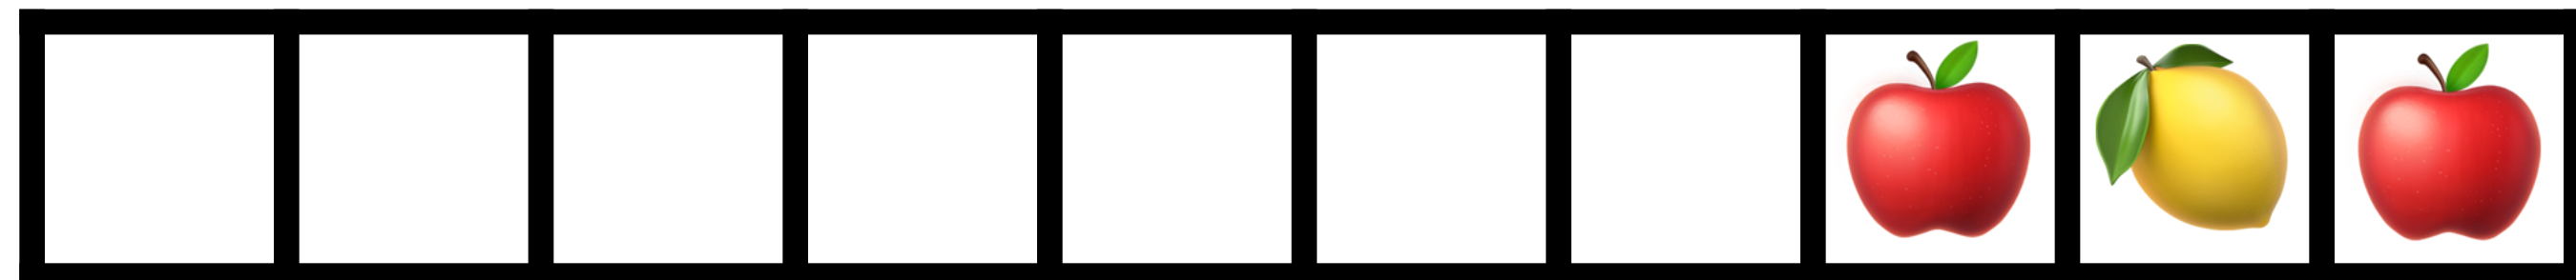
We can have a *ranking function*:



Review: priority queue

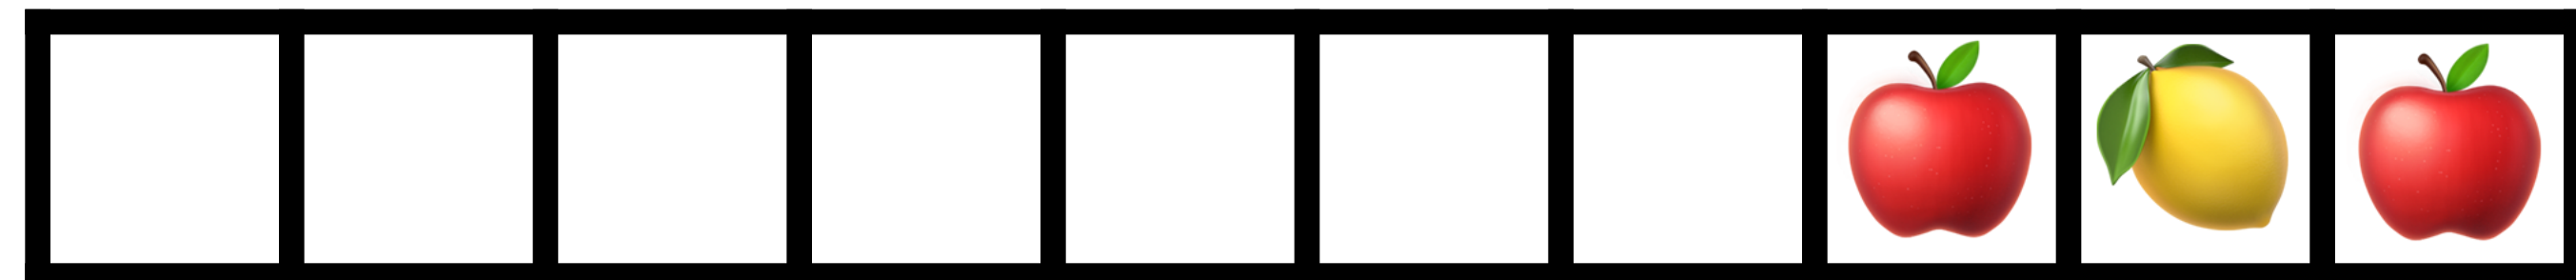
The priority need not be inherent to the item!

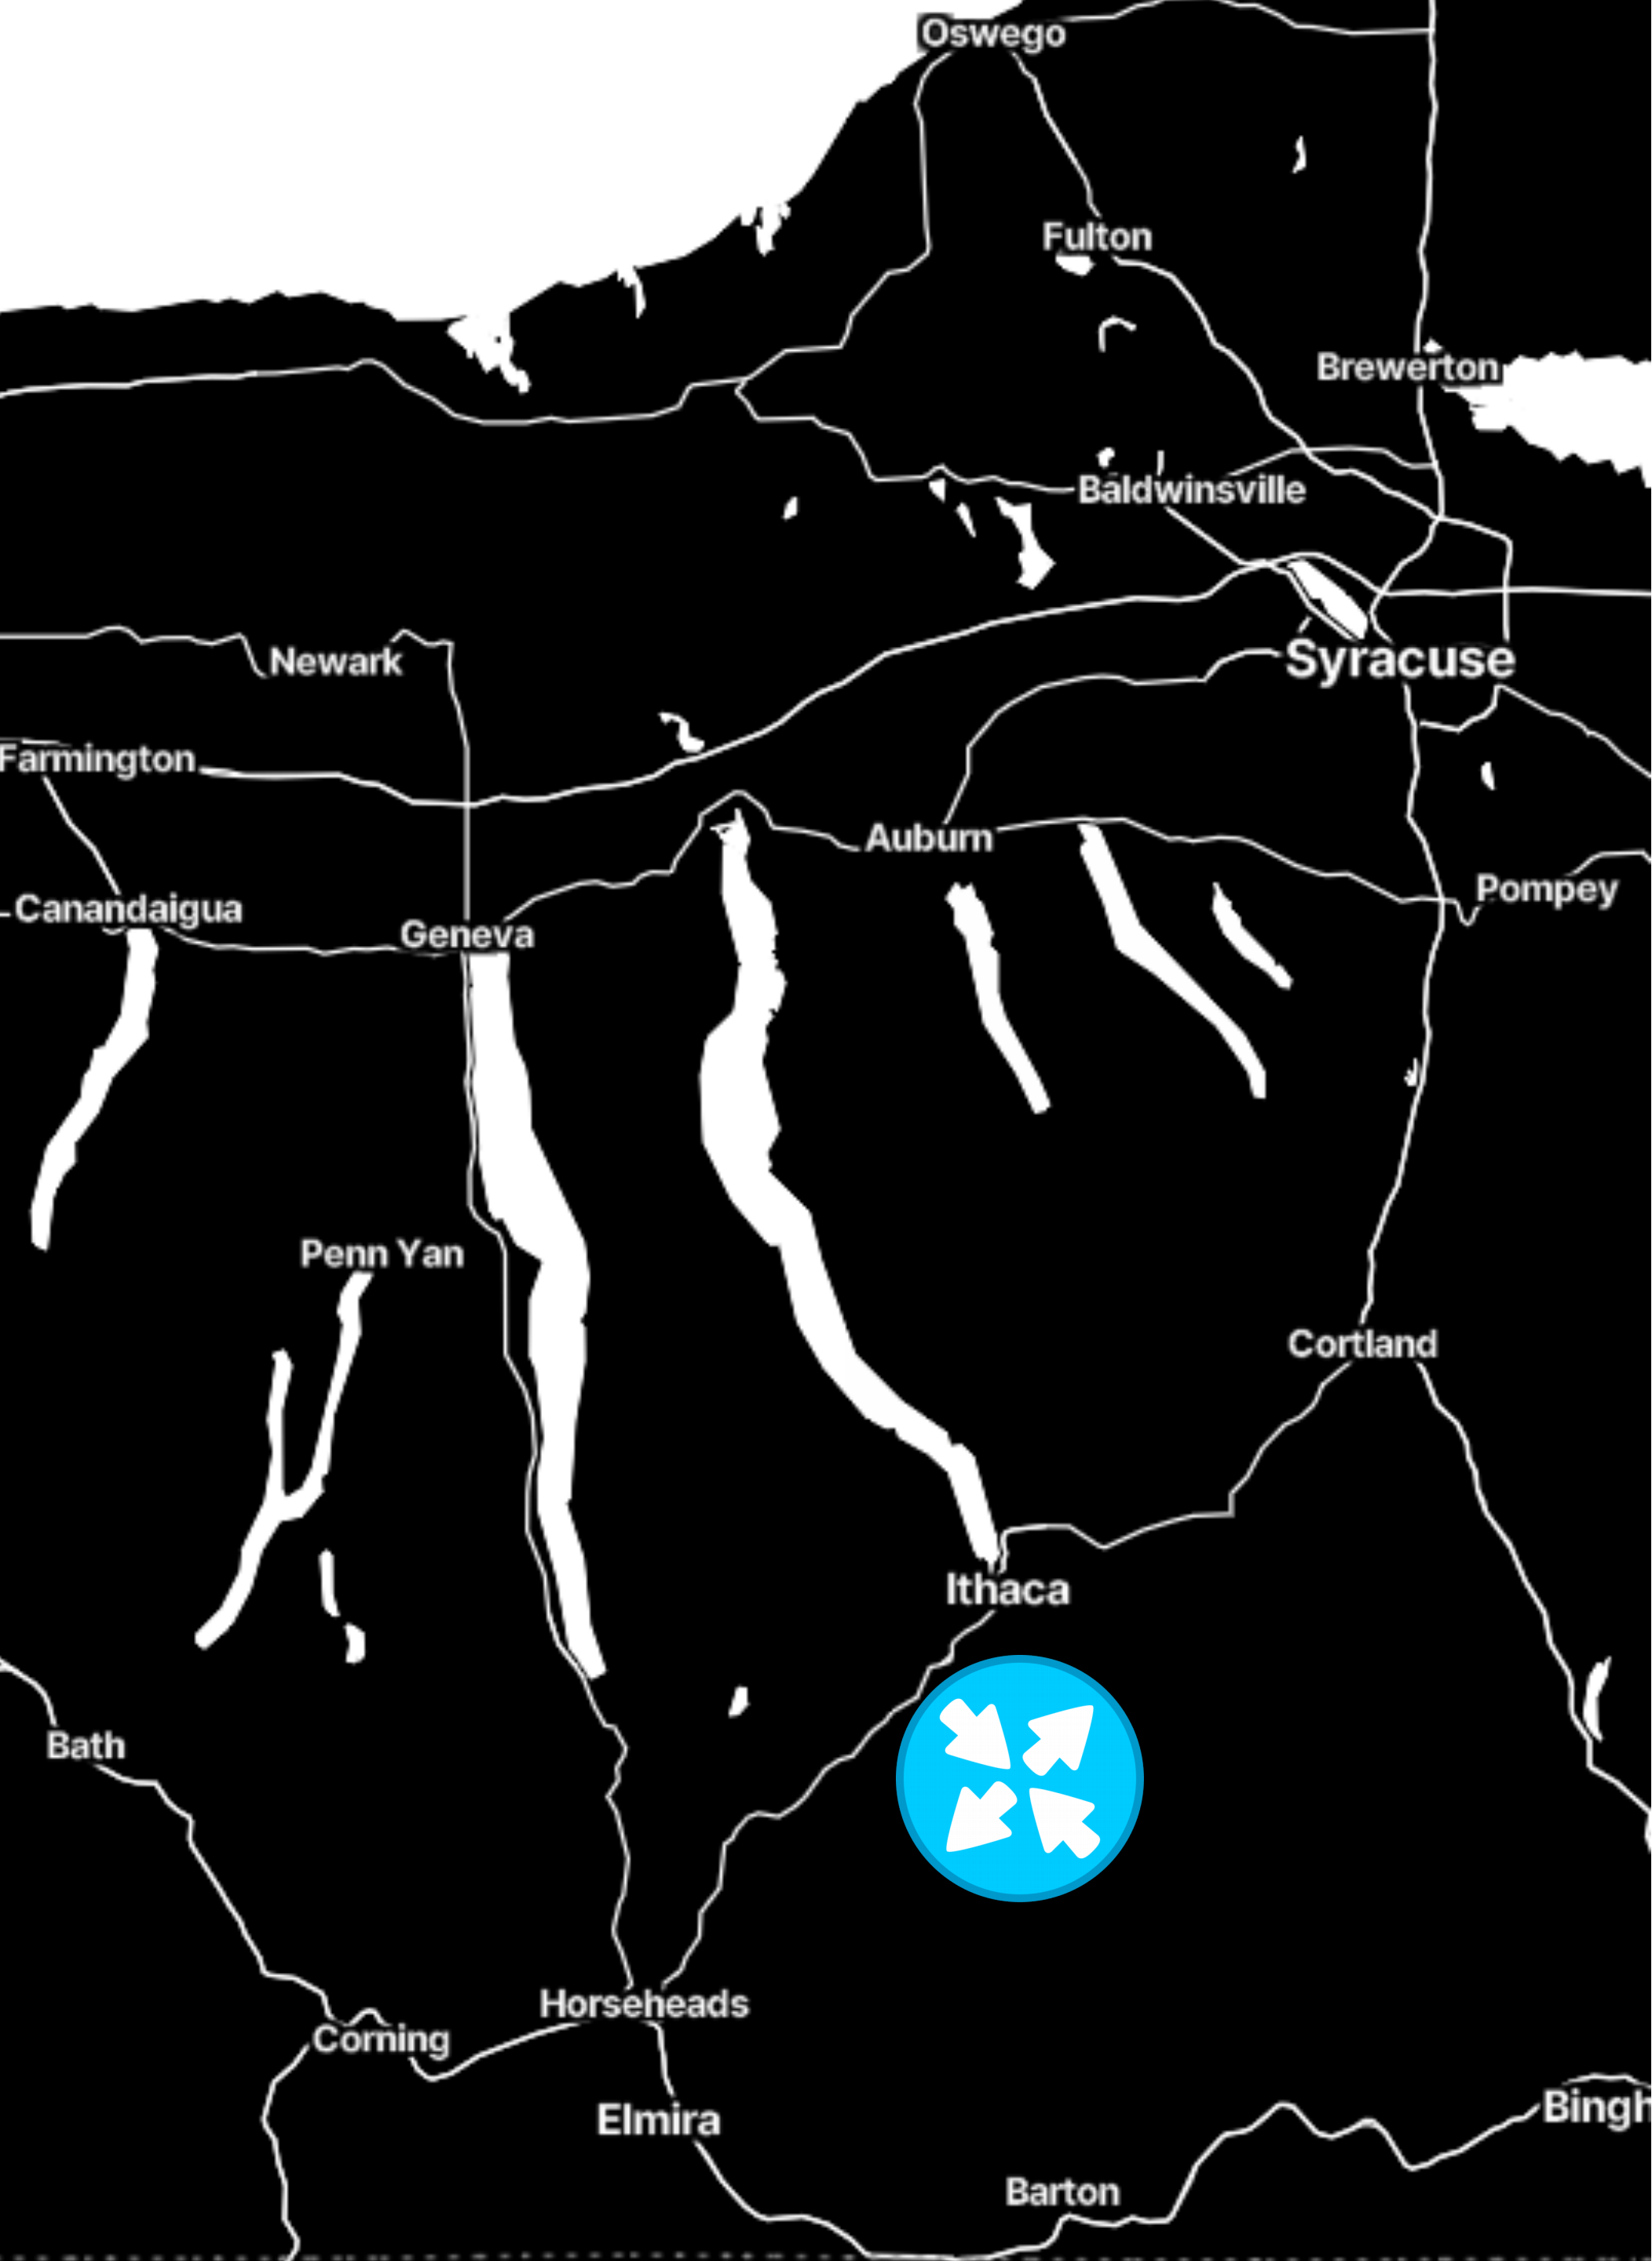
We can have a *ranking function*:

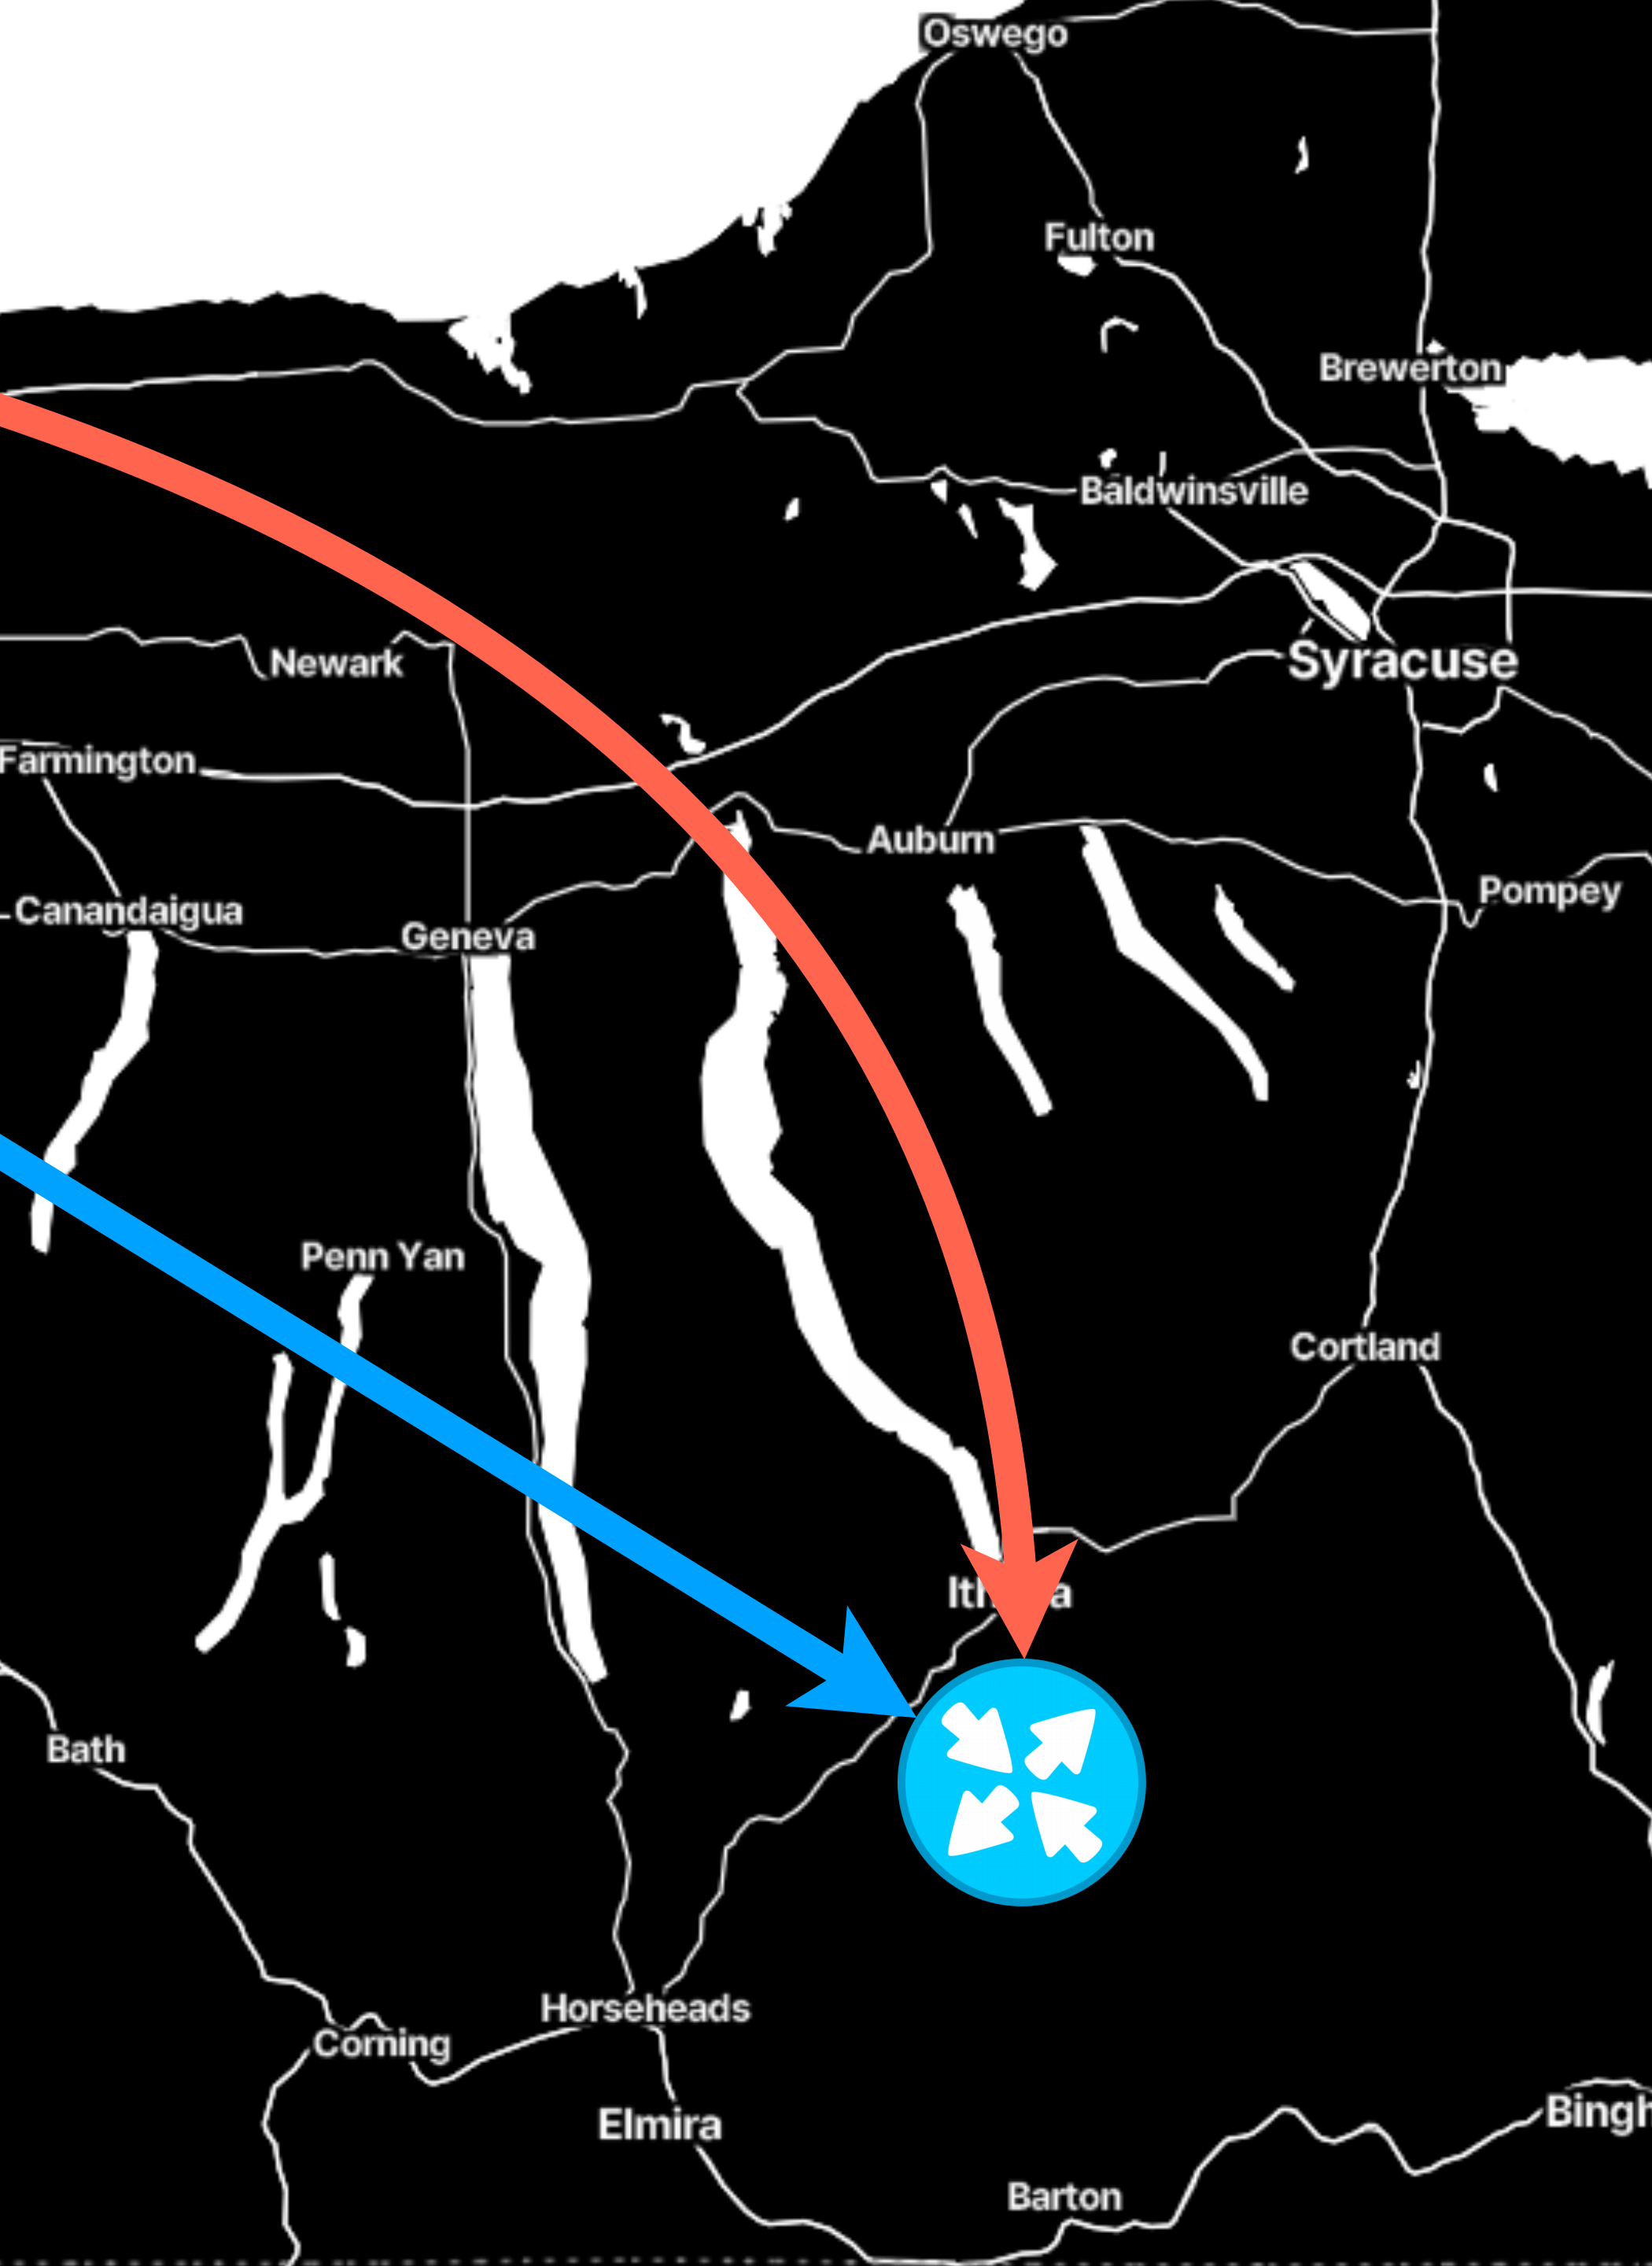


Introducing: PIFO

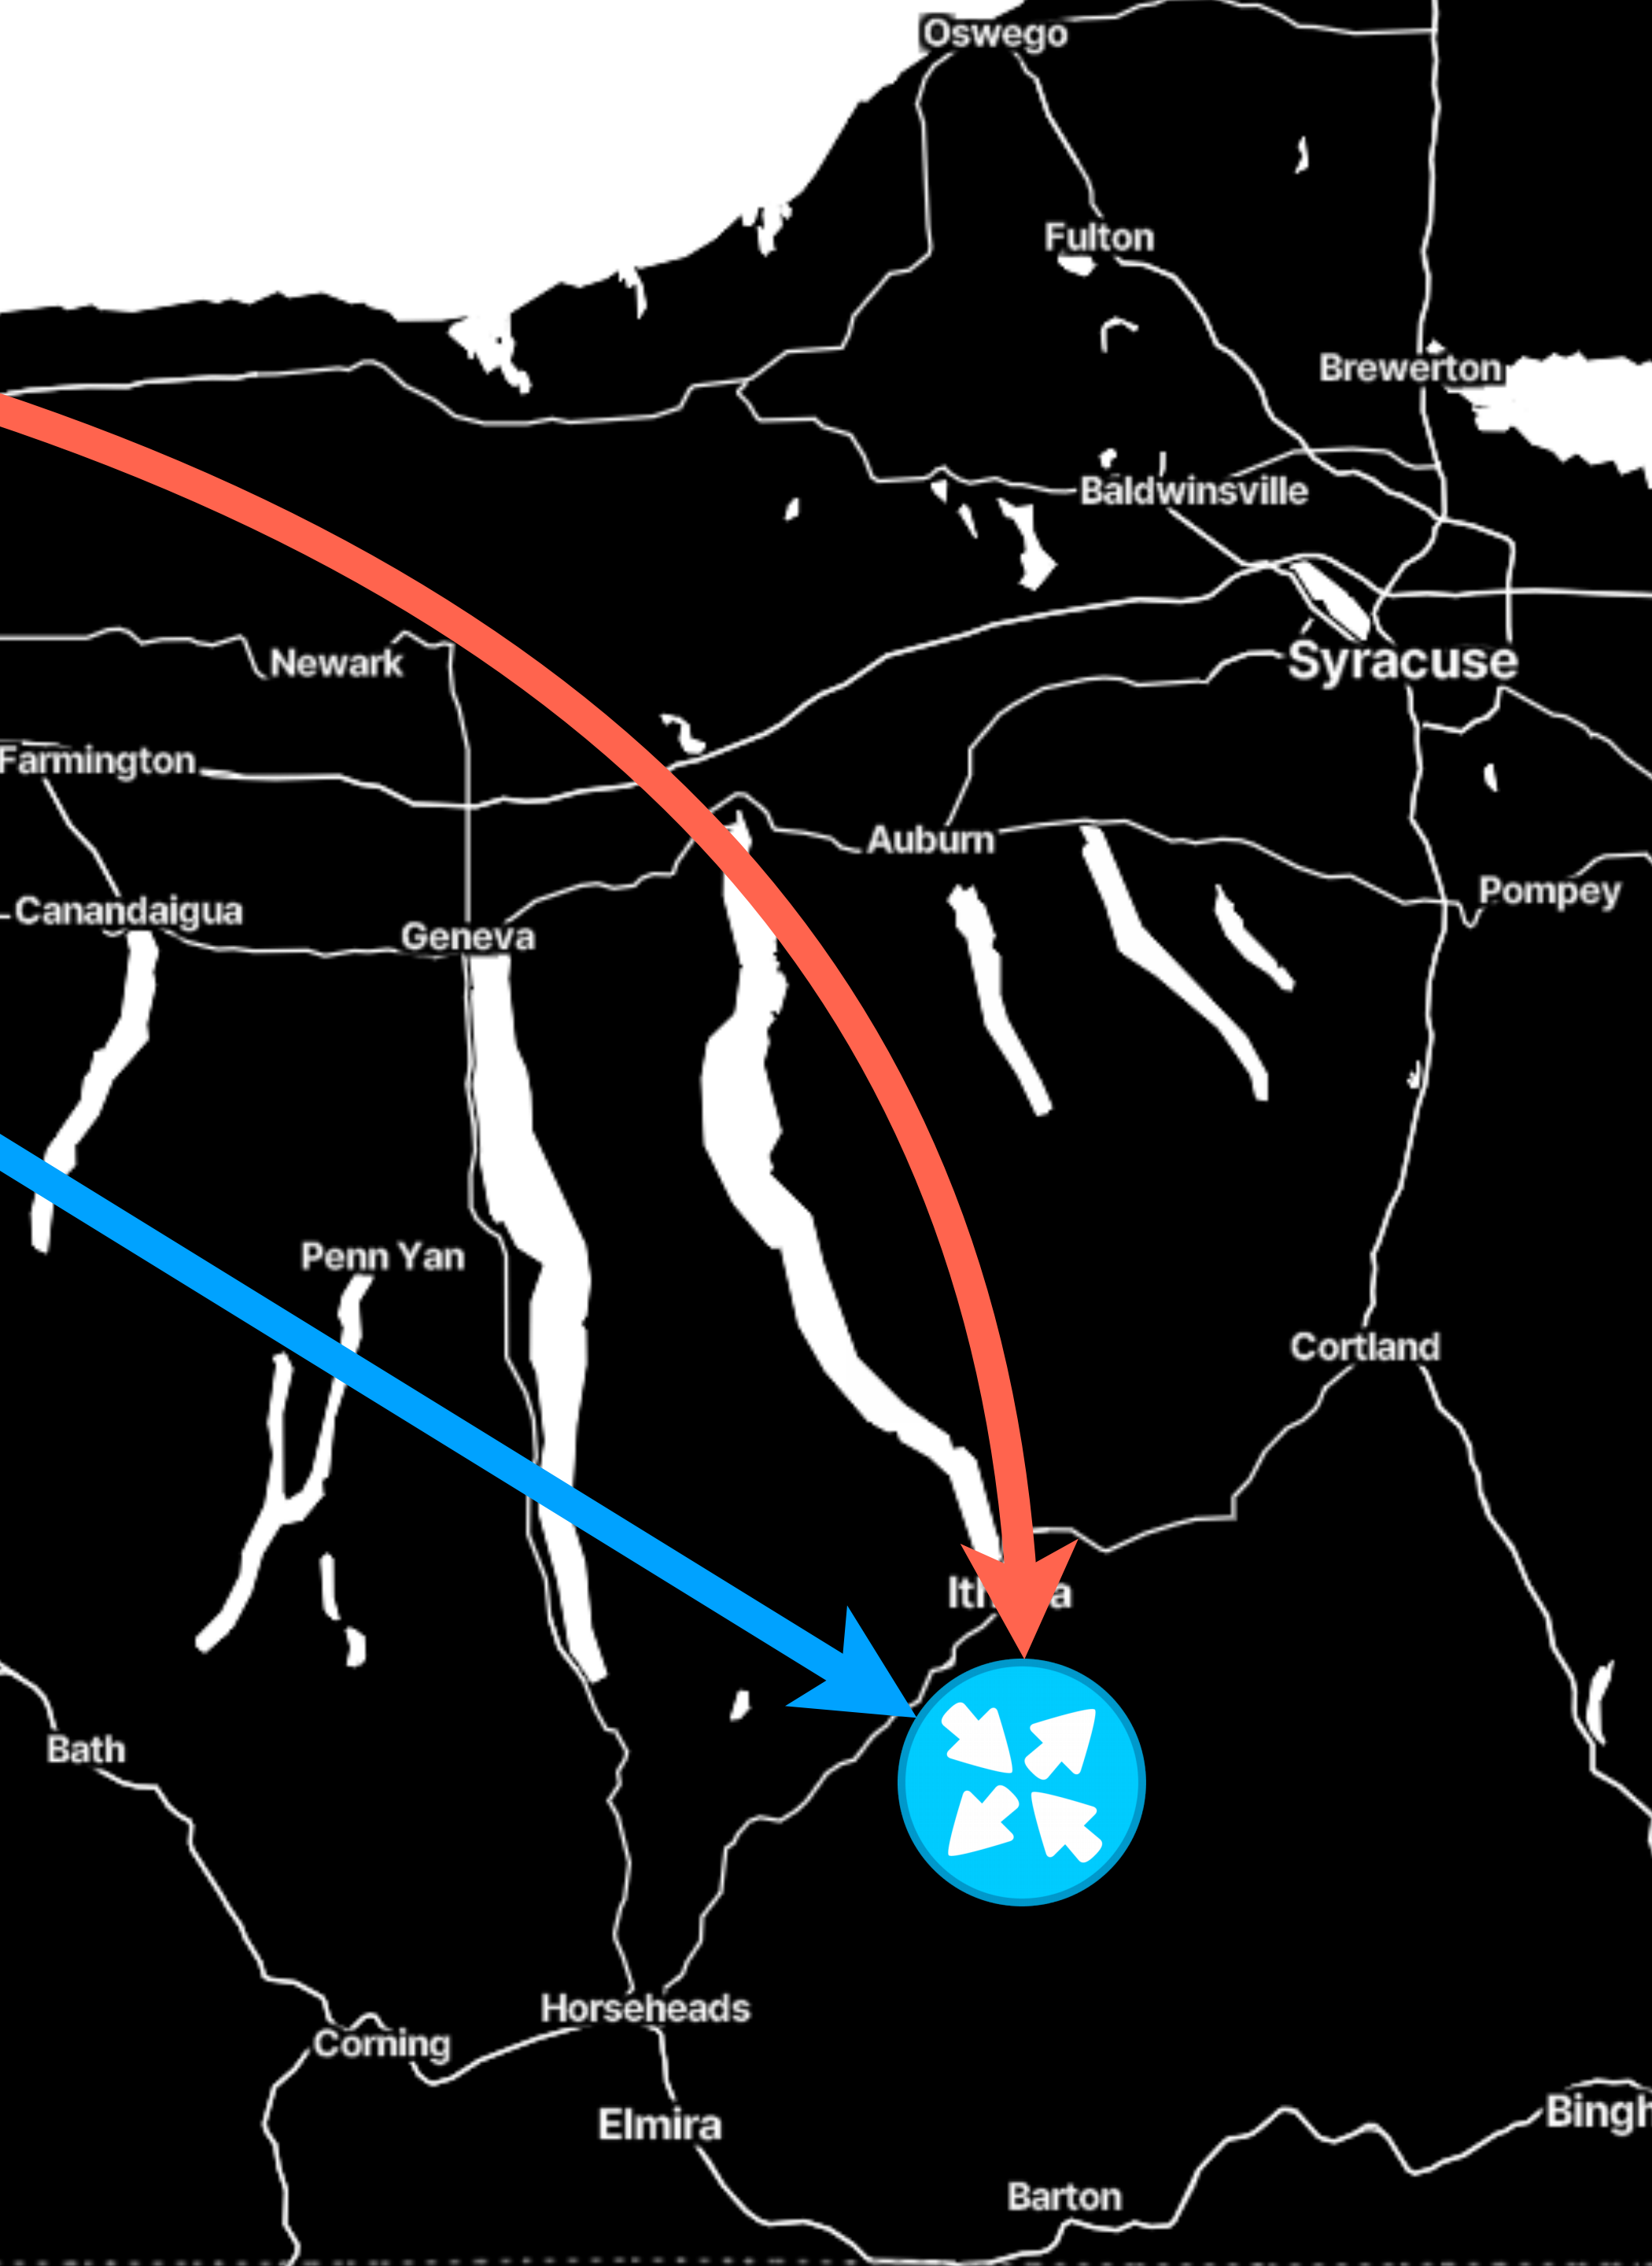
Just a PQ, with a ranking function,
but with *rank-ties* broken in FIFO order.





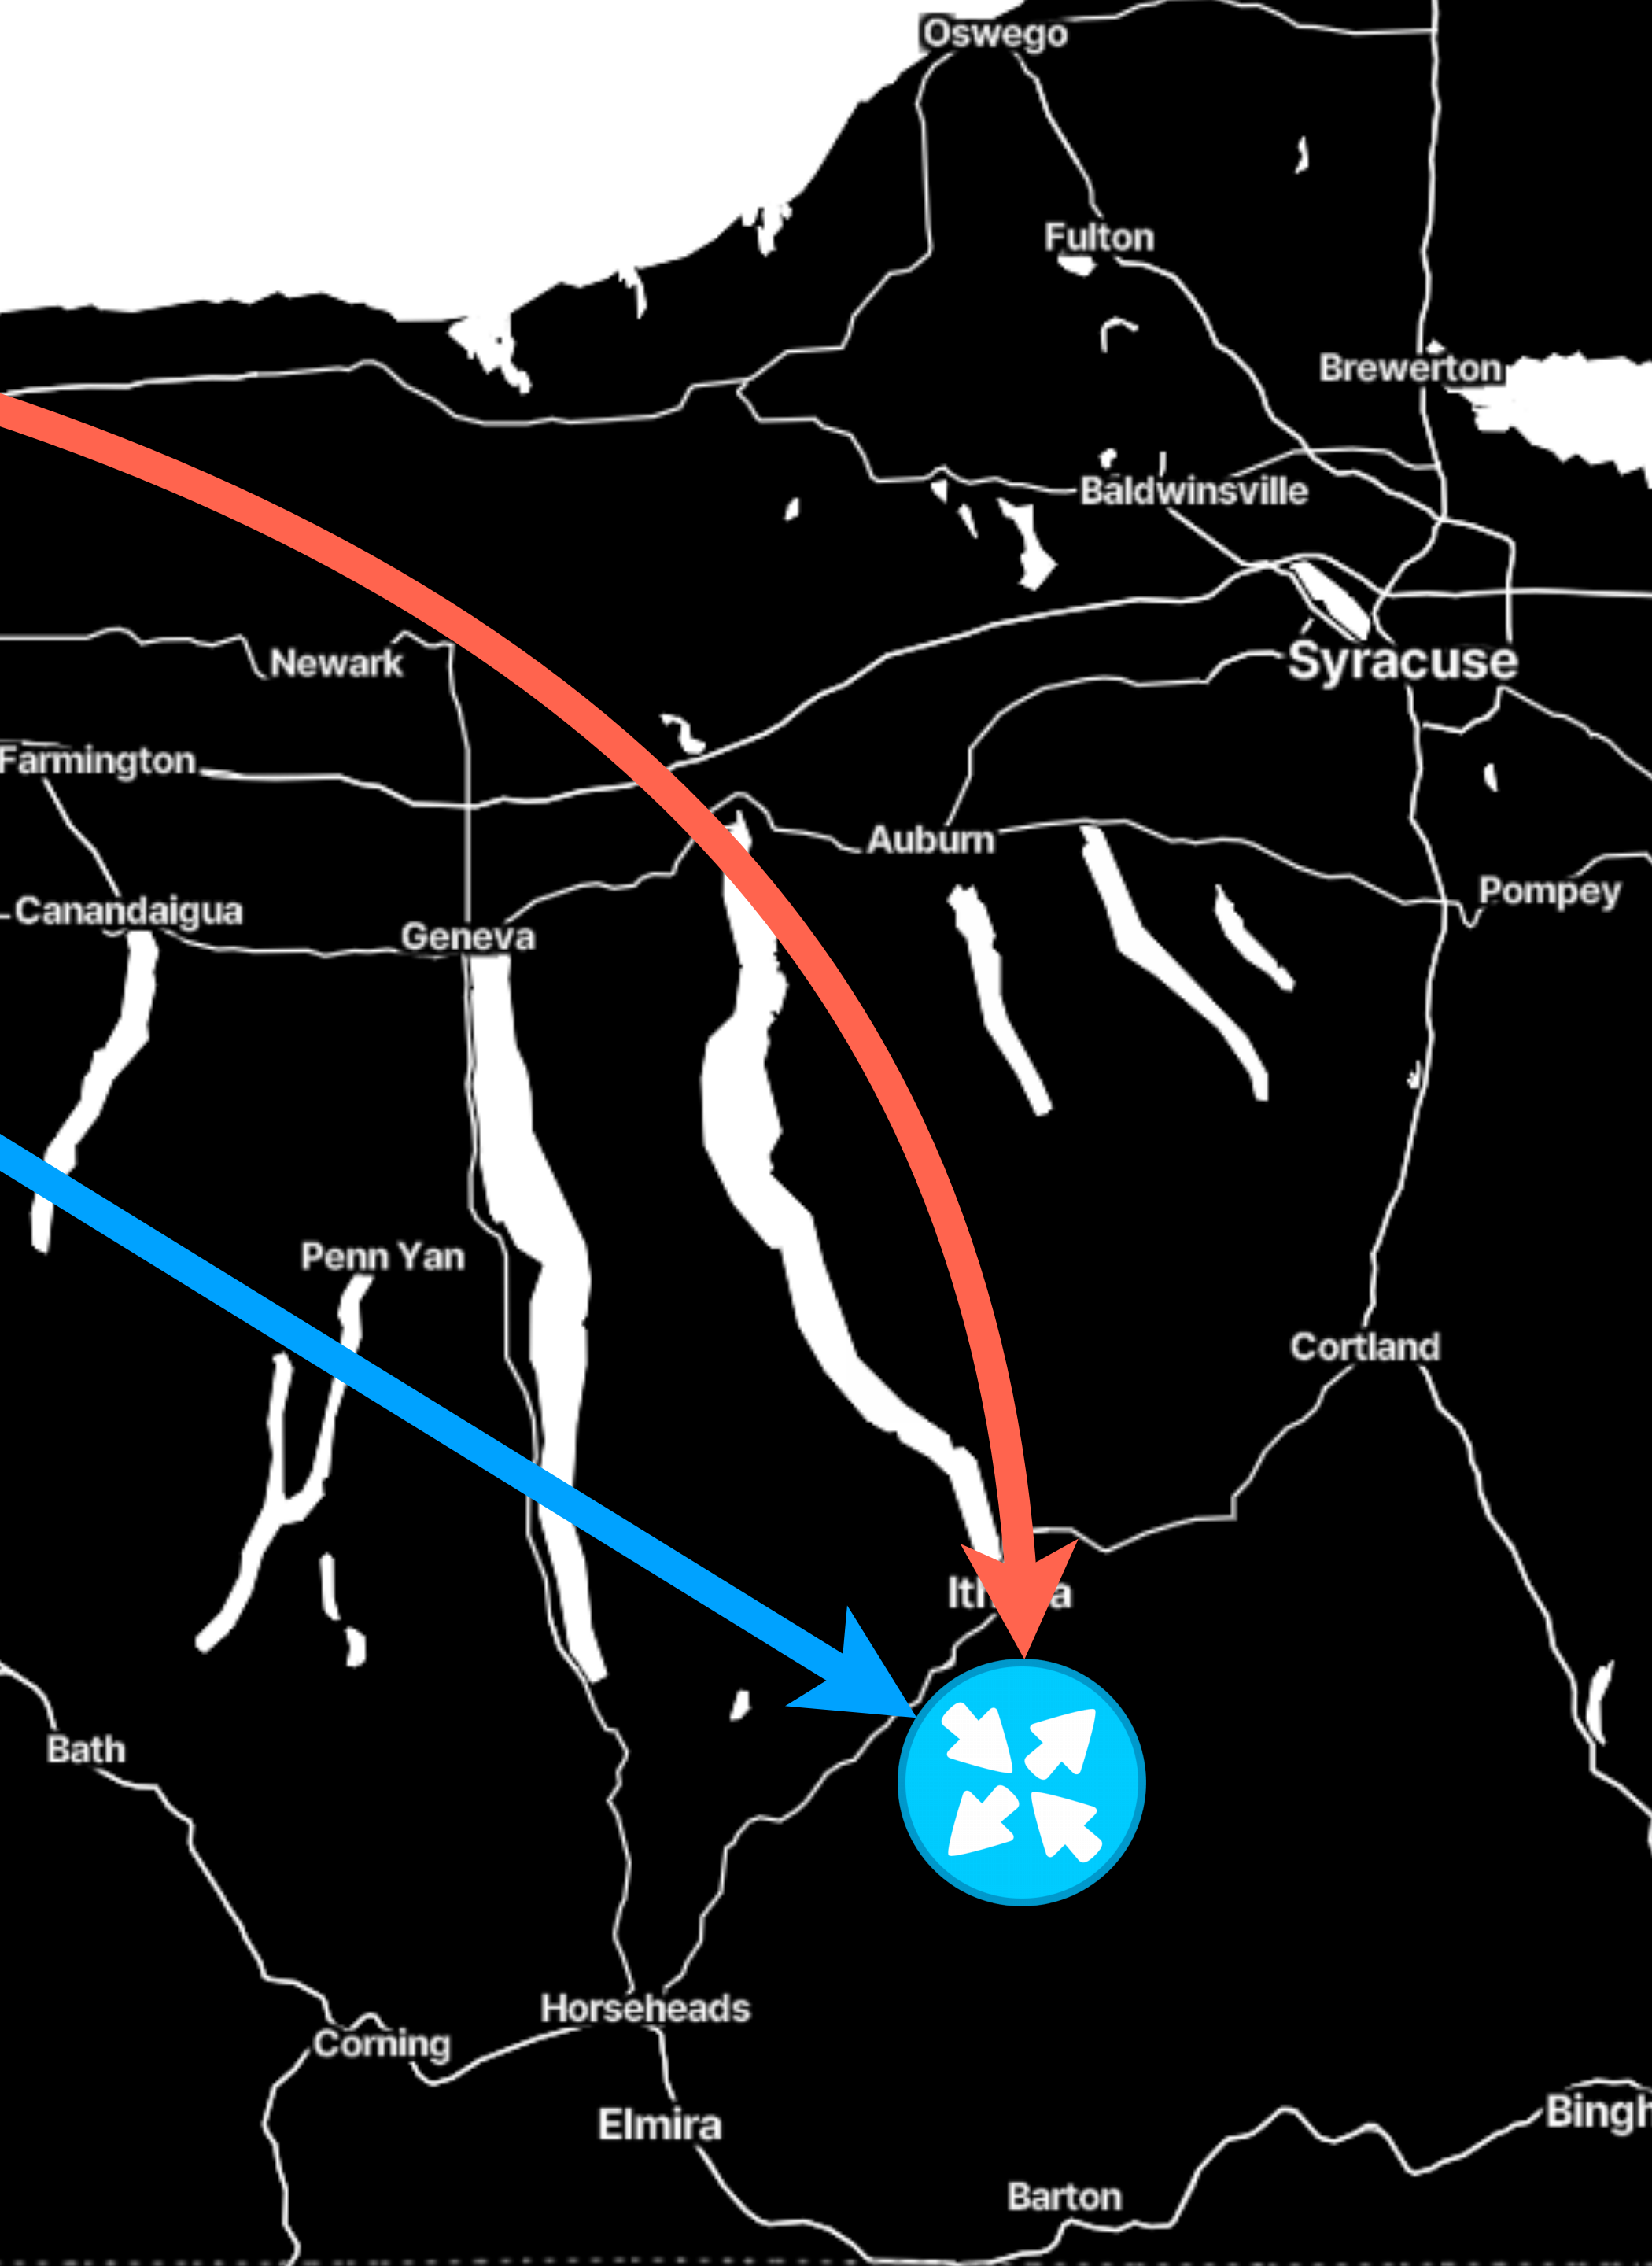


Baby goal:
interleave **R** and **B**.



Baby goal:
interleave **R** and **B**.

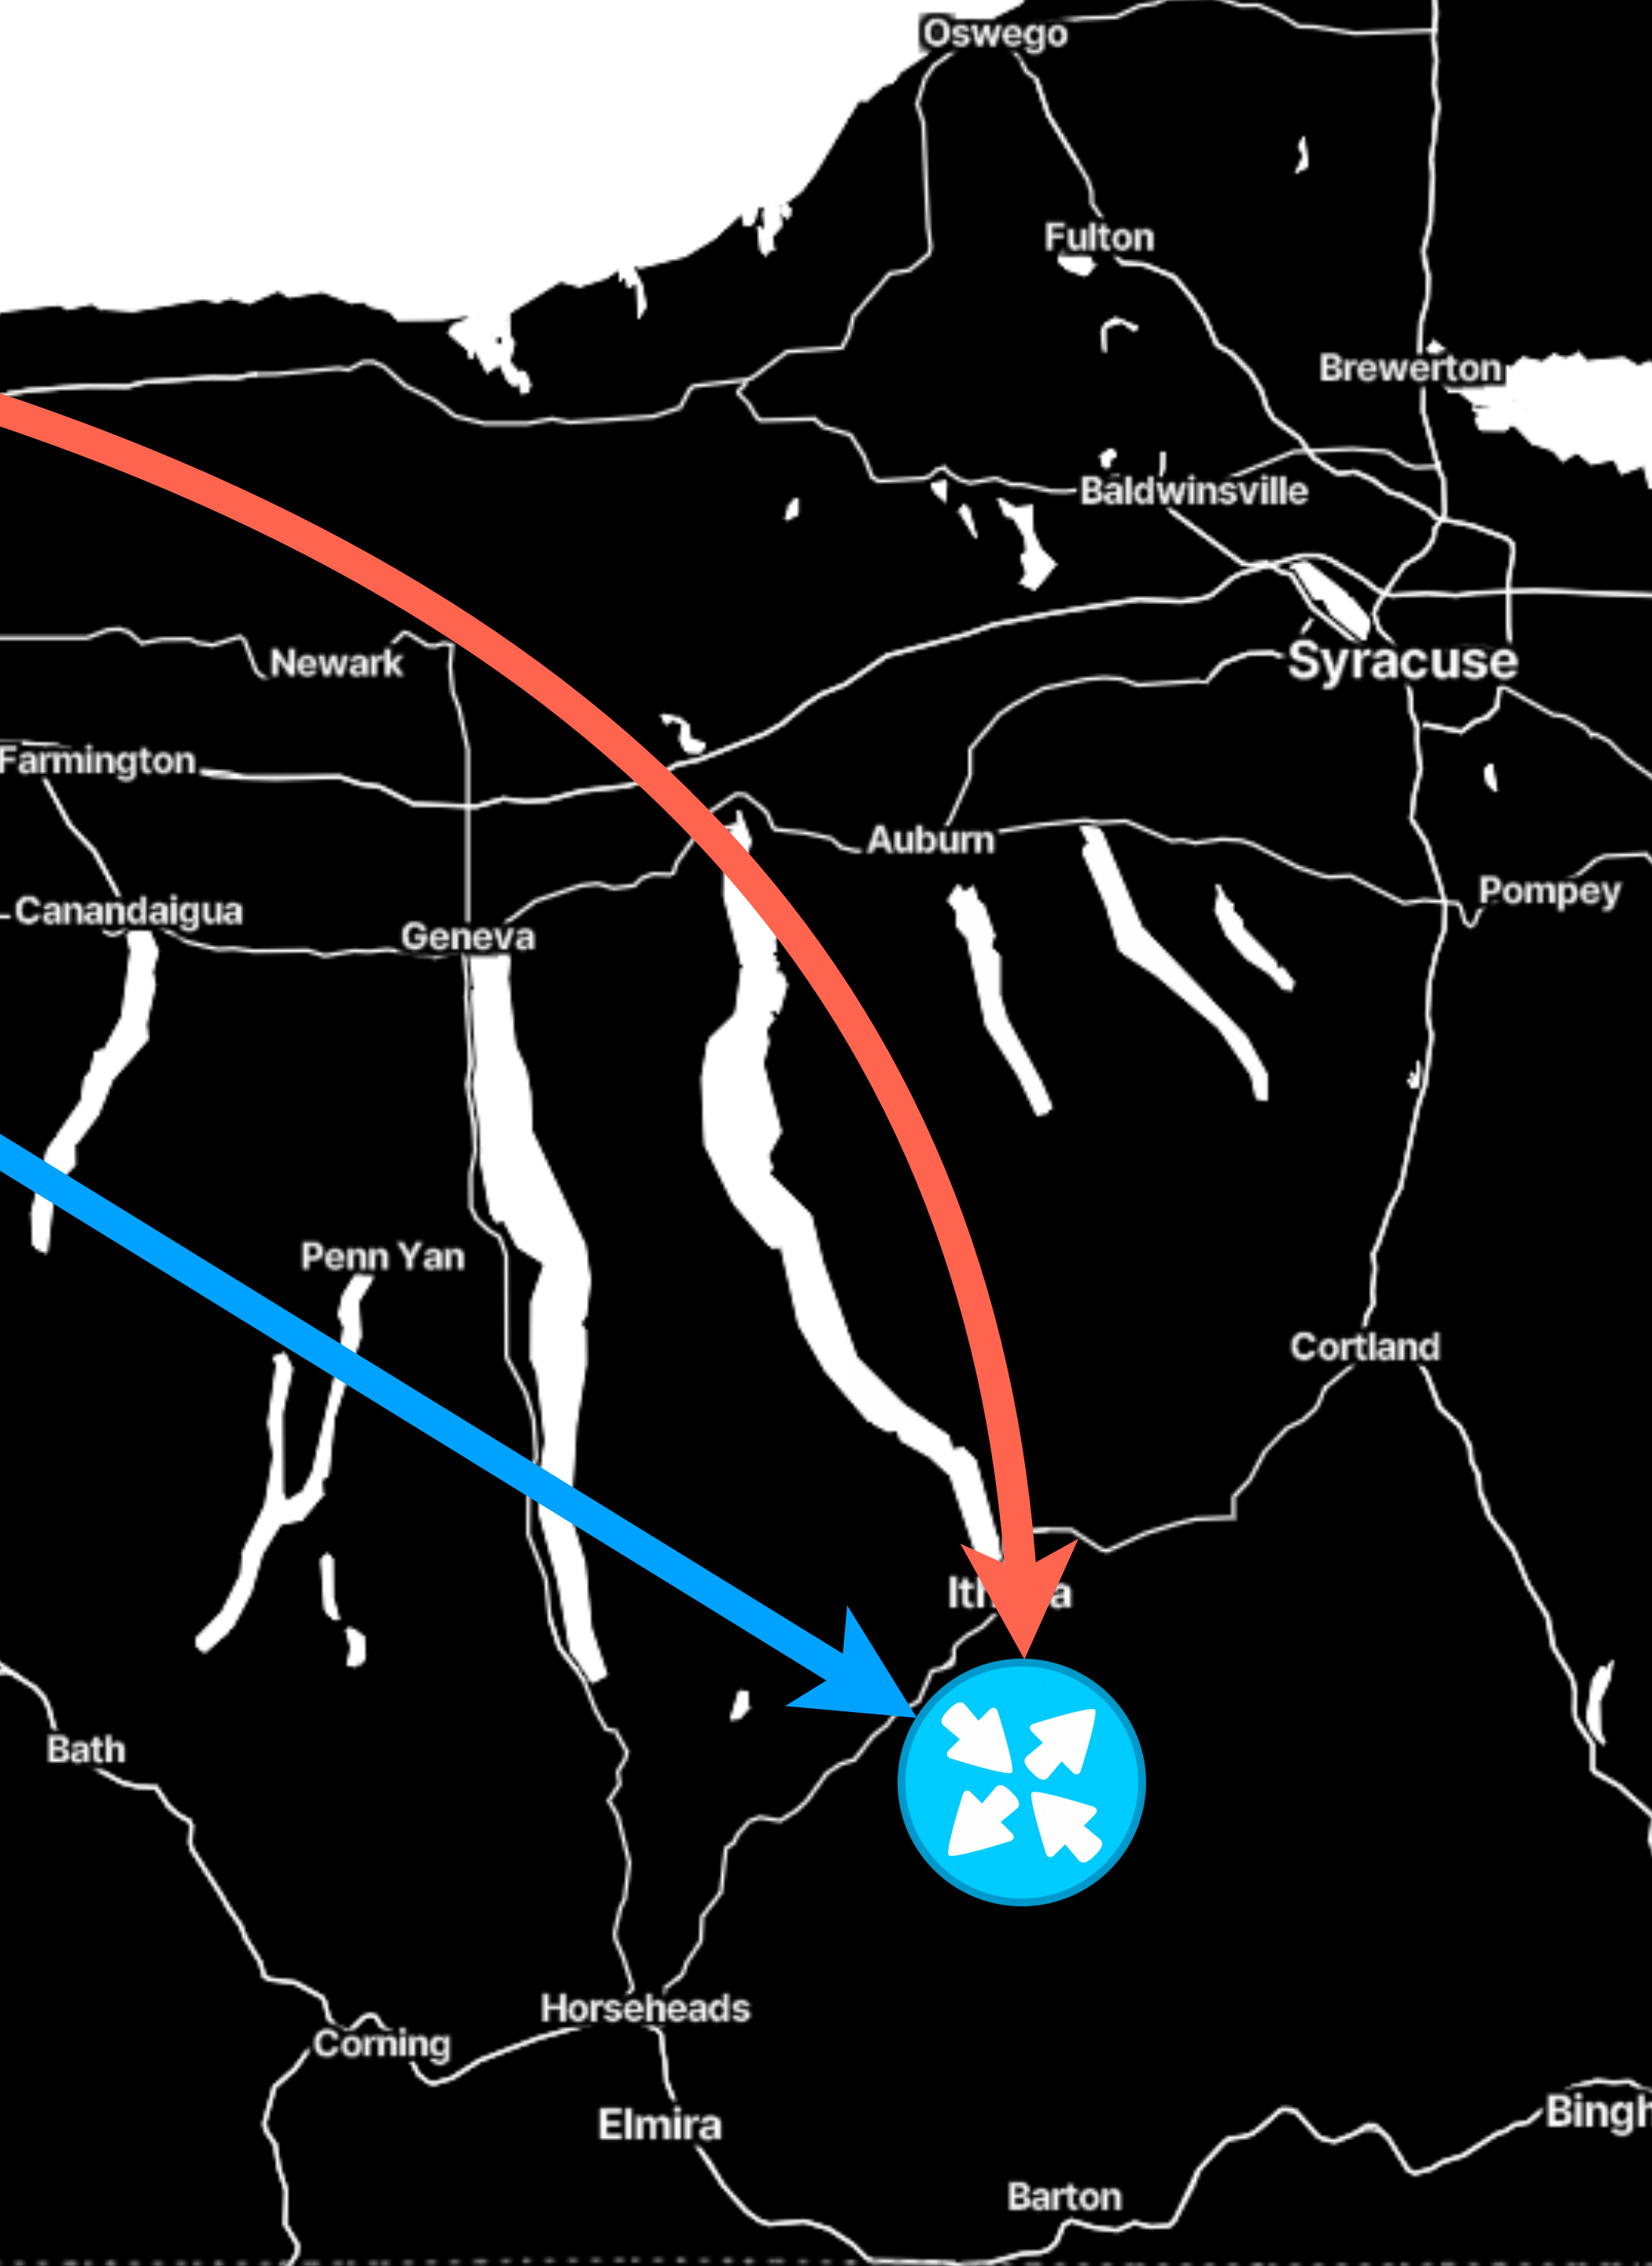
A PIFO will suffice.



Baby goal:
interleave **R** and **B**.

A PIFO will suffice.

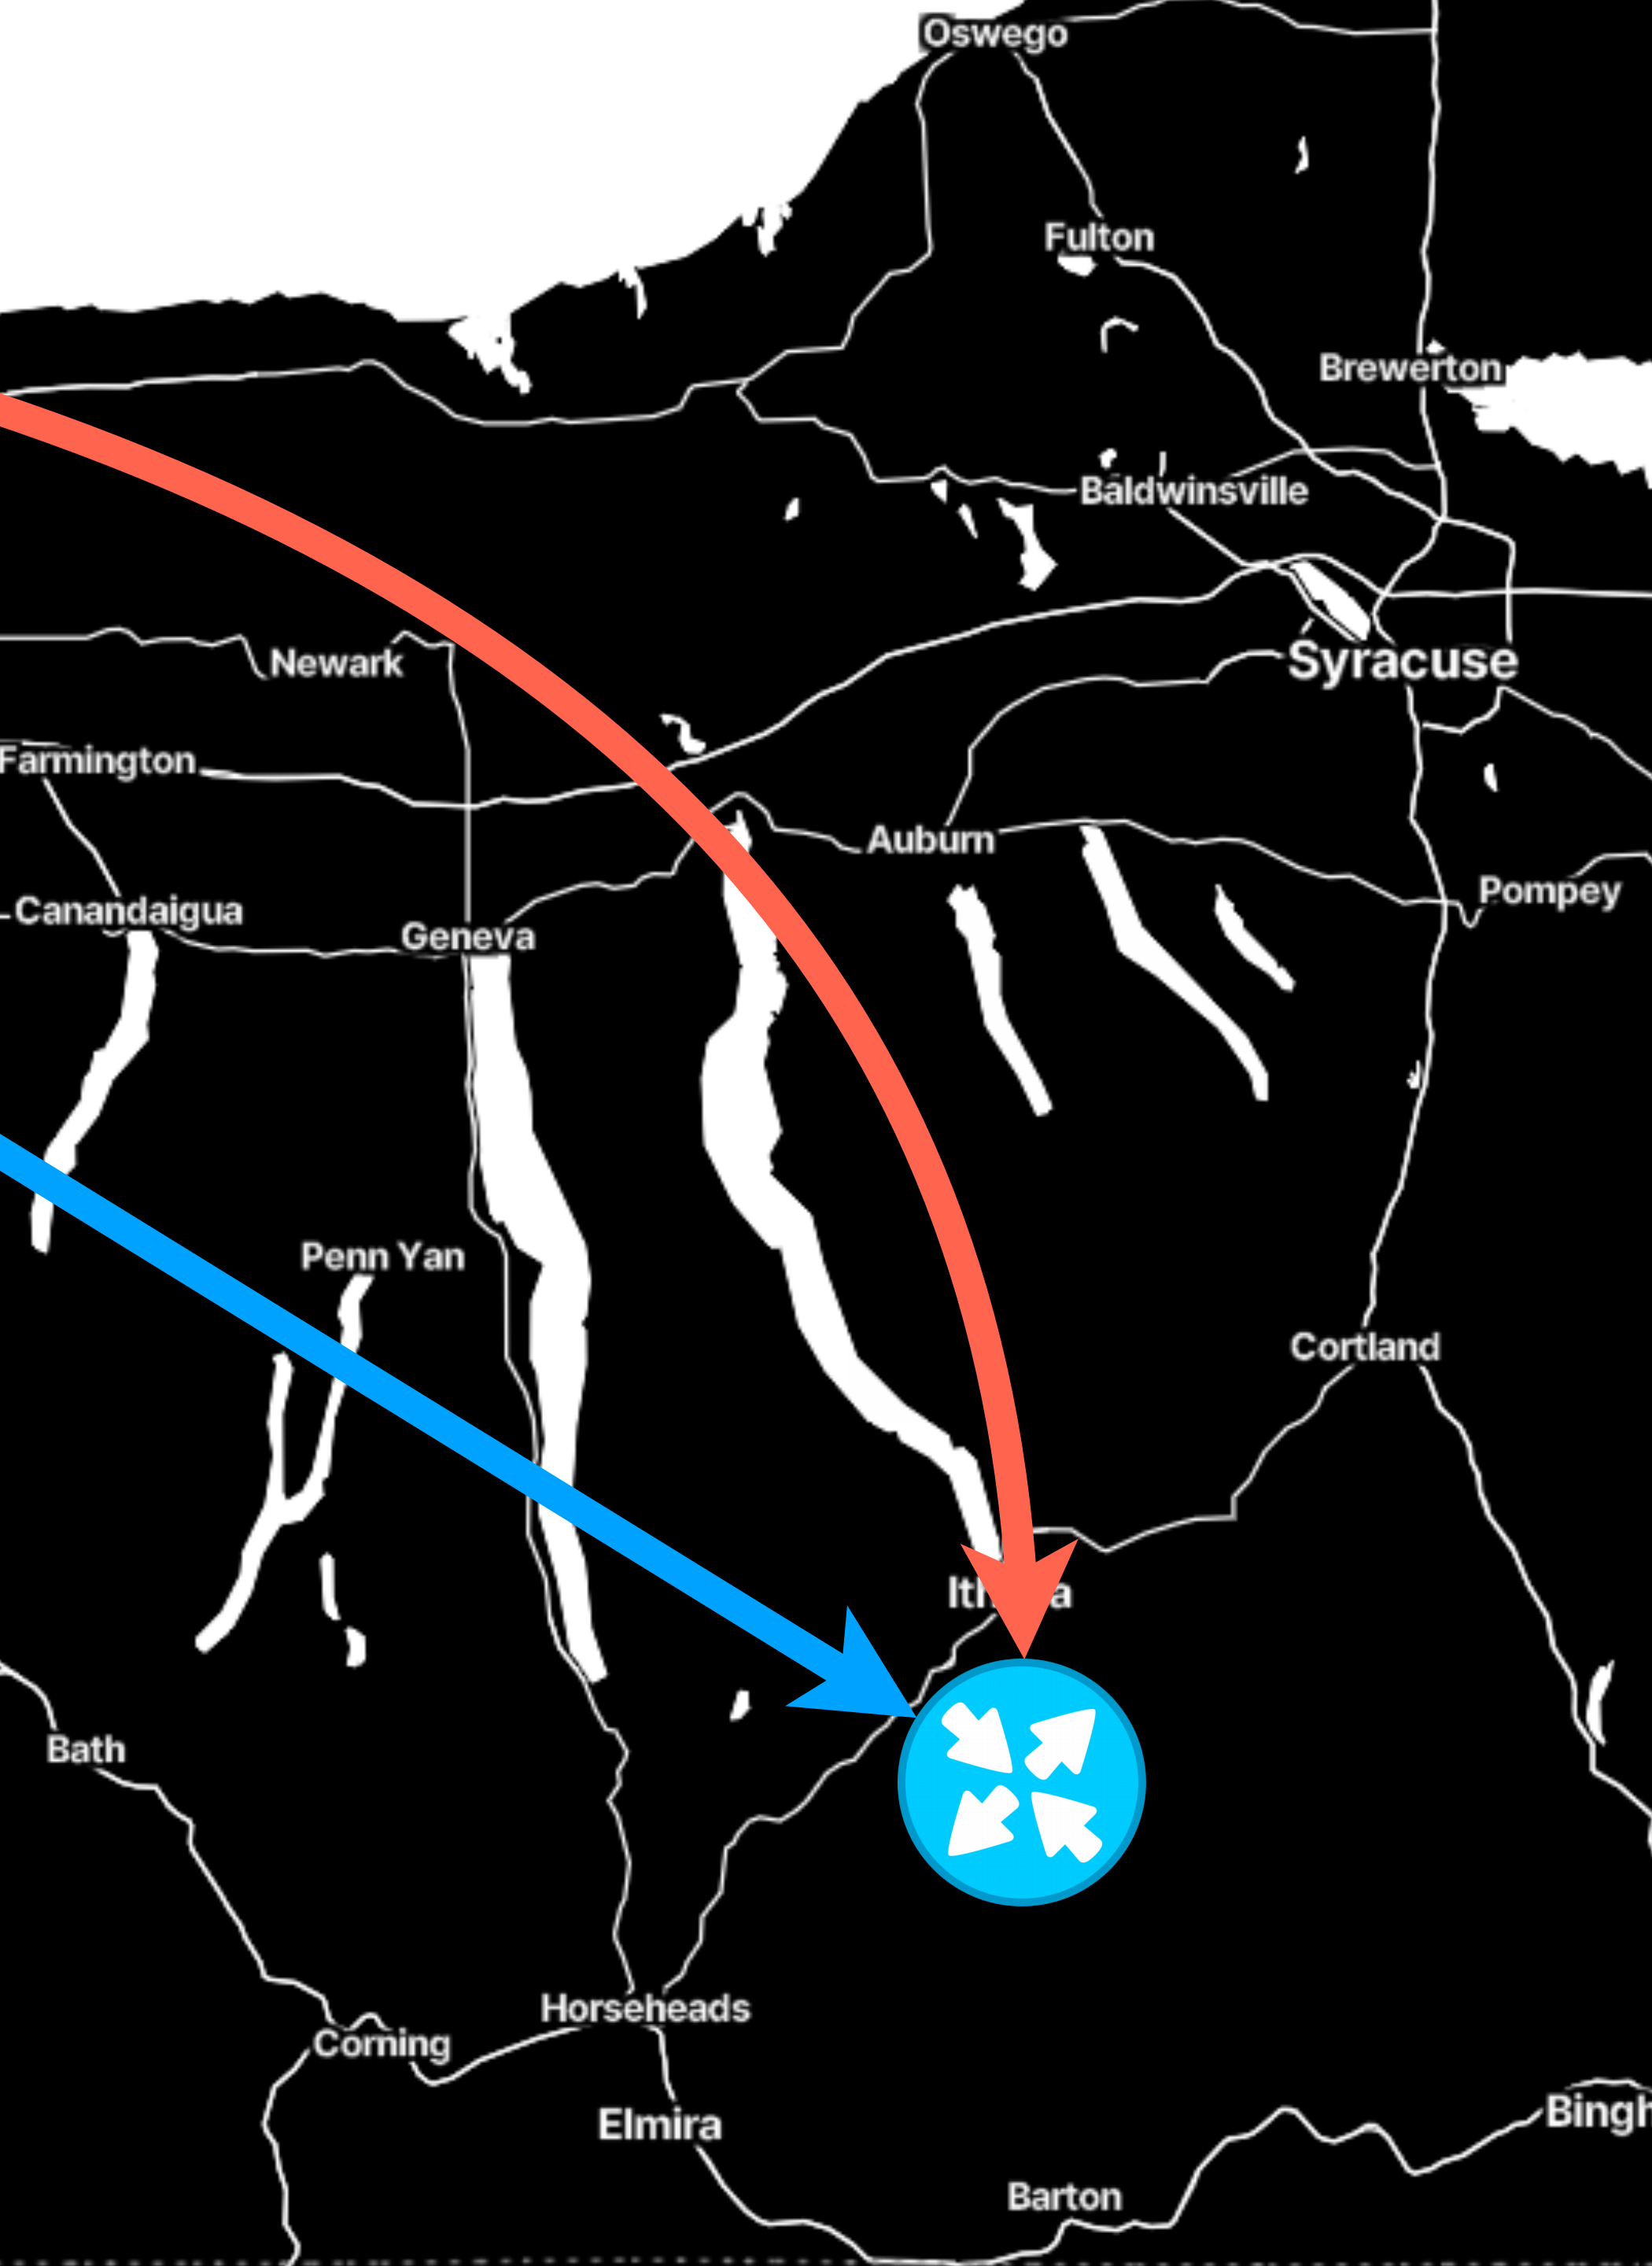
$B_n, \dots, B_1, (R, B)^*$



Baby goal:
interleave **R** and **B**.

A PIFO will suffice.

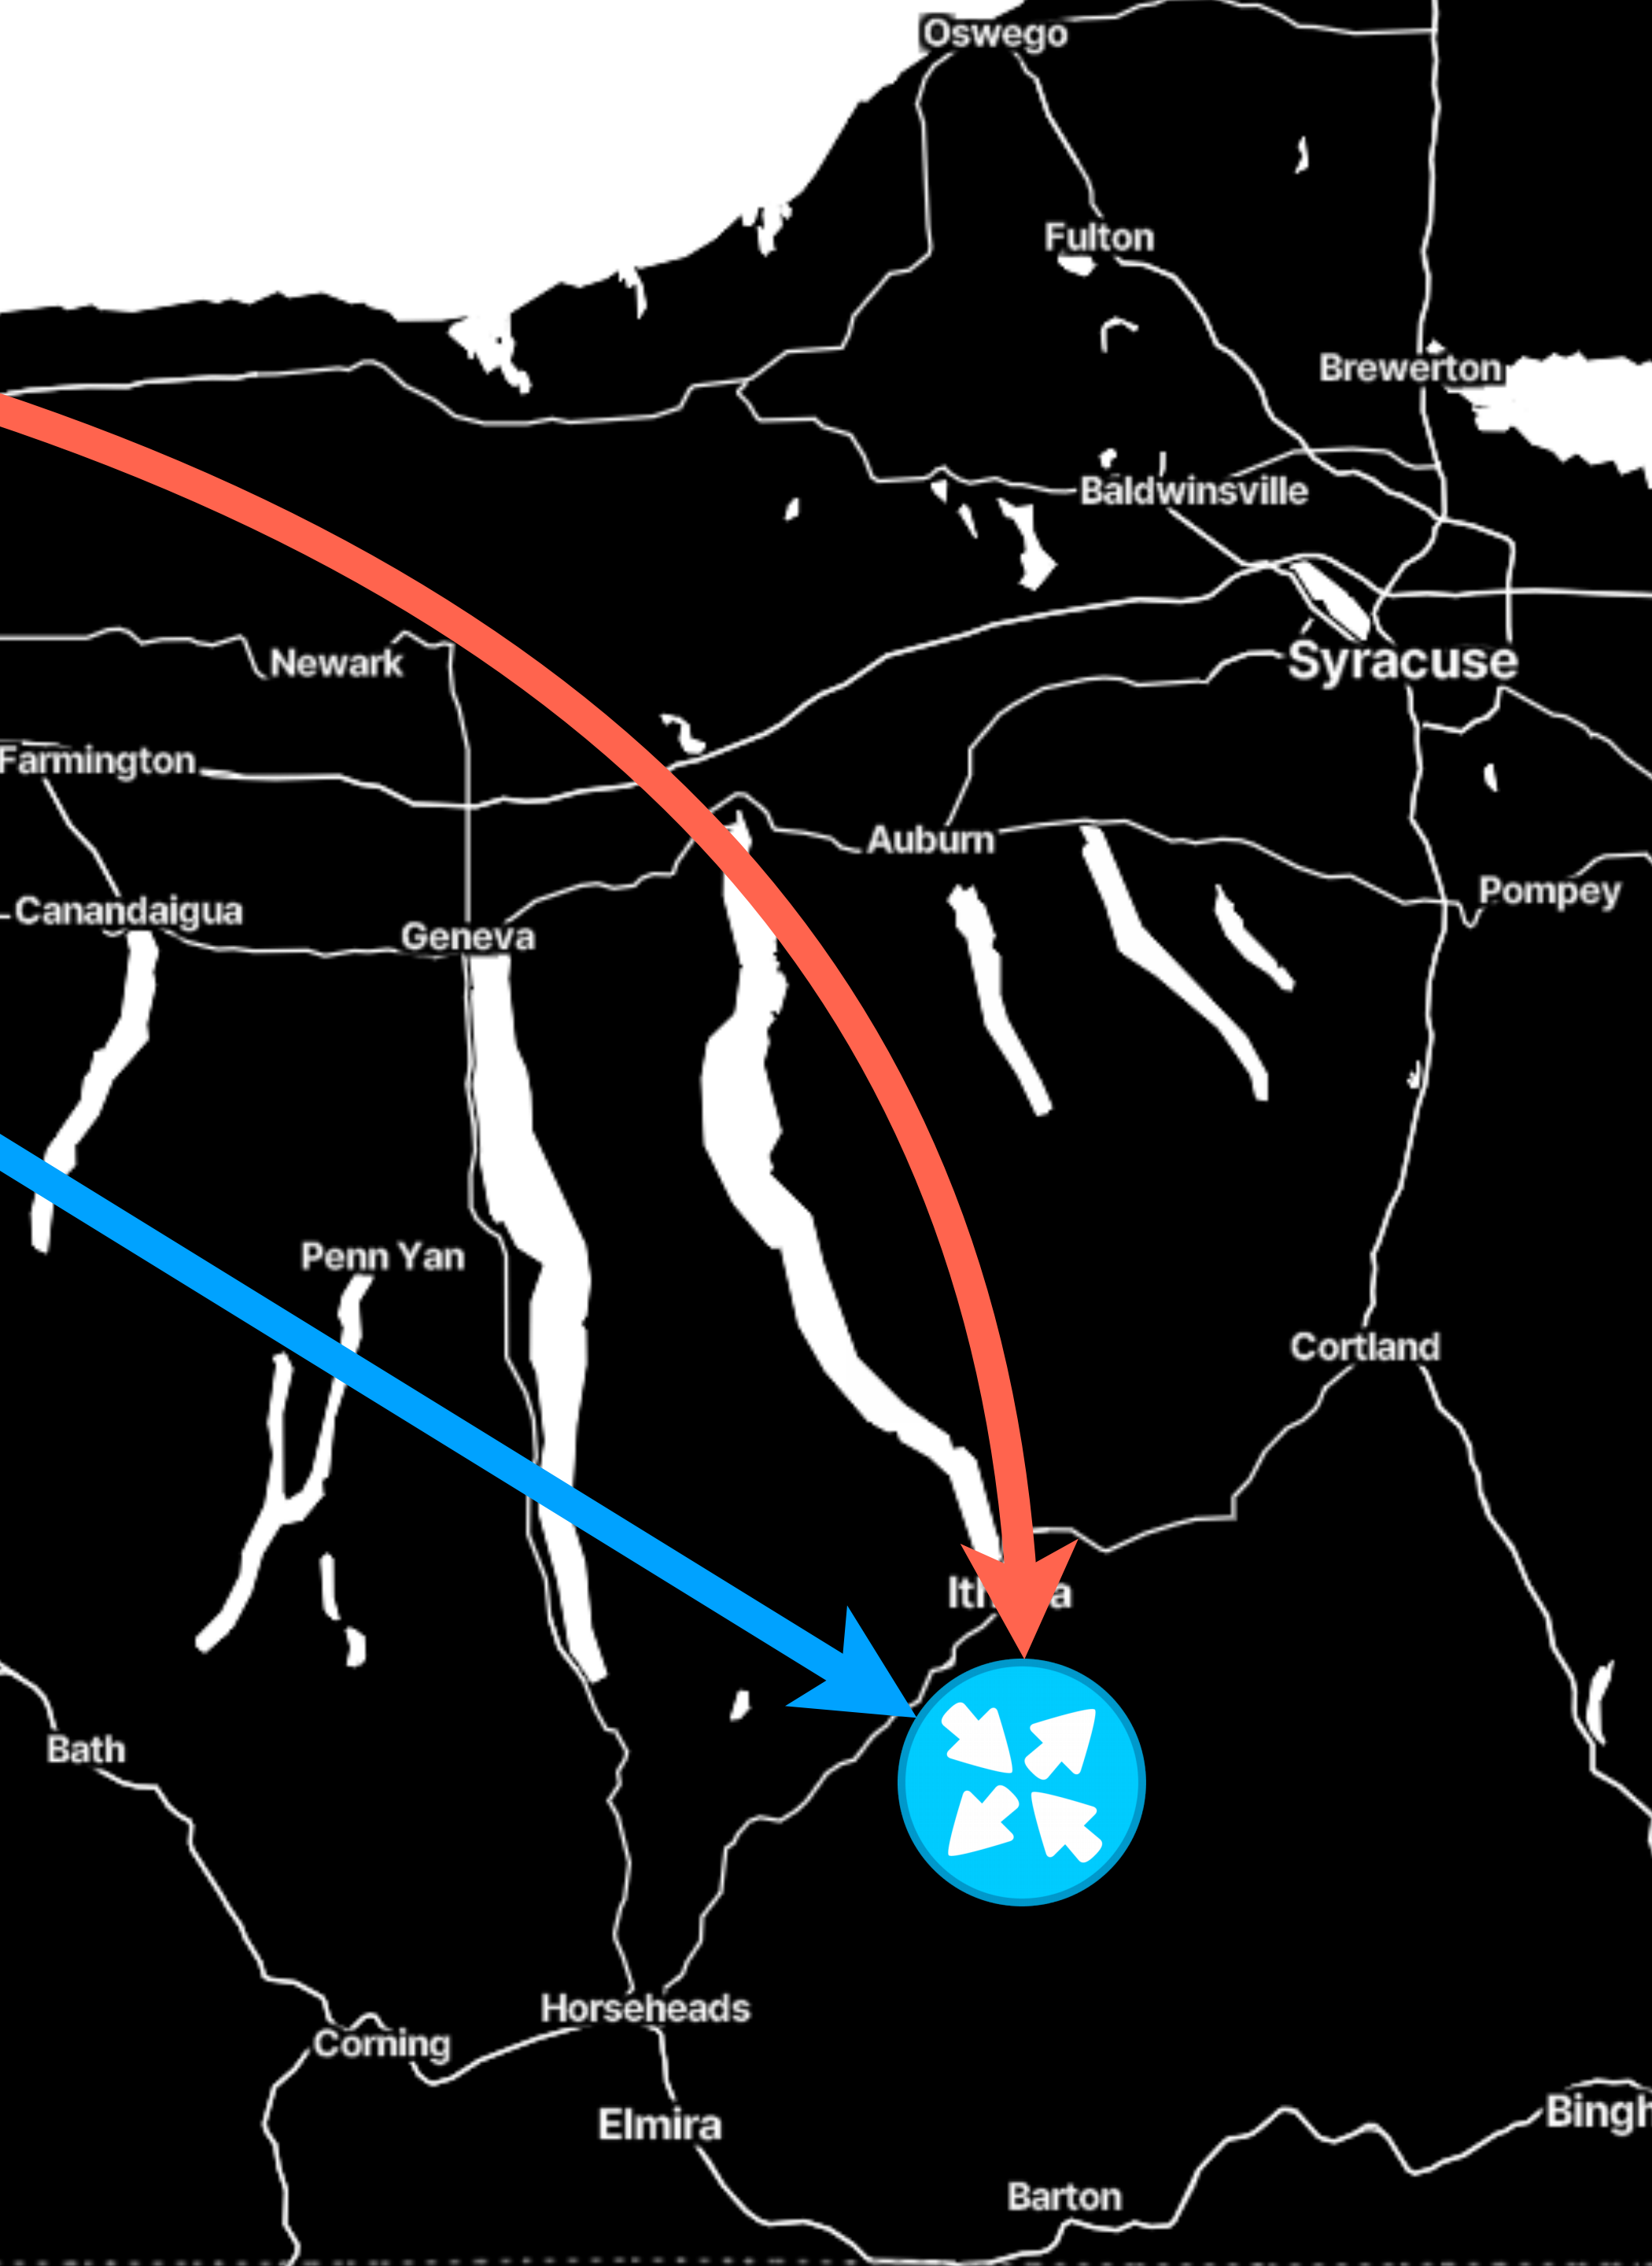
R \rightarrow $B_n, \dots, B_1, (\textcolor{red}{R}, \textcolor{blue}{B})^*$



Baby goal:
interleave **R** and **B**.

A PIFO will suffice.

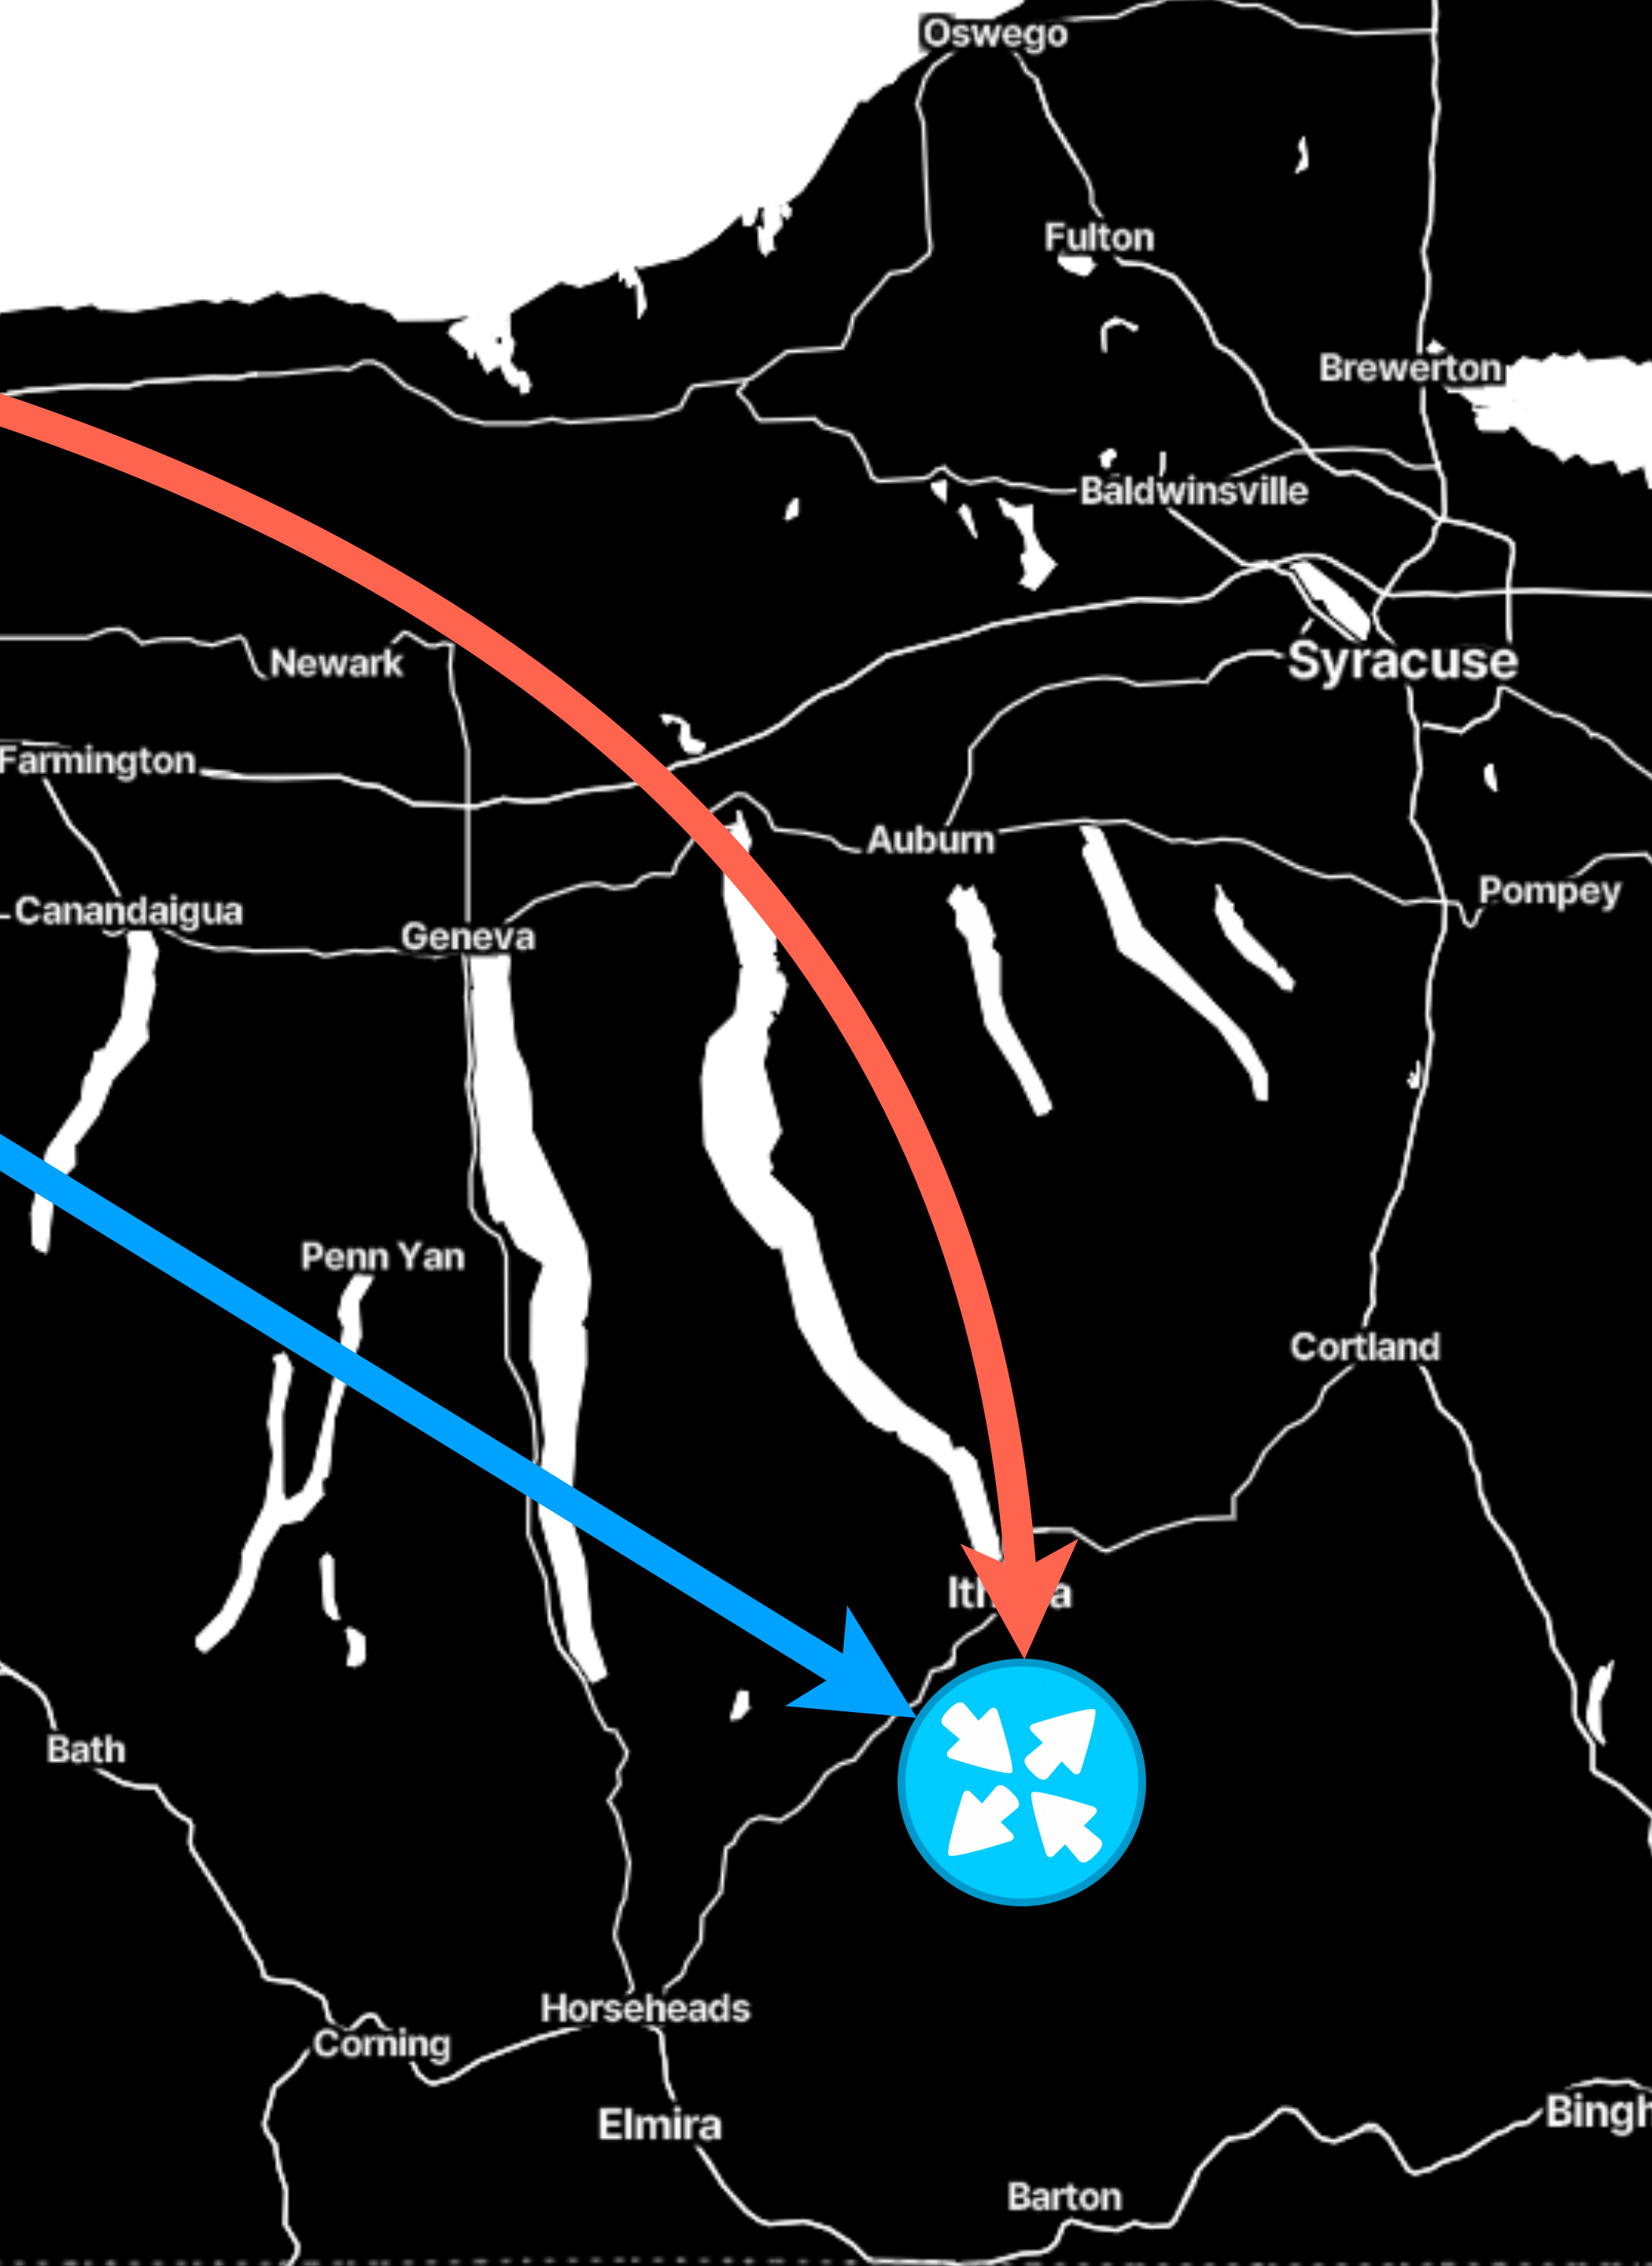
B
↪ $B_n, \dots, B_1, (\textcolor{red}{R}, \textcolor{blue}{B})^*$



Baby goal:
interleave **R** and **B**.

A PIFO will suffice.

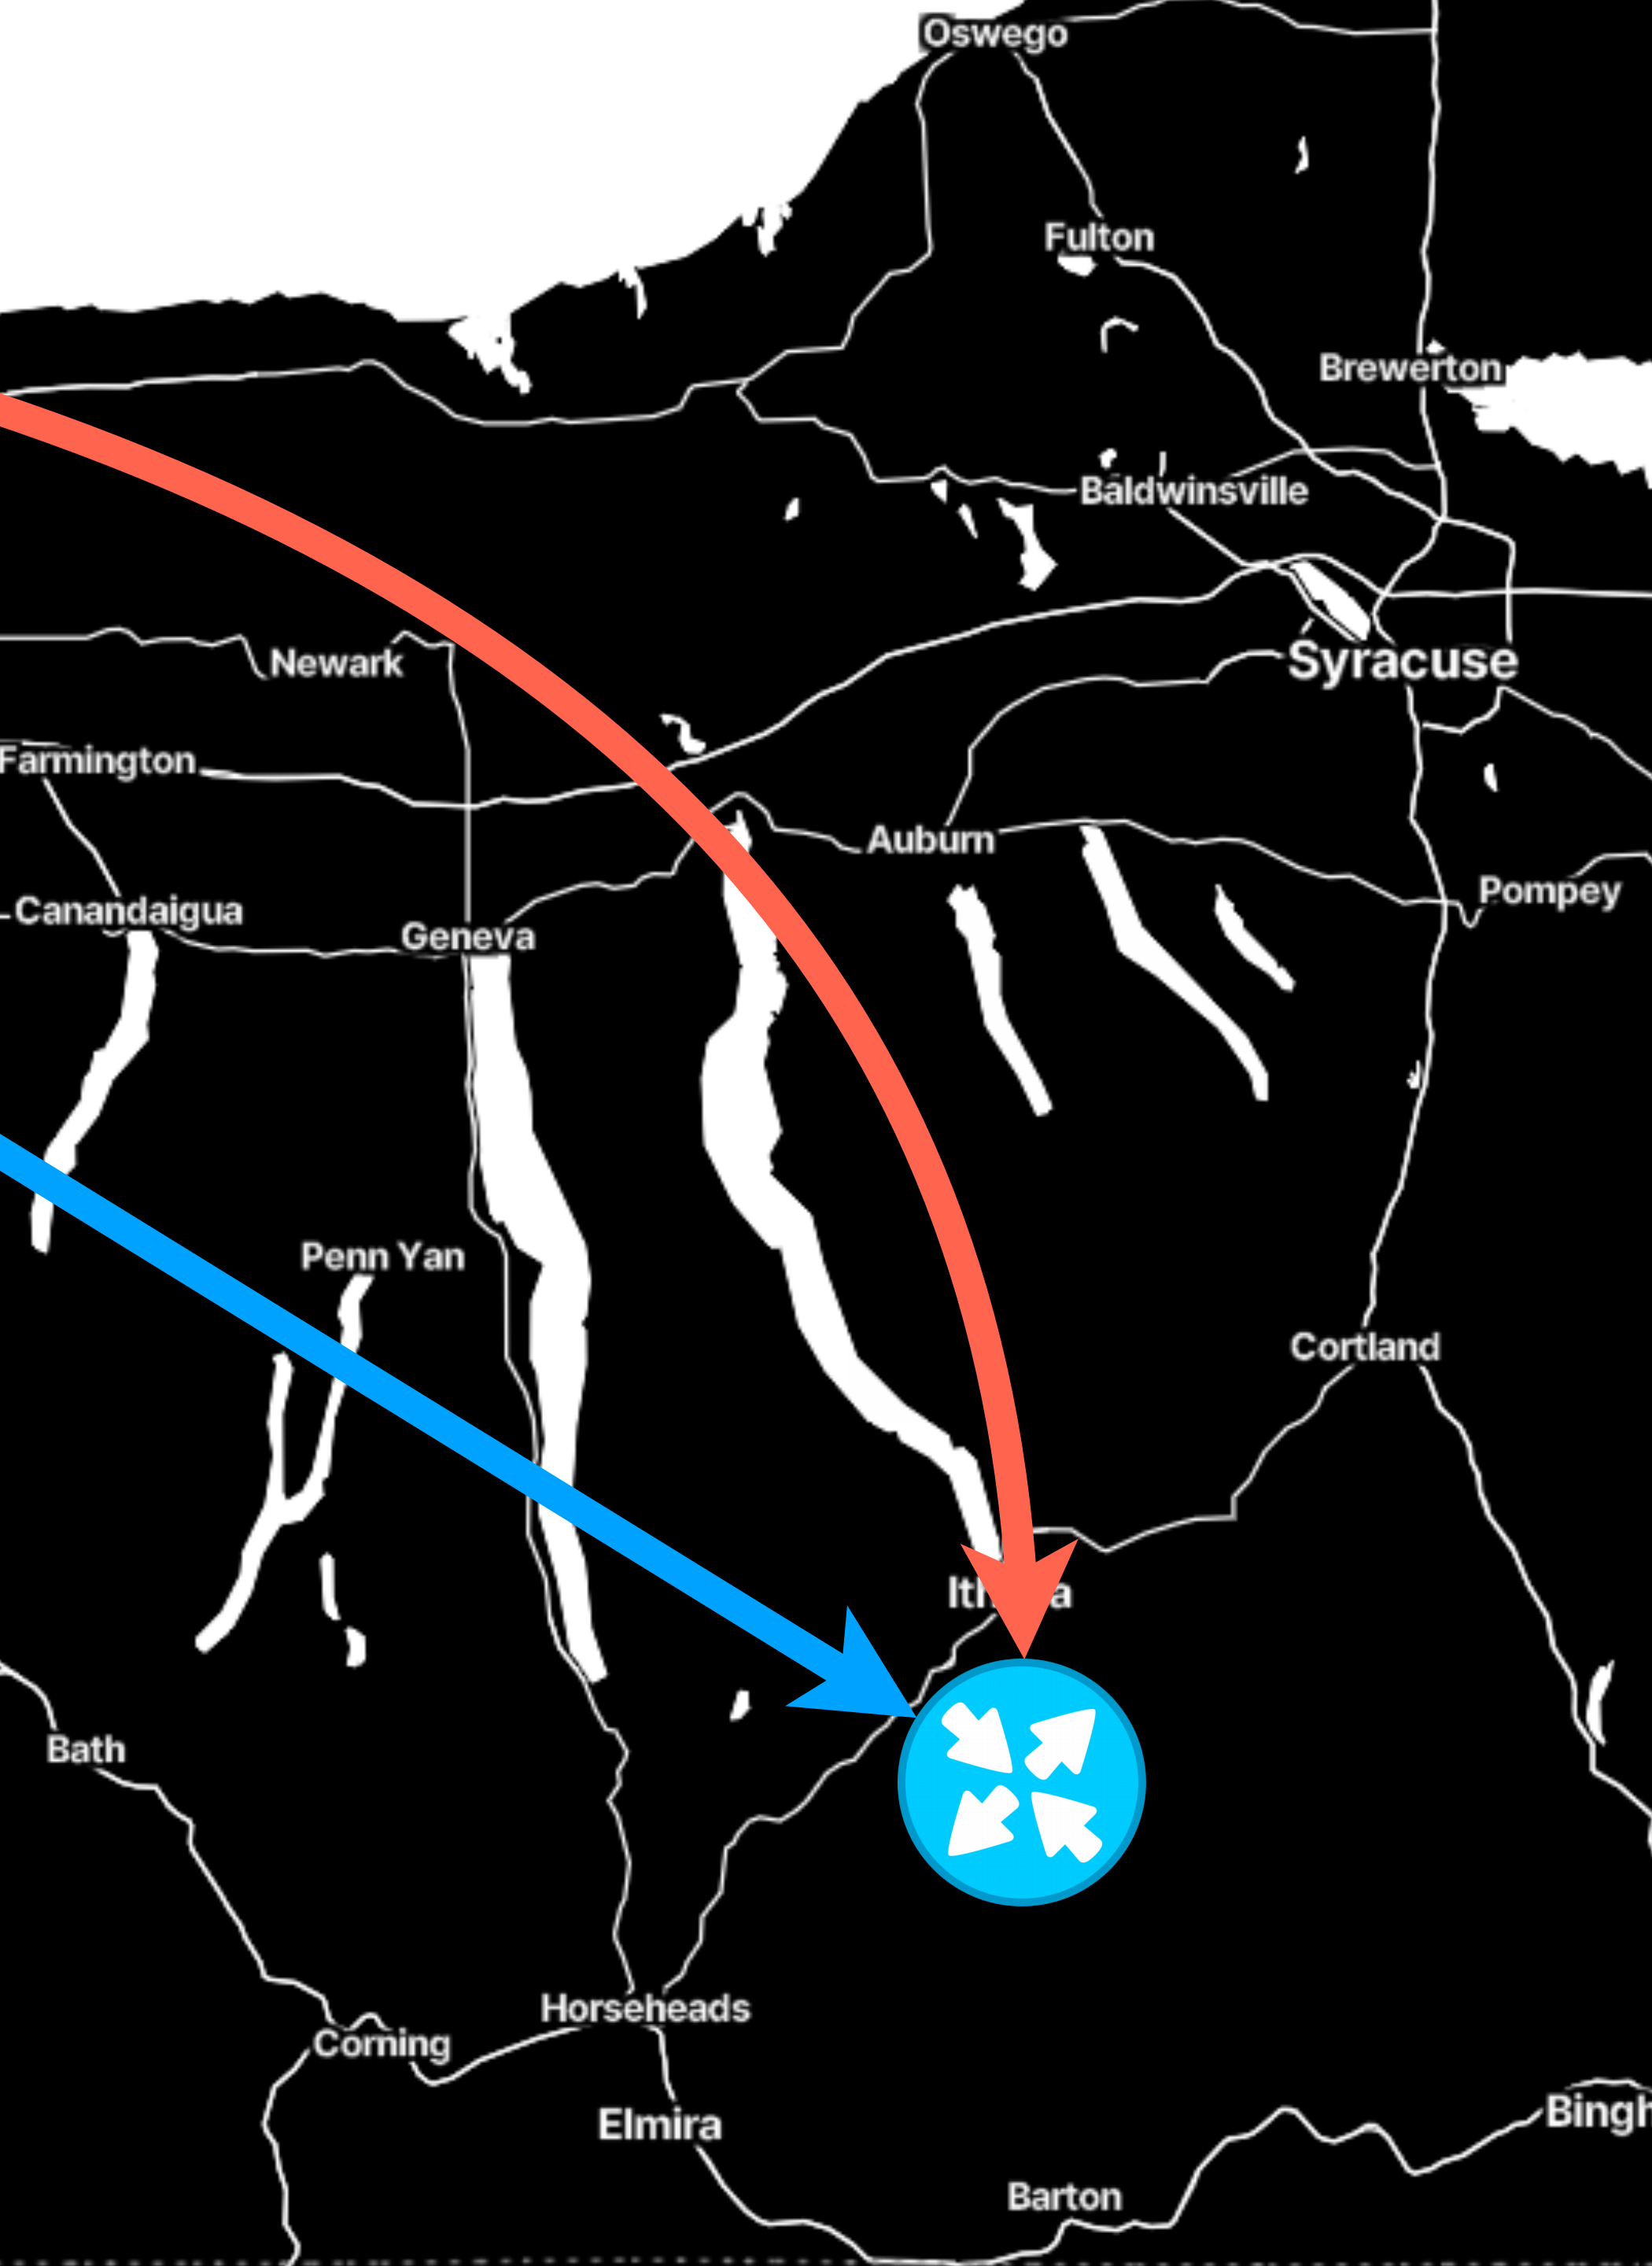
$B_n, \dots, B_1, (R, B)^*$



Baby goal:
interleave **R** and **B**.

A PIFO will suffice.

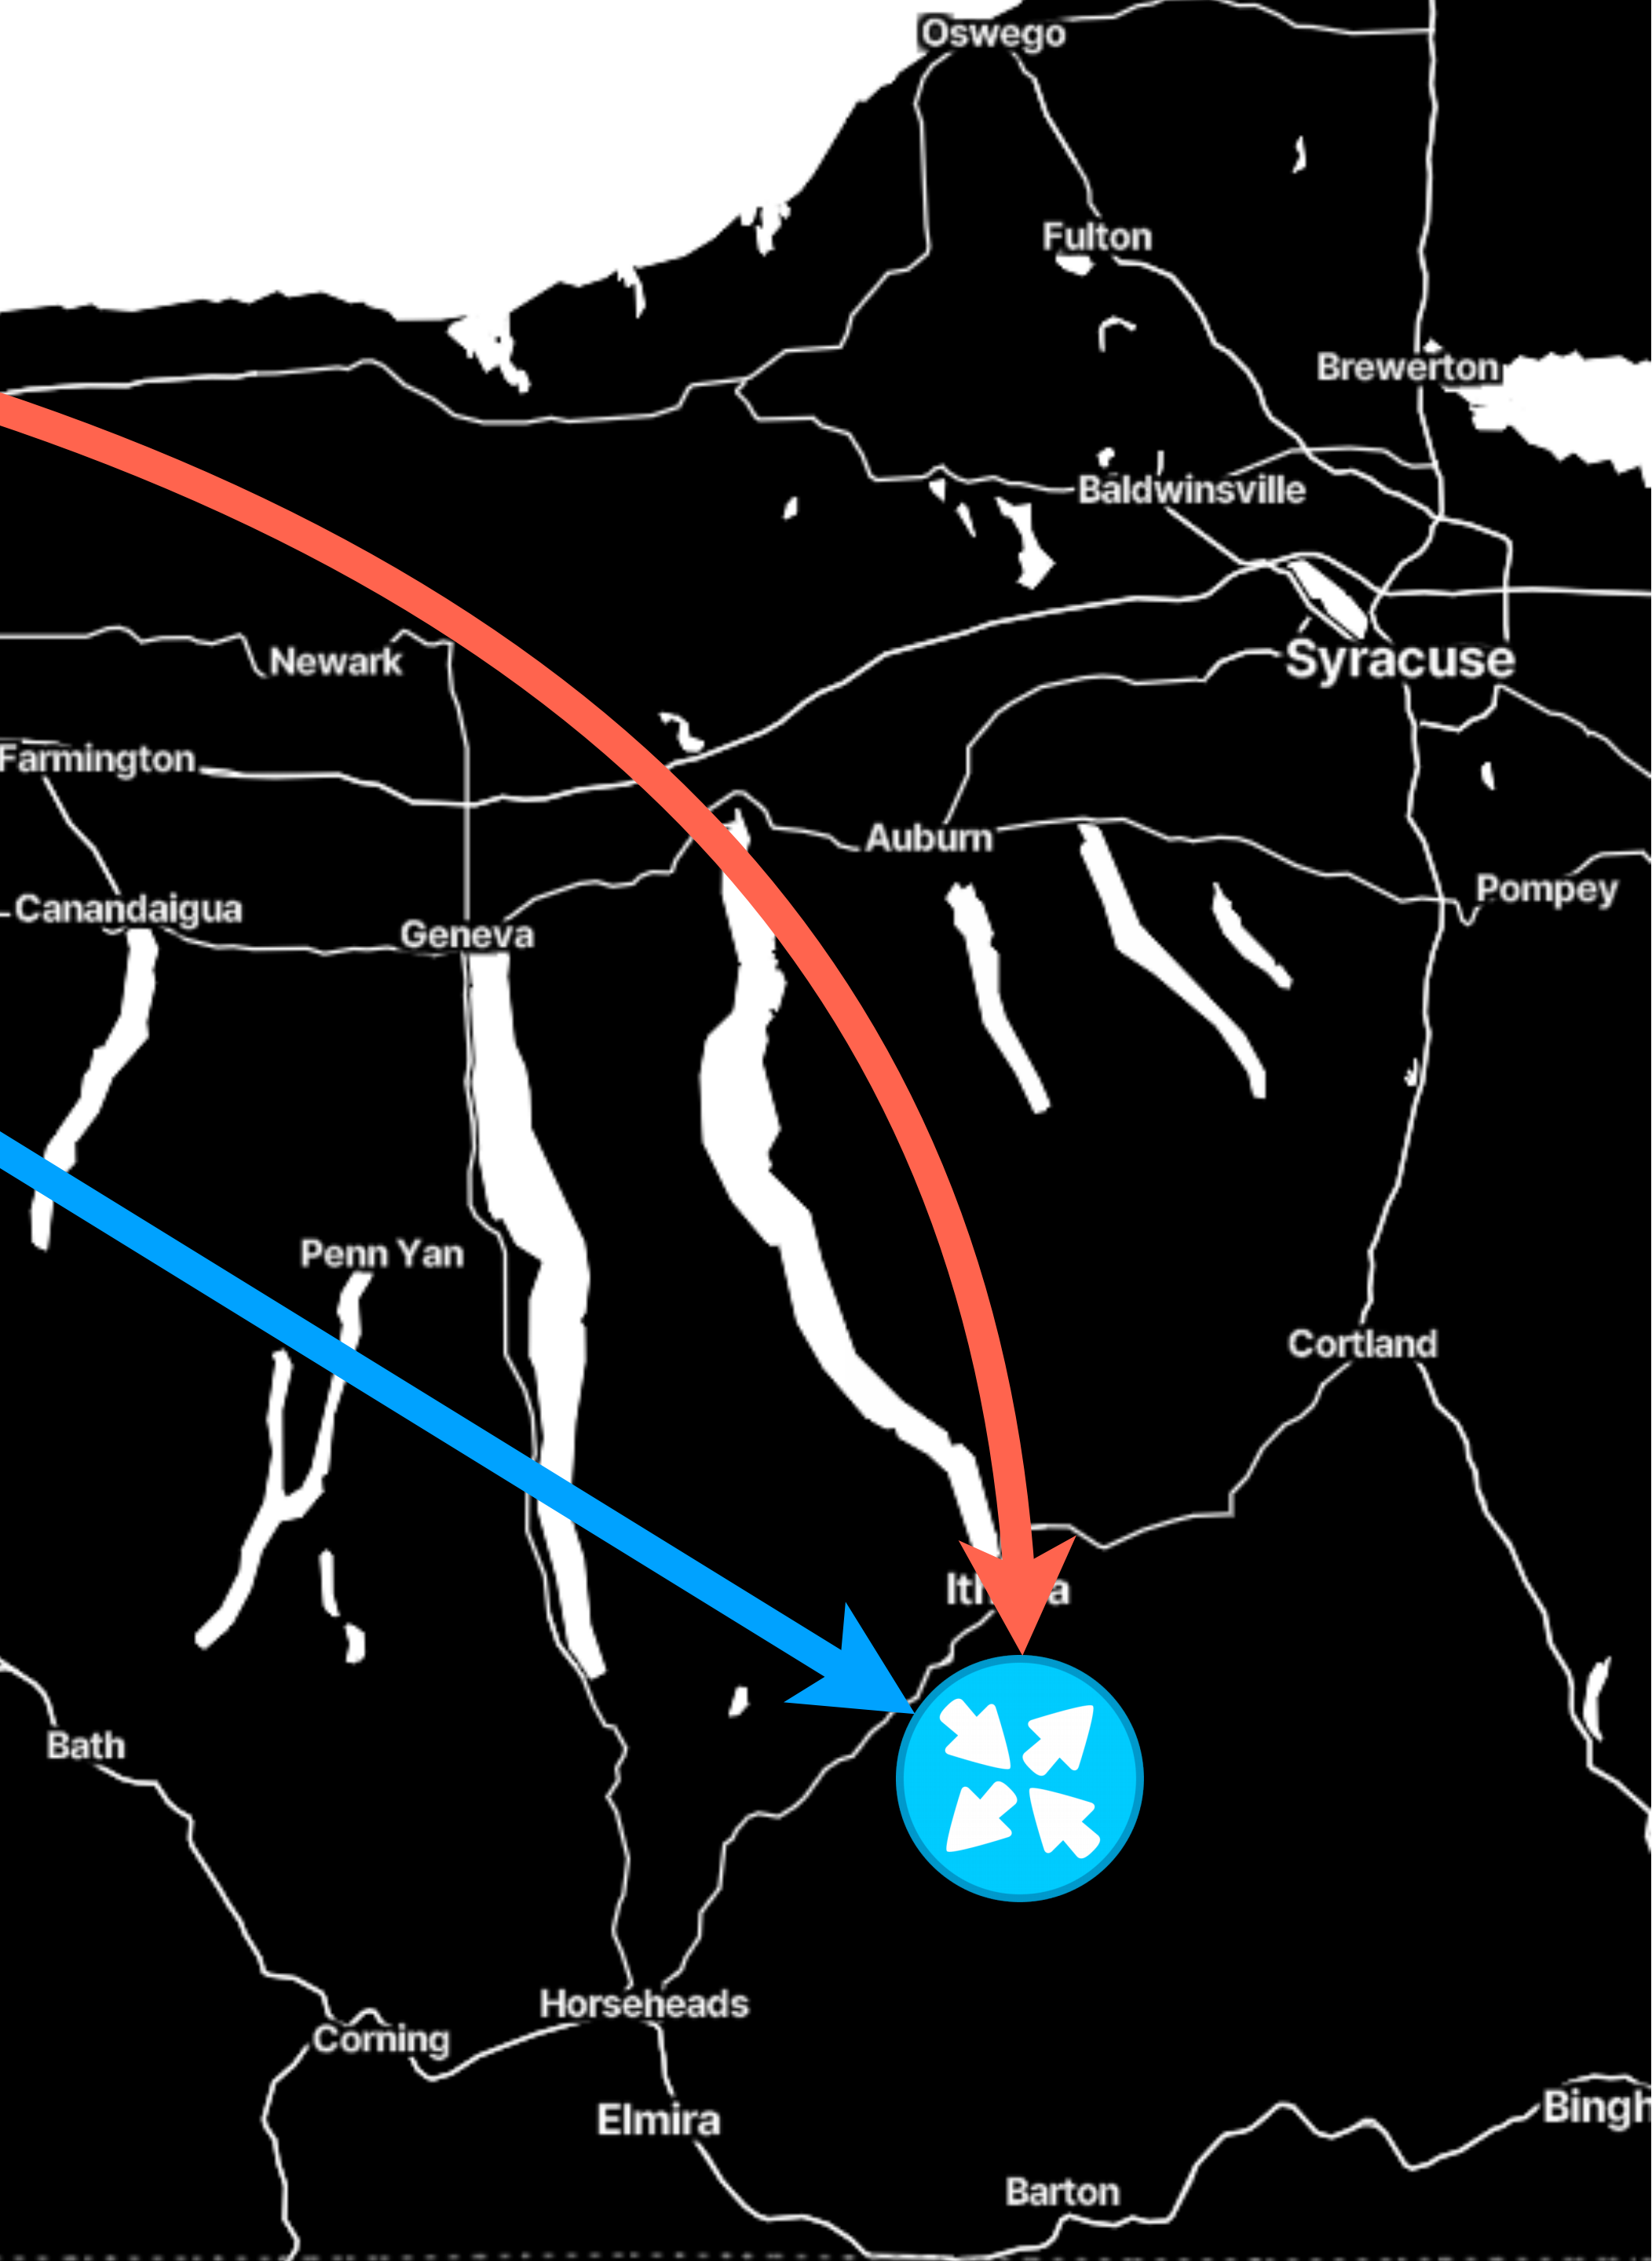
$$\begin{array}{l} B_n, \dots, B_1, (R, B)^* \\ R_n, \dots, R_1, (R, B)^* \end{array}$$



Baby goal:
interleave **R** and **B**.

A PIFO will suffice.

$$\begin{array}{l} B_n, \dots, B_1, (R, B)^* \\ R_n, \dots, R_1, (R, B)^* \\ (R, B)^* \end{array}$$



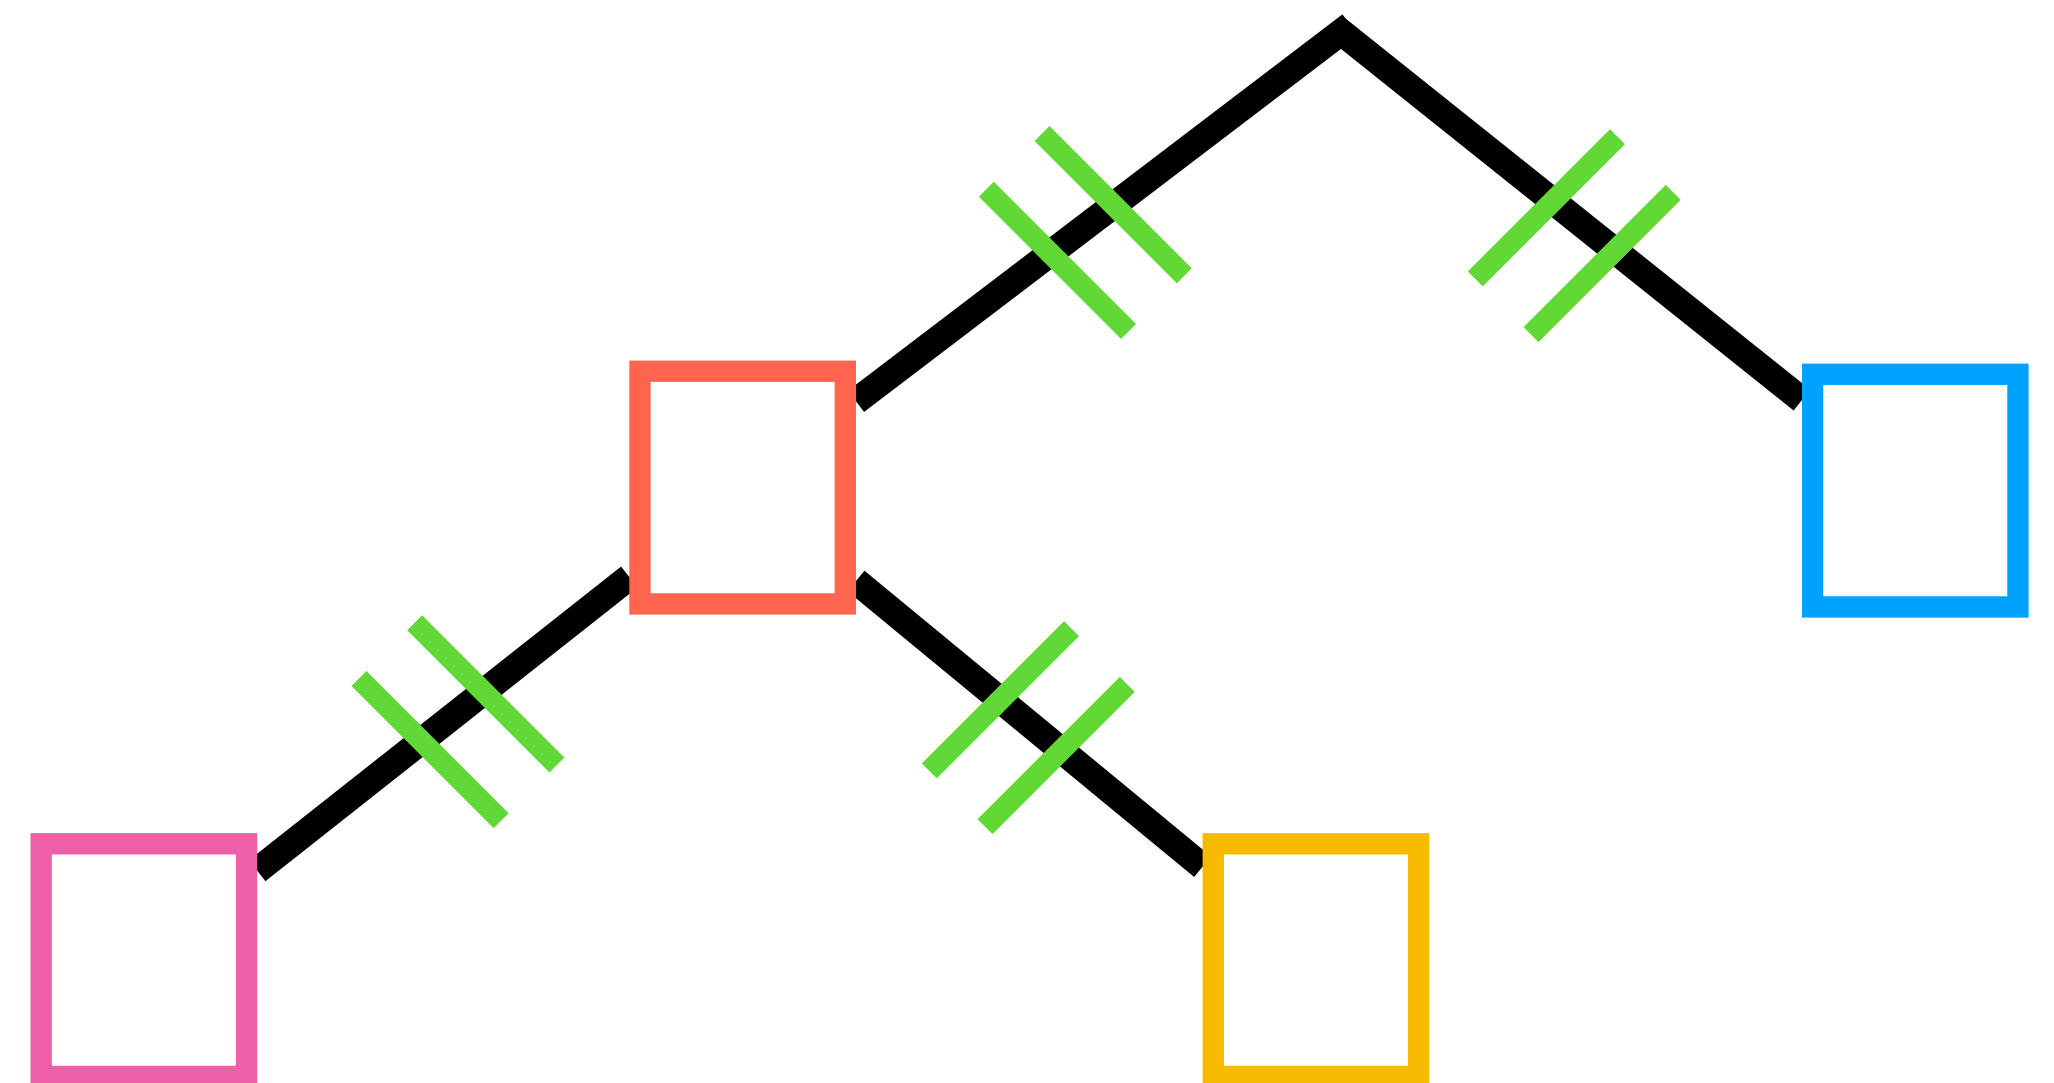




Goal:
interleave **R** and **B**;
interleave **P** and **T**.



Goal:
interleave **R** and **B**;
interleave **P** and **T**.



Goal:

interleave **R** and **B**;

interleave **P** and **T**.

Goal:

interleave **R** and **B**;

interleave **P** and **T**.

B₃, **B**₂, **P**₂, **B**₁, **P**₁

Goal:

interleave **R** and **B**;

interleave **P** and **T**.

T₁ \longrightarrow **B**₃, **B**₂, **P**₂, **B**₁, **P**₁

Goal:

interleave **R** and **B**;

interleave **P** and **T**.

T₁ → **B**₃, **B**₂, **P**₂, **B**₁, **P**₁

B₃, **B**₂, **P**₂, **T**₁, **B**₁, **P**₁

Goal:

interleave **R** and **B**;

interleave **P** and **T**.

$T_1 \longrightarrow B_3, B_2, P_2, B_1, P_1$

$B_3, B_2, P_2, T_1, B_1, P_1$





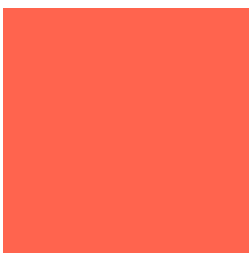

$B_3, B_2, P_2, B_1, T_1, P_1$

Goal:

interleave **R** and **B**;

interleave **P** and **T**.

$T_1 \longrightarrow B_3, B_2, P_2, B_1, P_1$

B_3	B_2			B_1	
B_3	B_2		B_1		

Goal:
interleave **R** and **B**;
interleave **P** and **T**.

B₃, **B**₂, **P**₂, **T**₁, **B**₁, **P**₁
B₃, **B**₂, **P**₂, **B**₁, **T**₁, **P**₁

T₁ → **B**₃, **B**₂, **P**₂, **B**₁, **P**₁

Goal:
interleave **R** and **B**;
interleave **P** and **T**.

~~B₃, B₂, P₂, T₁, B₁, P₁
B₃, B₂, P₂, B₁, T₁, P₁~~

T₁ → B₃, B₂, P₂, B₁, P₁

Goal:
interleave **R** and **B**;
interleave **P** and **T**.

~~$B_3, B_2, P_2, T_1, B_1, P_1$
 $B_3, B_2, P_2, B_1, T_1, P_1$~~

$T_1 \longrightarrow B_3, B_2, P_2, B_1, P_1$

$B_3, P_2, B_2, T_1, B_1, P_1$

Goal:
interleave **R** and **B**;
interleave **P** and **T**.

~~$B_3, B_2, P_2, T_1, B_1, P_1$
 $B_3, B_2, P_2, B_1, T_1, P_1$~~

$T_1 \longrightarrow B_3, B_2, P_2, B_1, P_1$

$B_3, \blacksquare, B_2, \blacksquare, B_1, \blacksquare$

Goal:
interleave **R** and **B**;
interleave **P** and **T**.

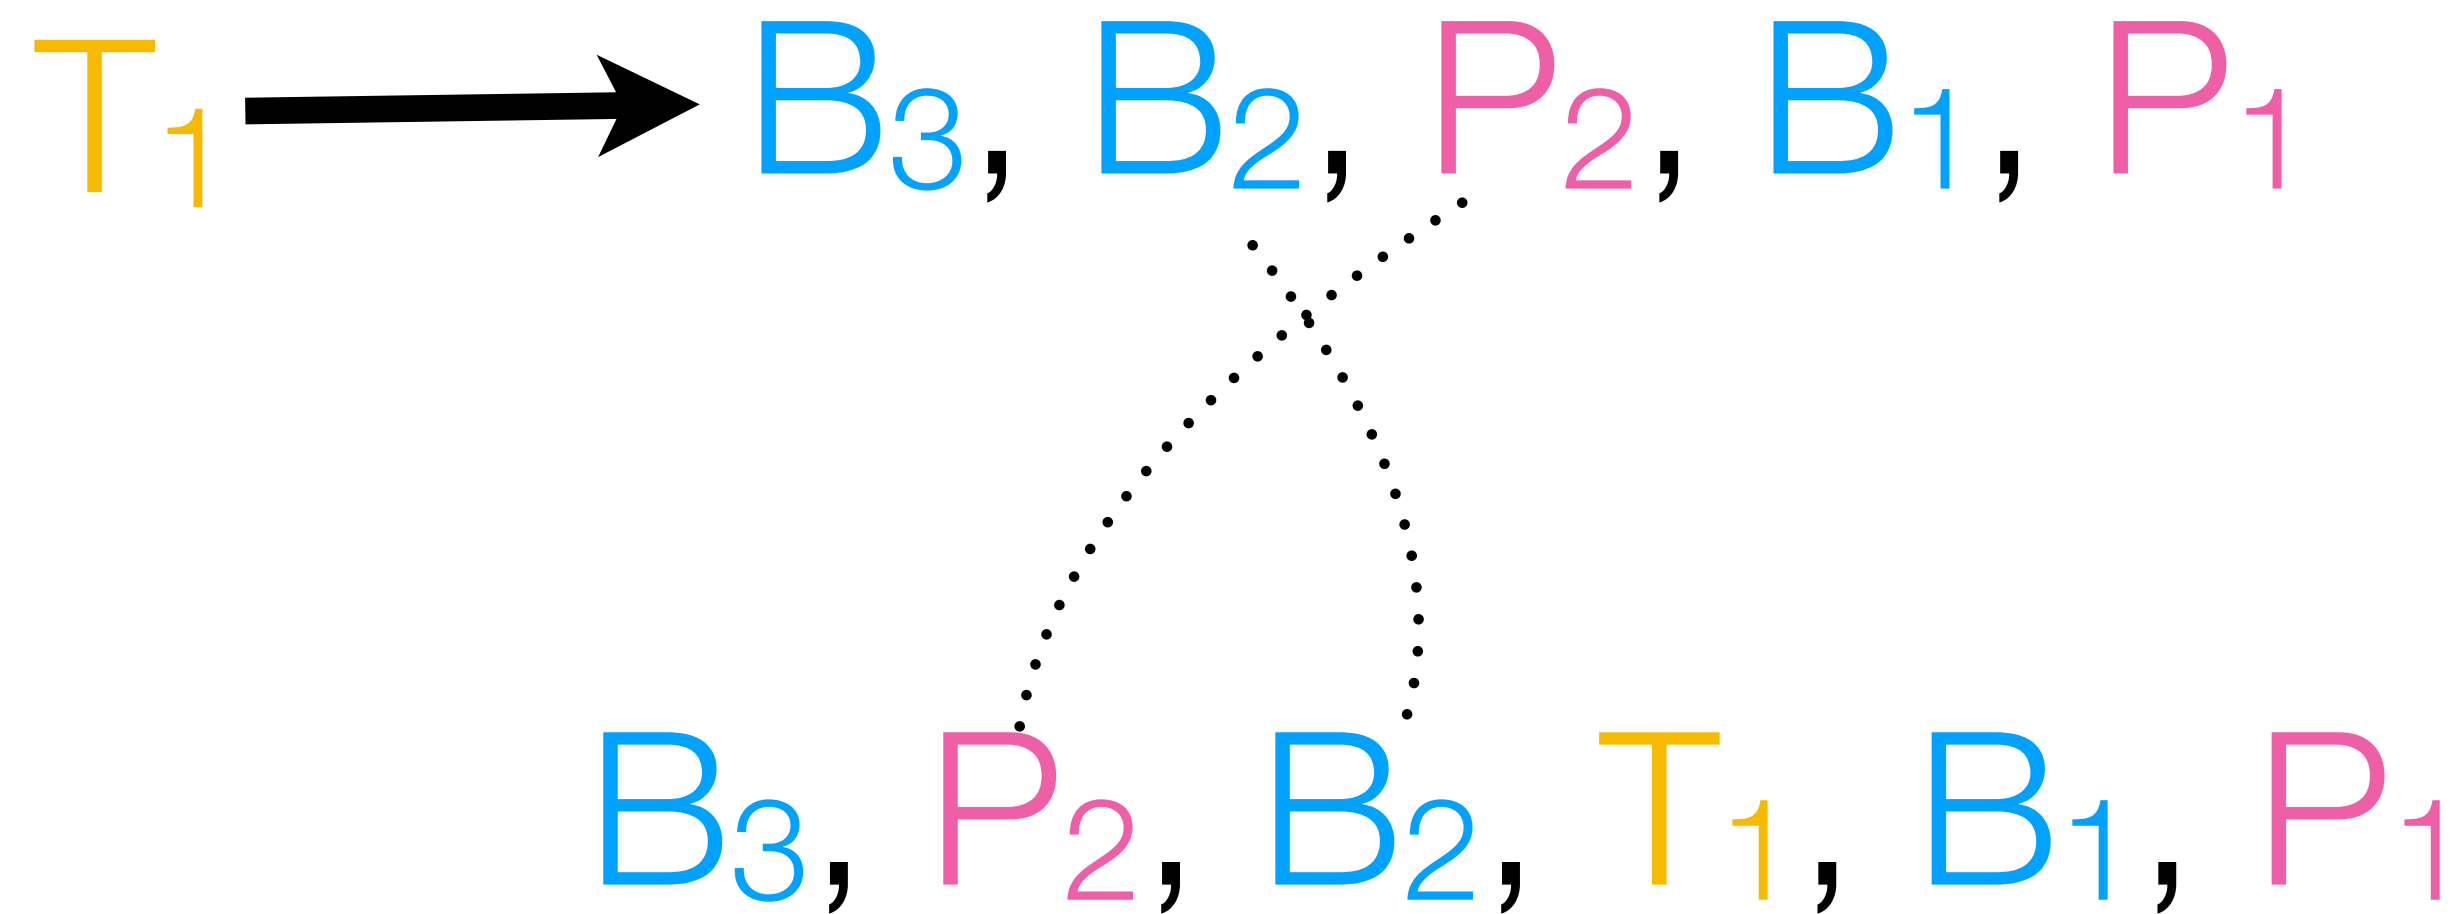
~~$B_3, B_2, P_2, T_1, B_1, P_1$
 $B_3, B_2, P_2, B_1, T_1, P_1$~~

$T_1 \longrightarrow B_3, B_2, P_2, B_1, P_1$

$B_3, P_2, B_2, T_1, B_1, P_1$

Goal:
interleave **R** and **B**;
interleave **P** and **T**.

~~$B_3, B_2, P_2, T_1, B_1, P_1$
 $B_3, B_2, P_2, B_1, T_1, P_1$~~

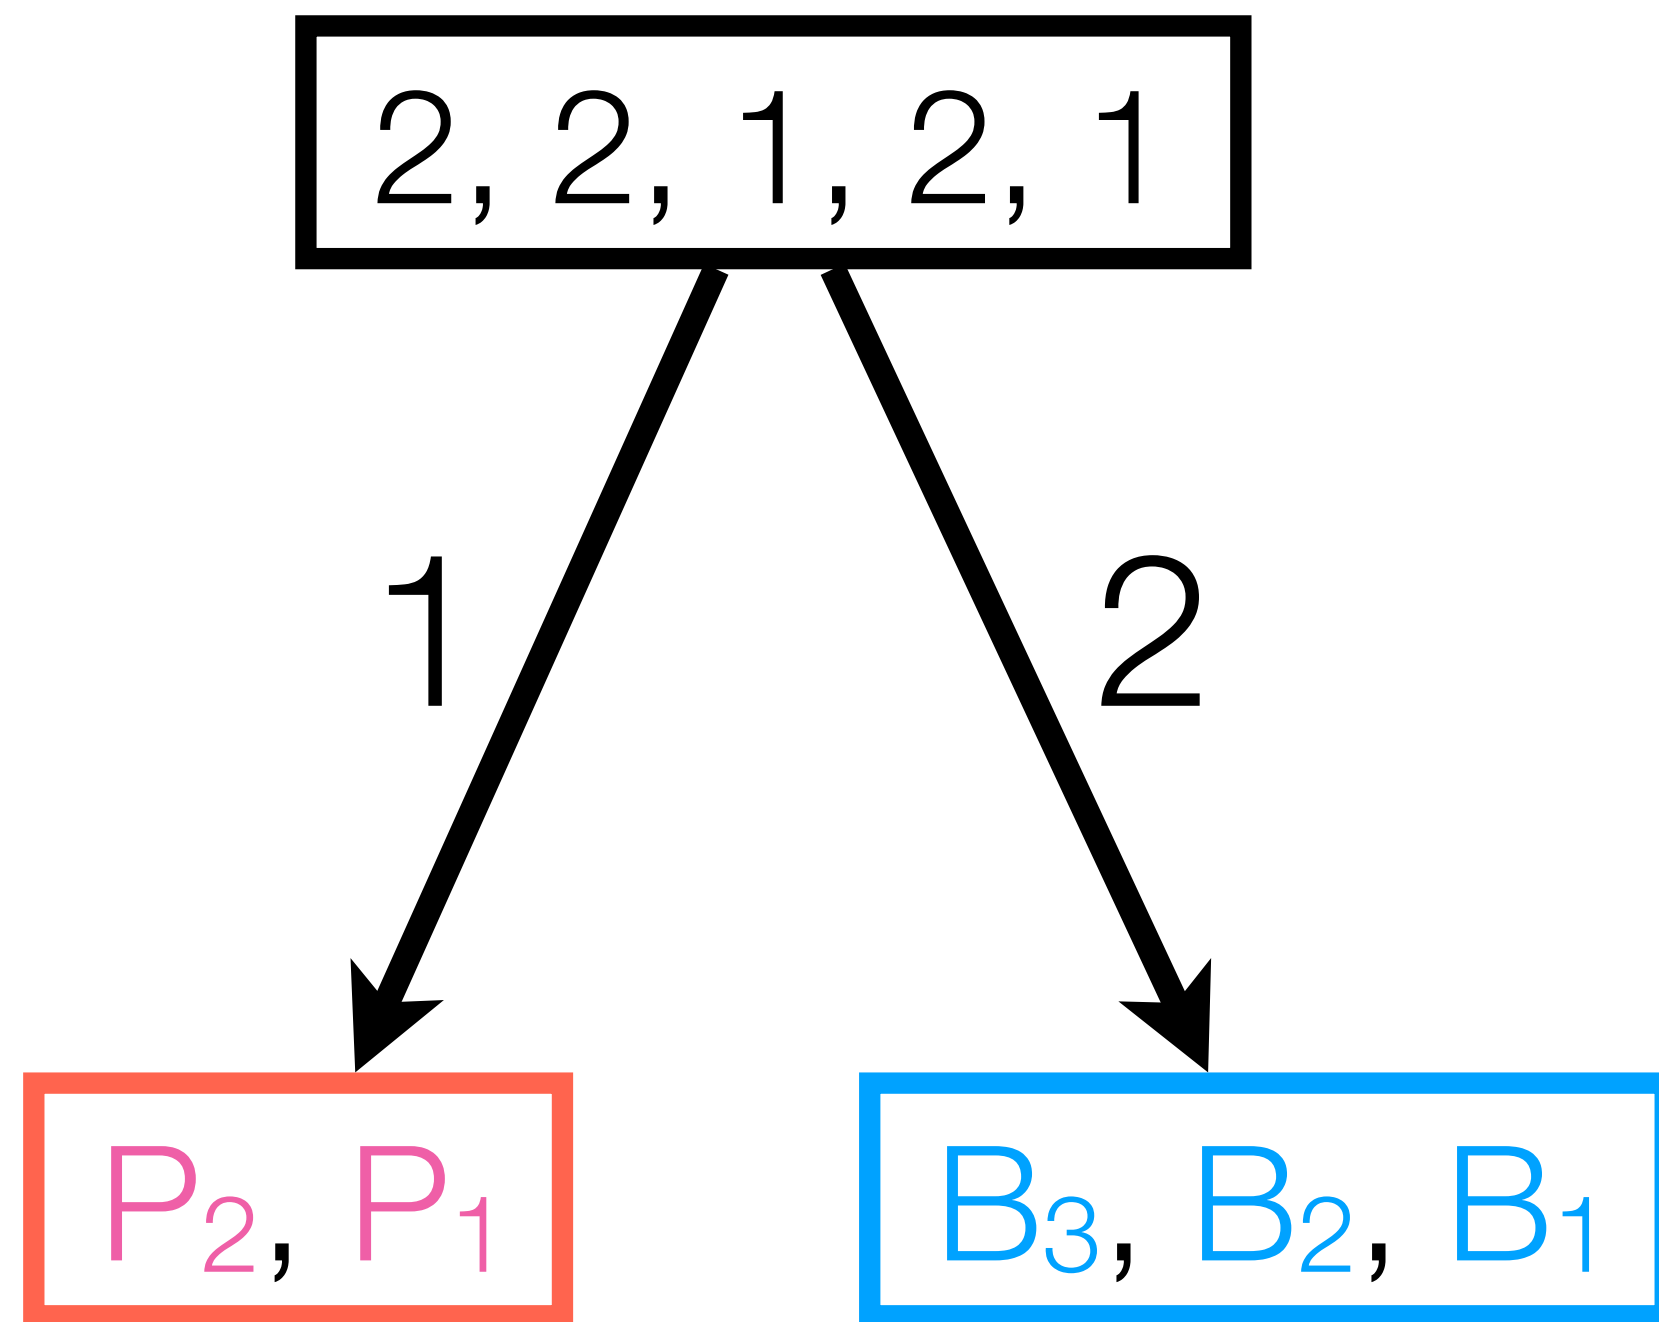


Enqueueing a packet can
require the reordering of
buffered packets.

No PIFO can do this.

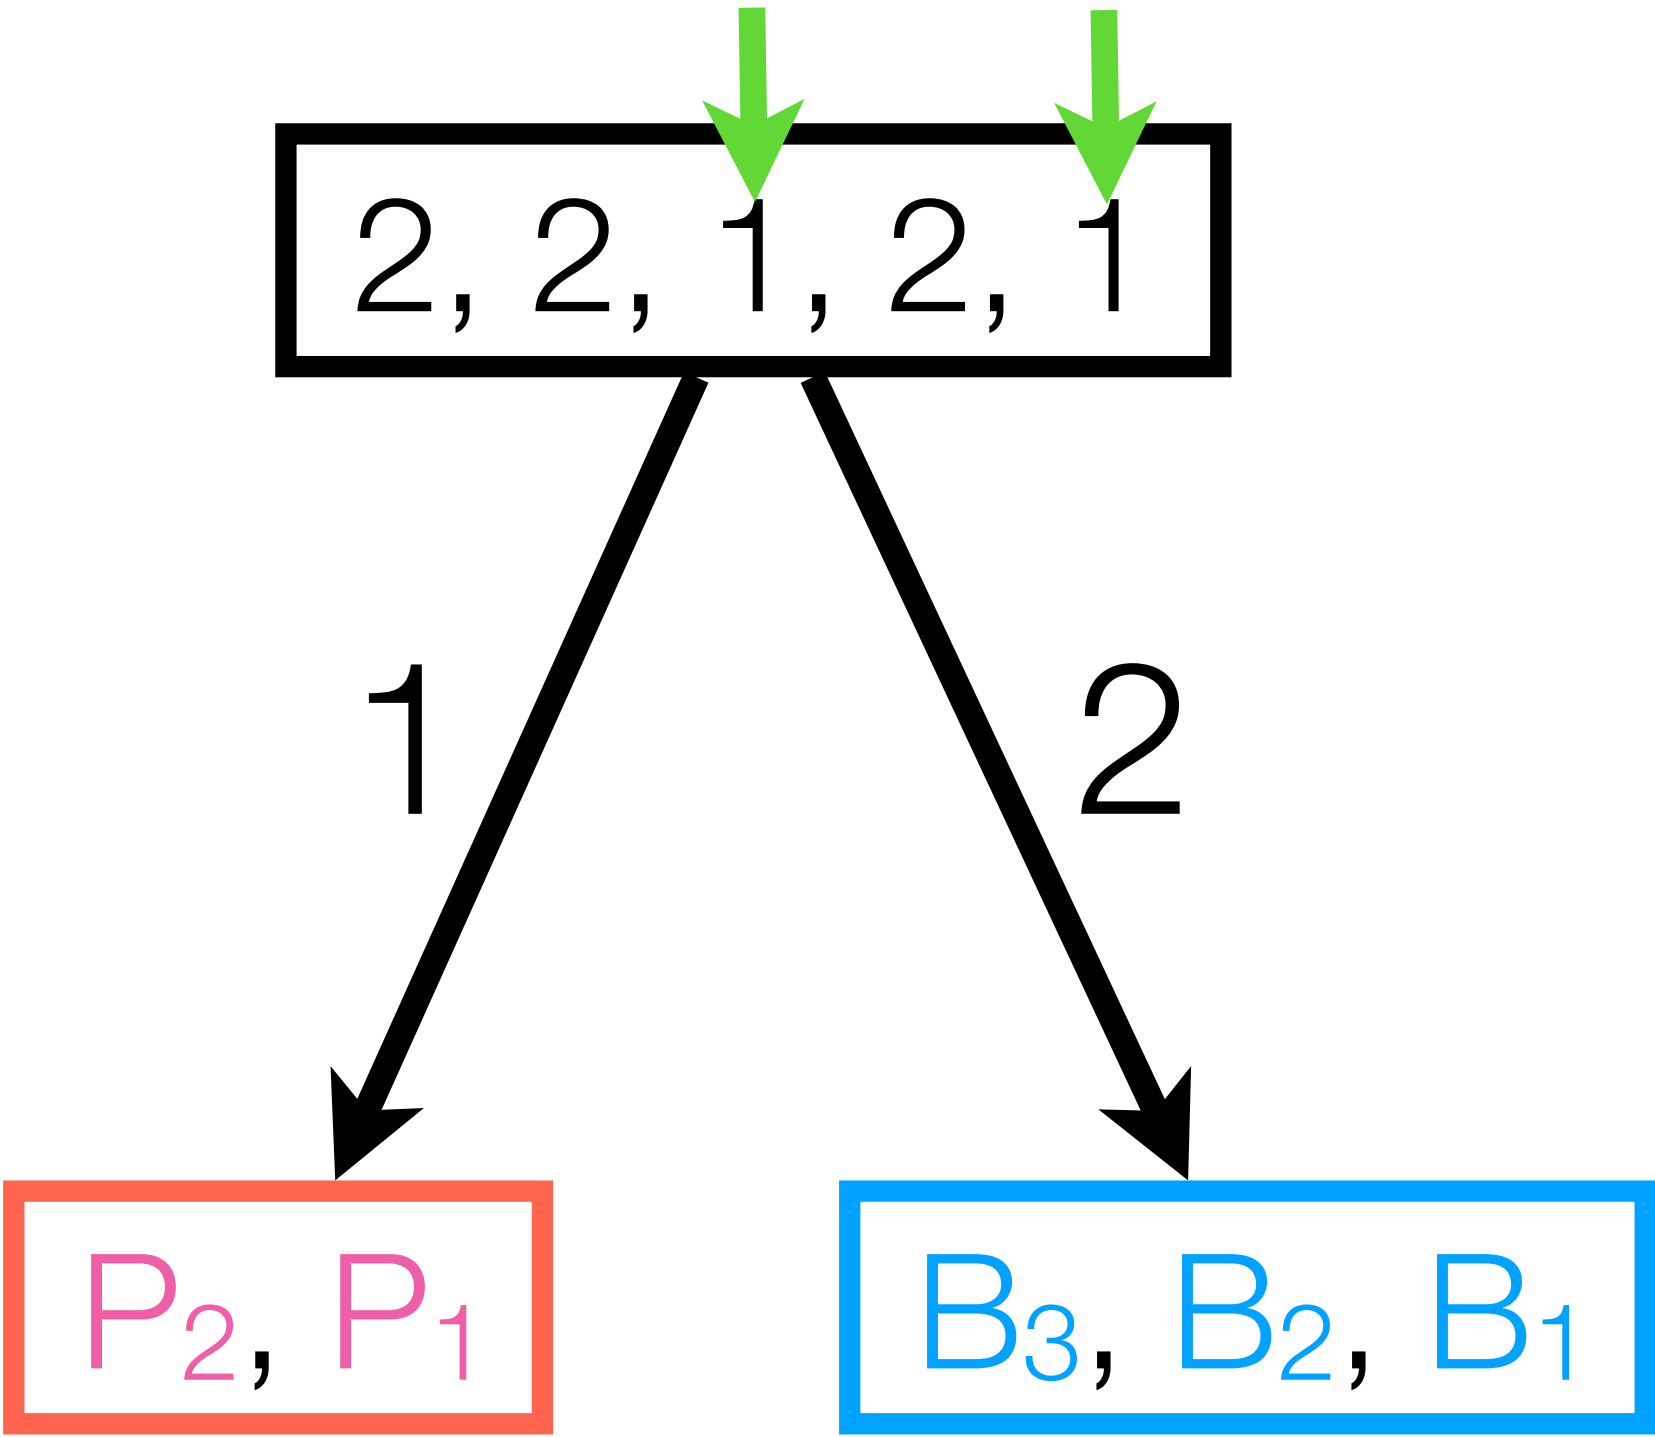
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



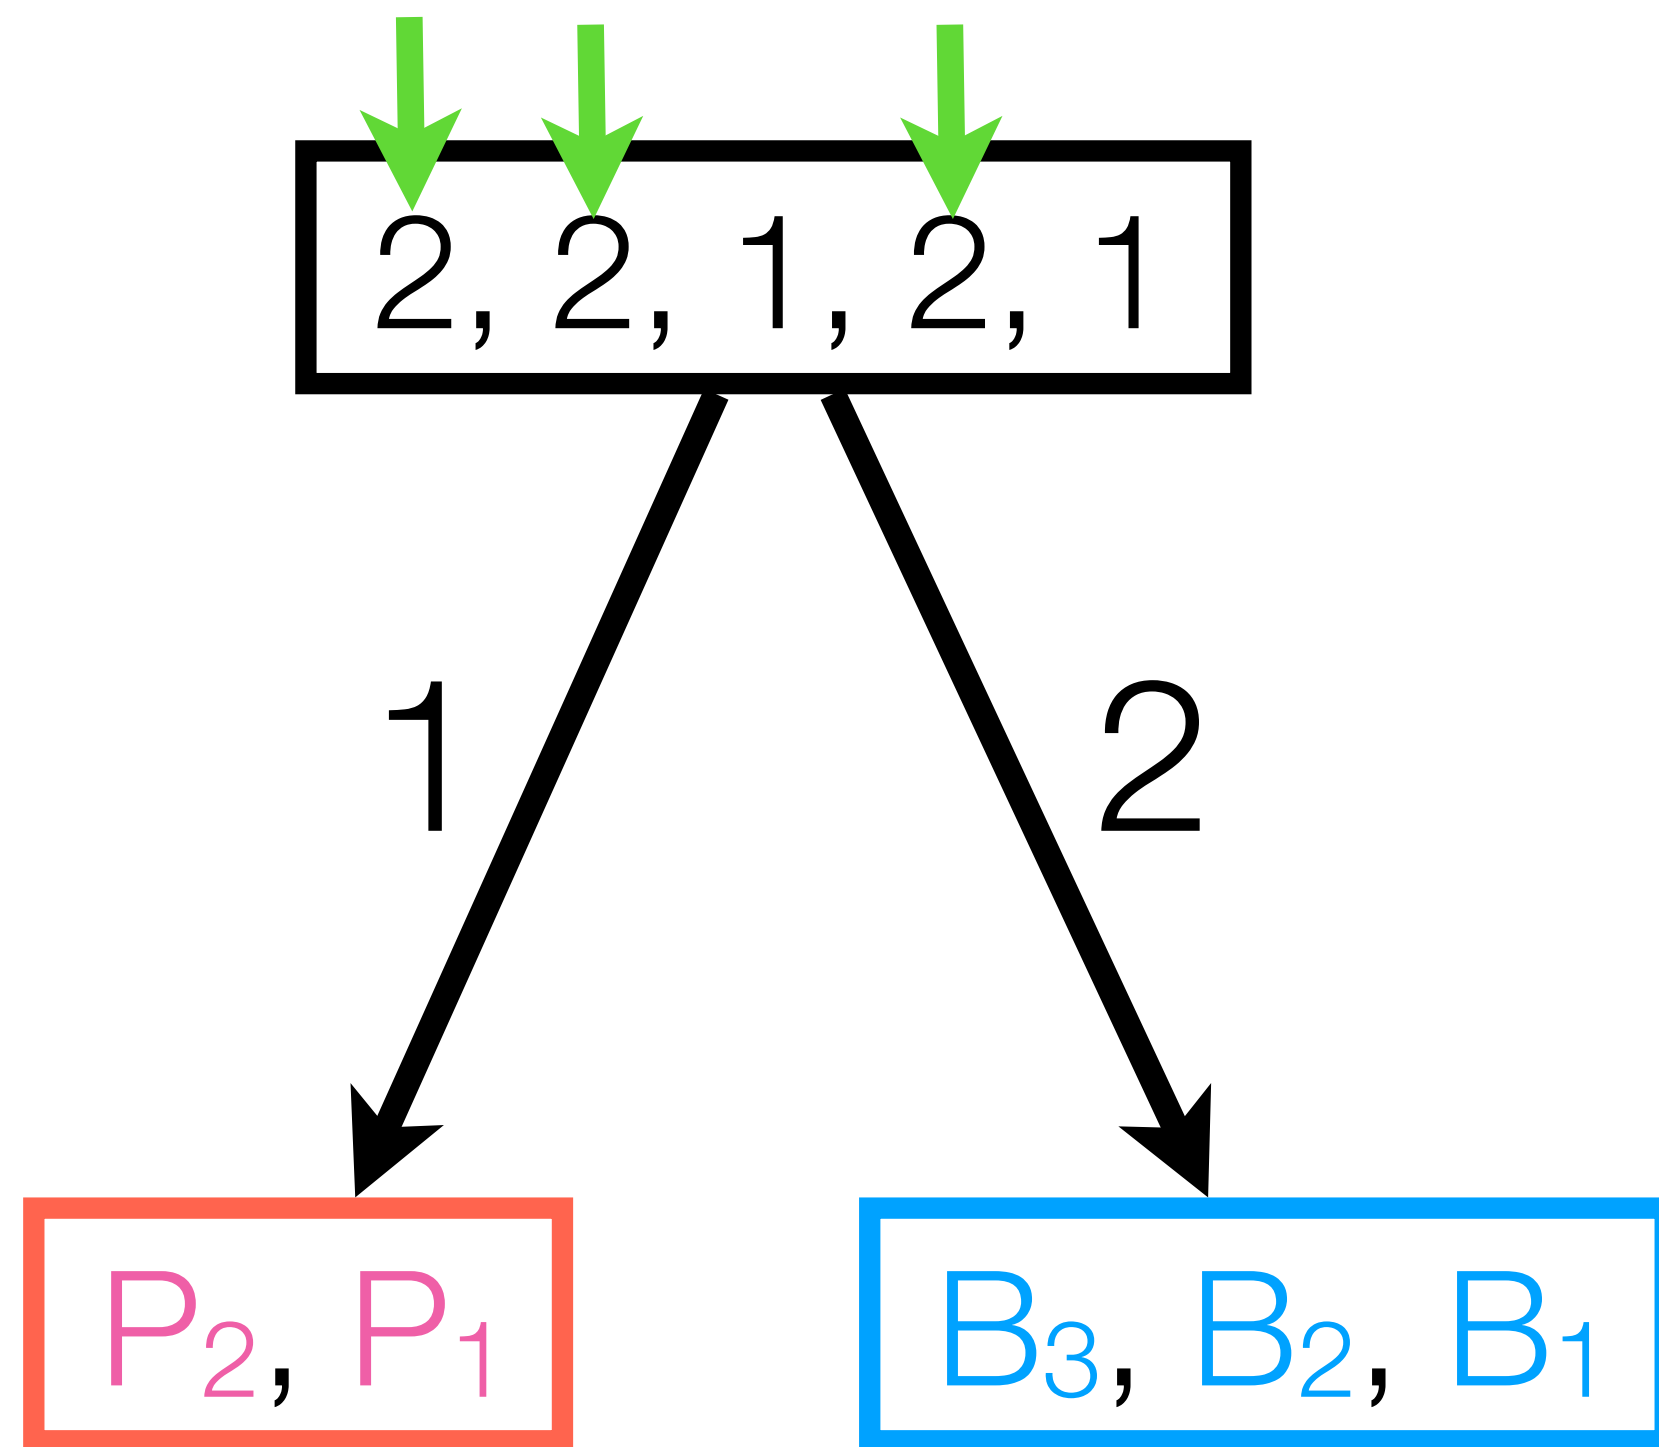
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



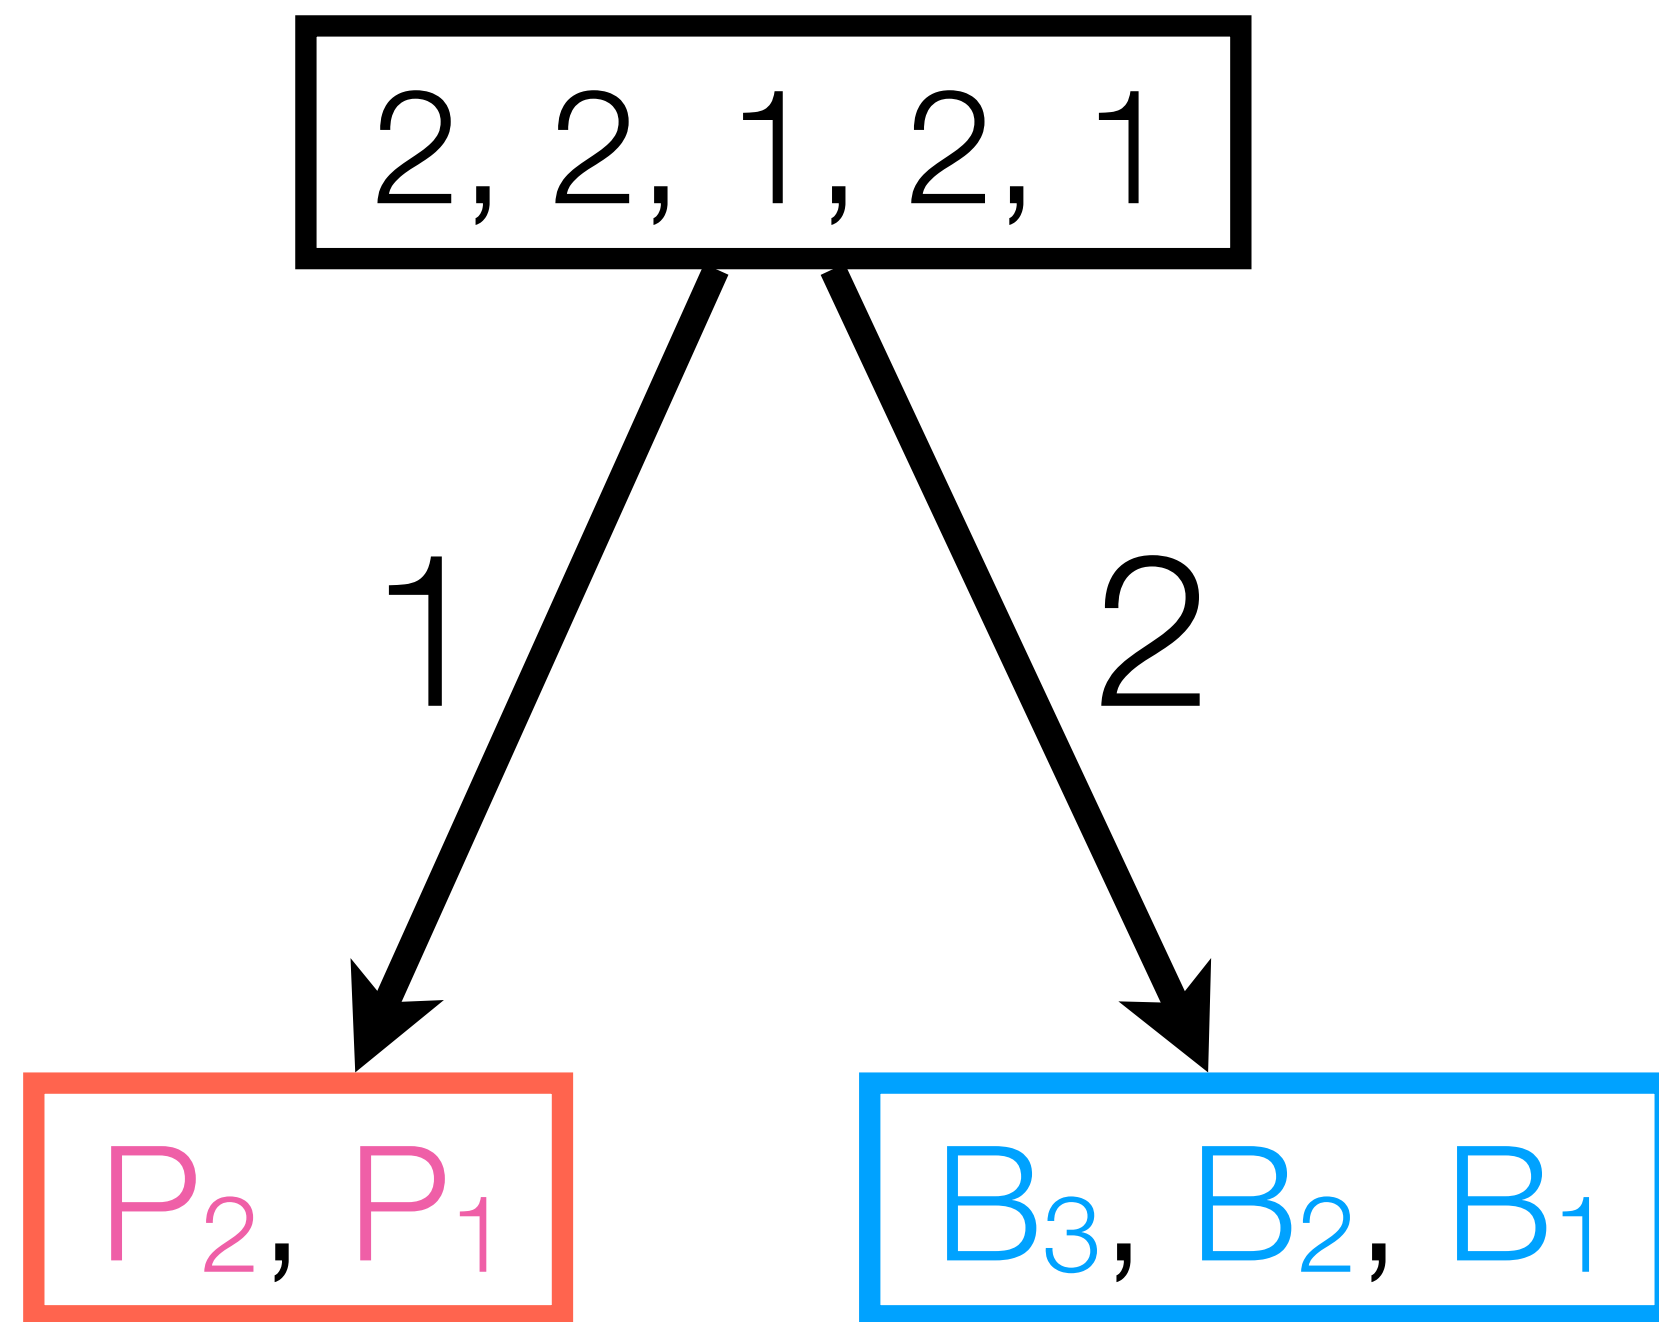
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



Introducing: PIFO trees

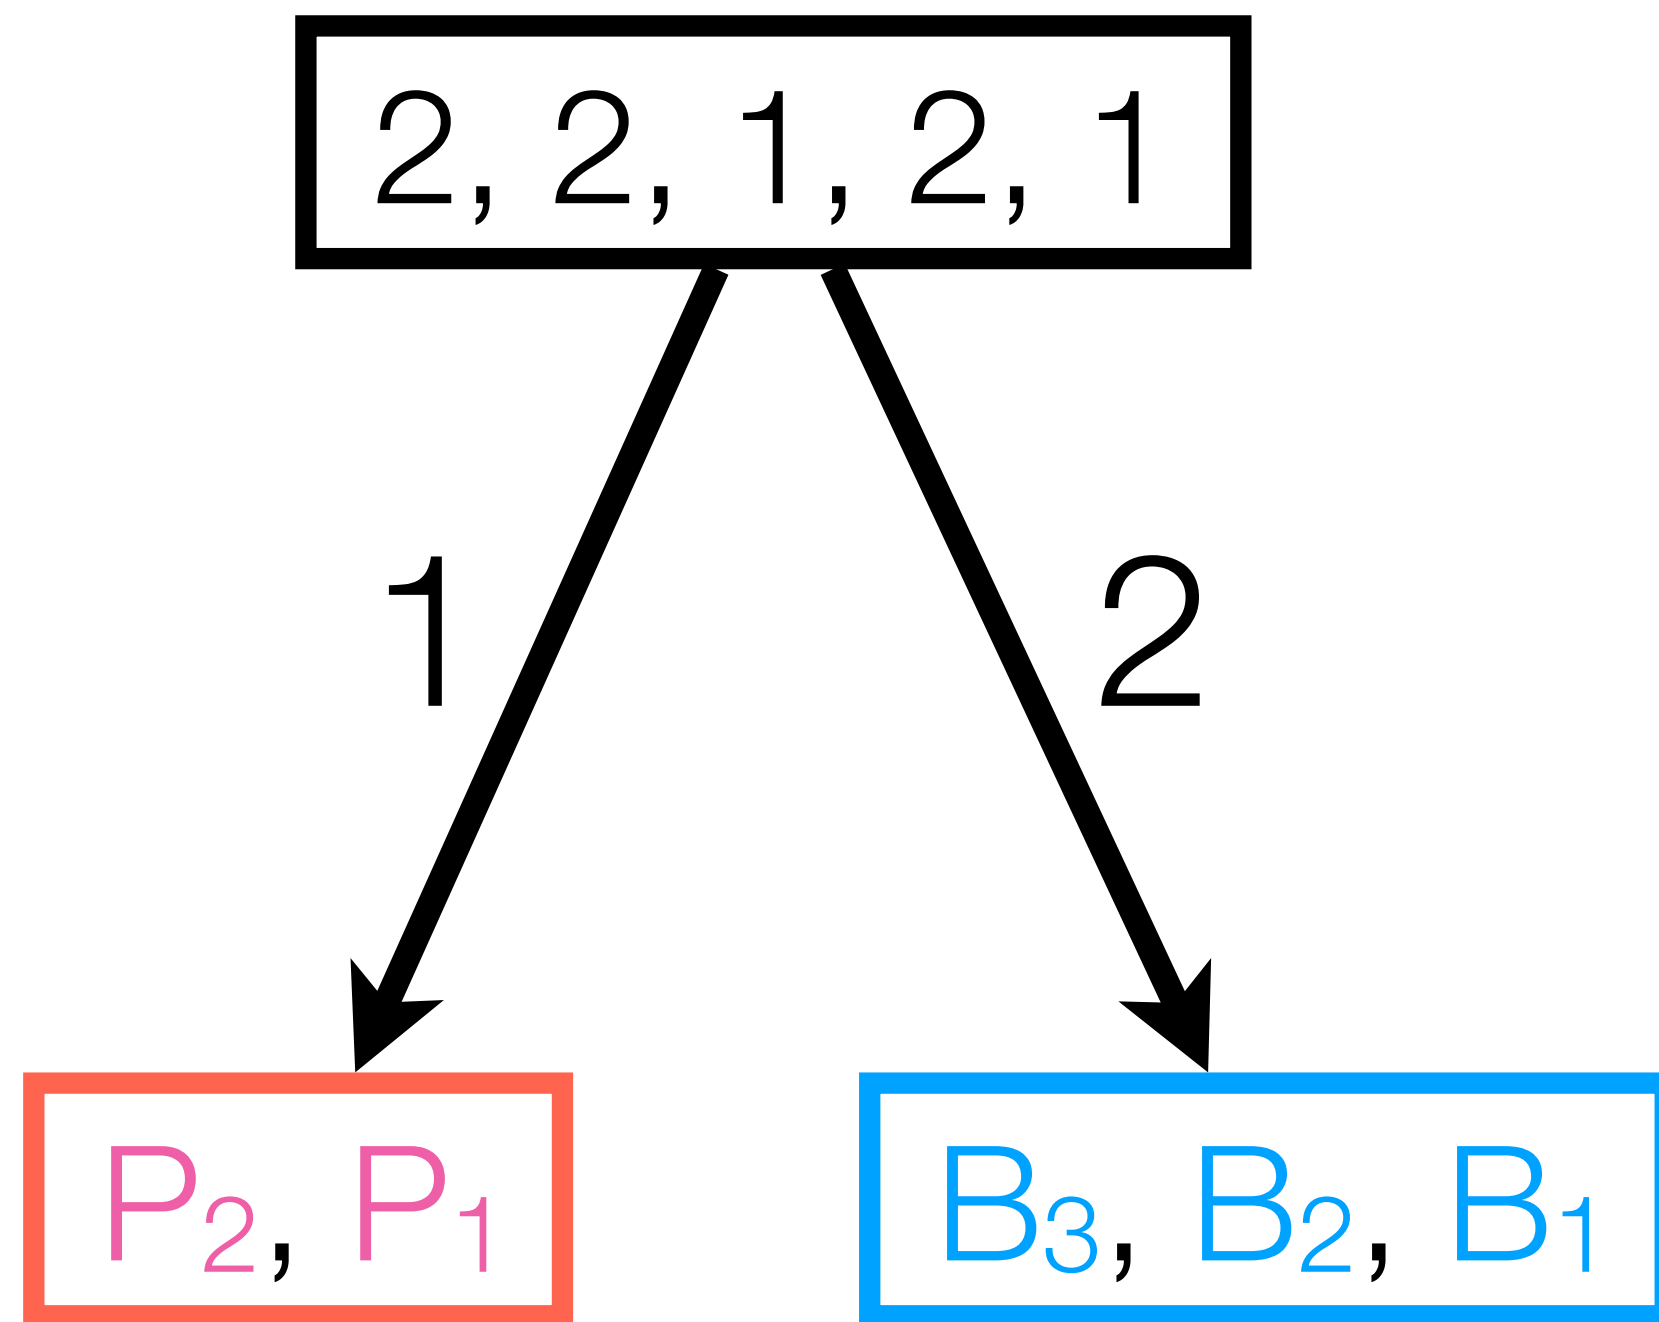
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!

Introducing: PIFO trees

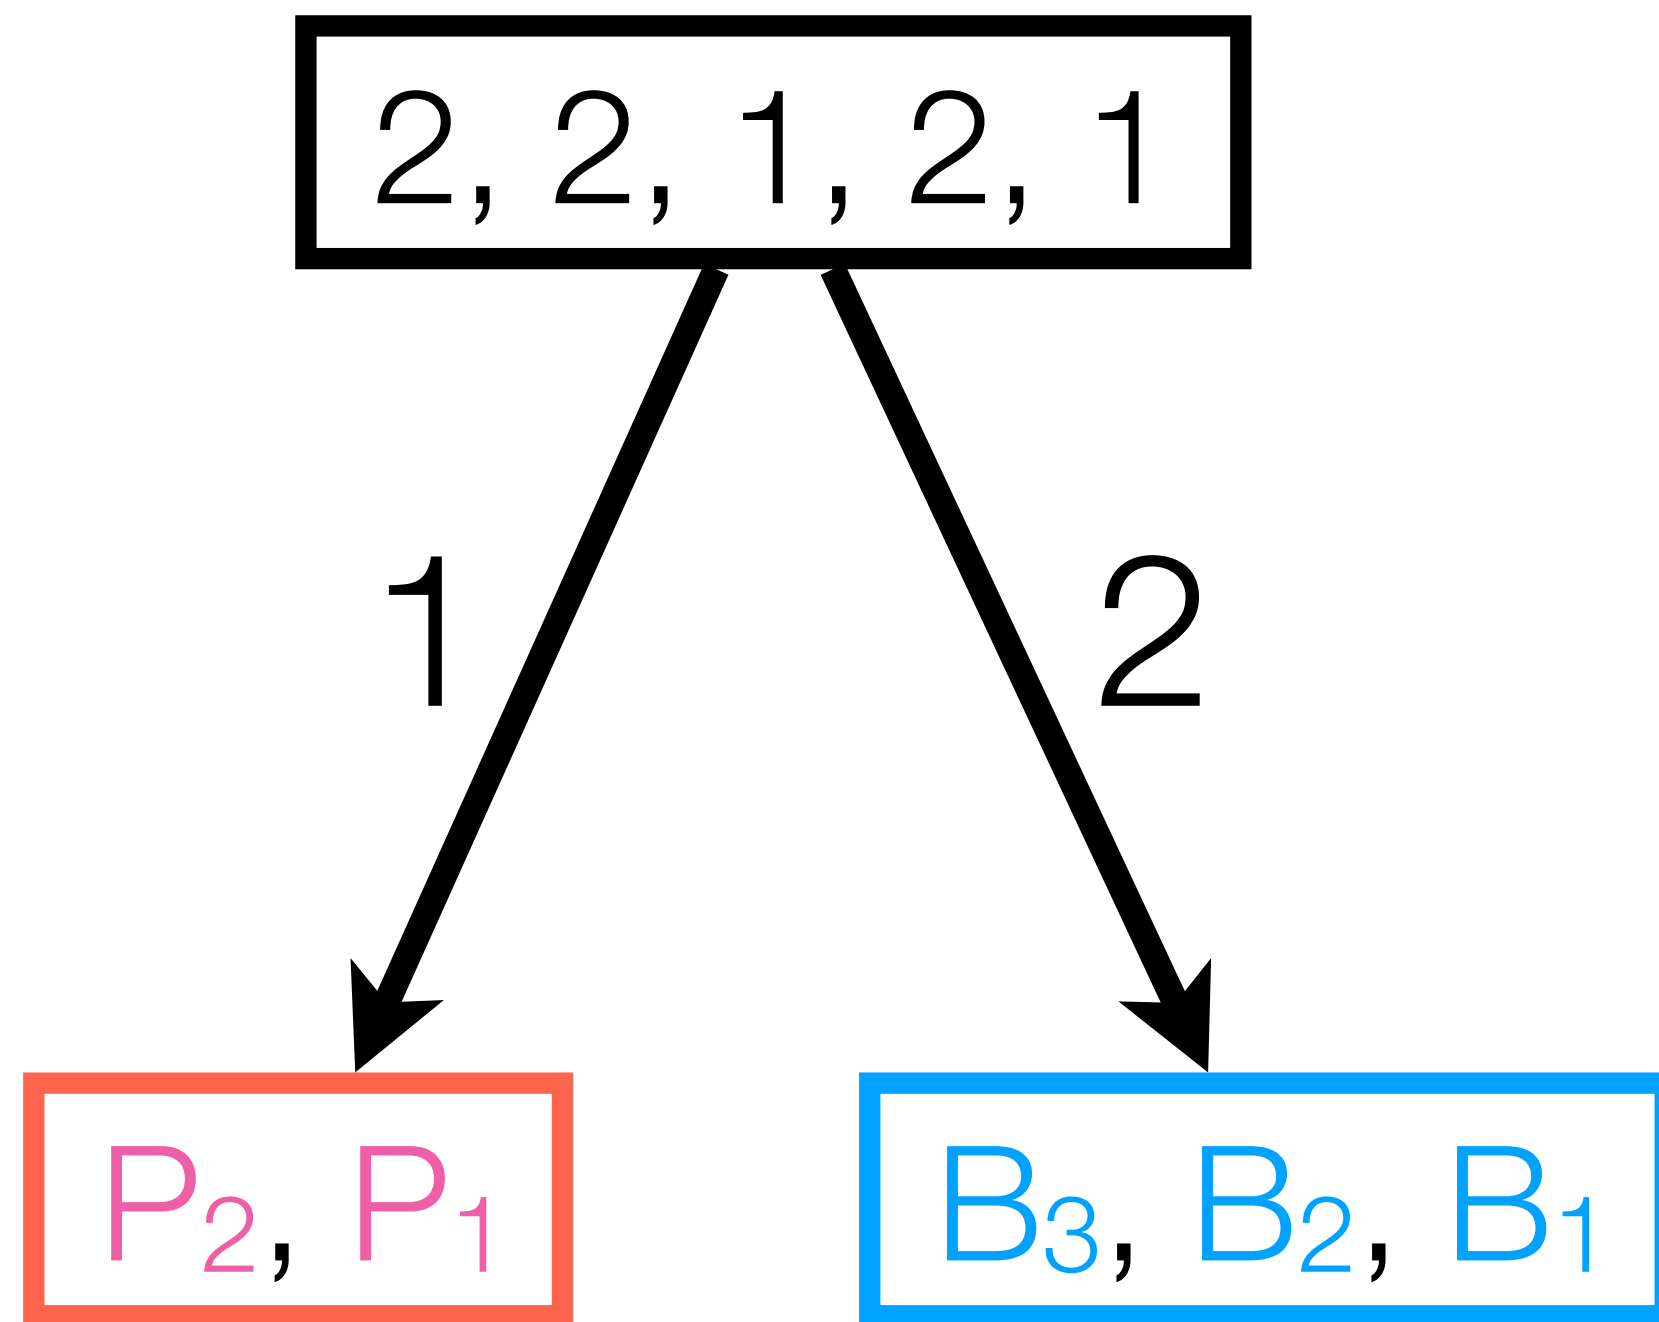
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

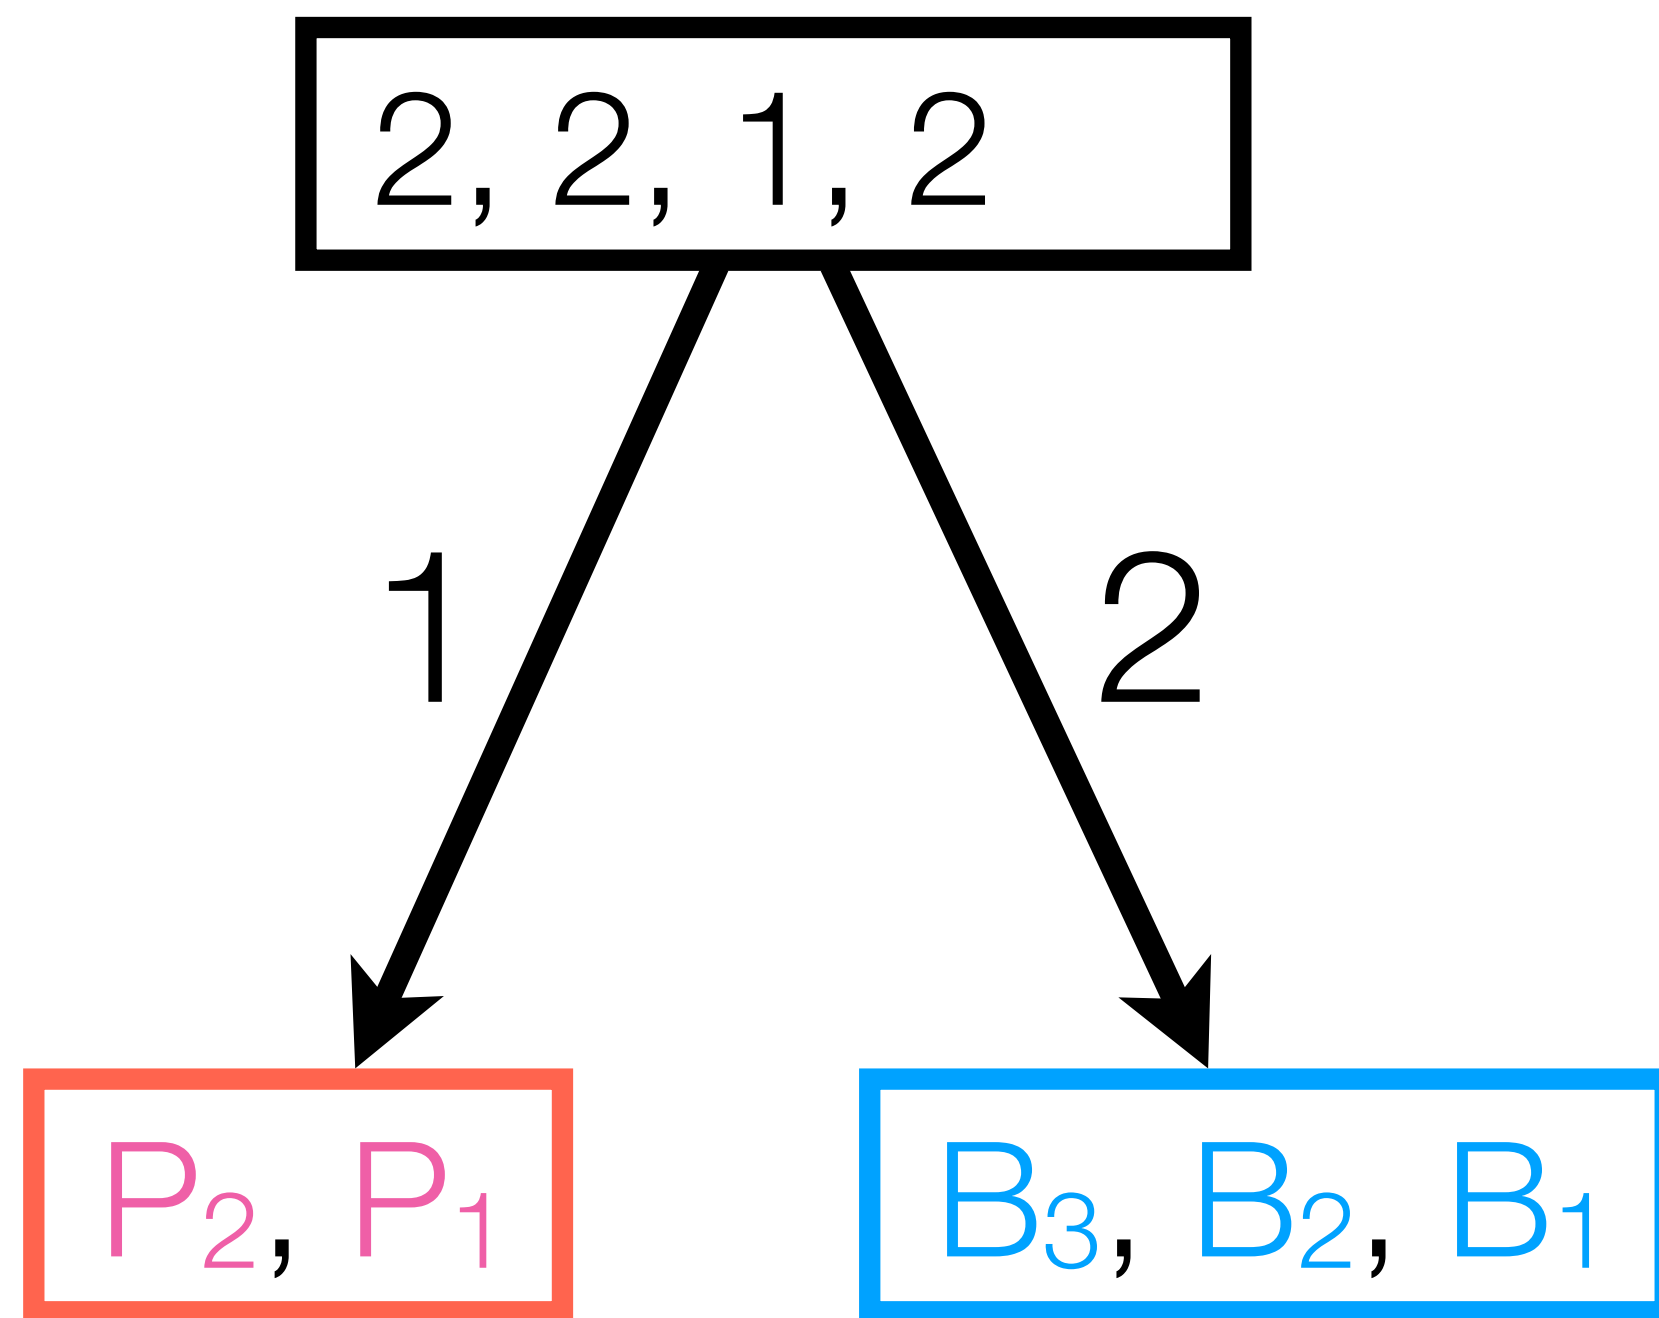
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

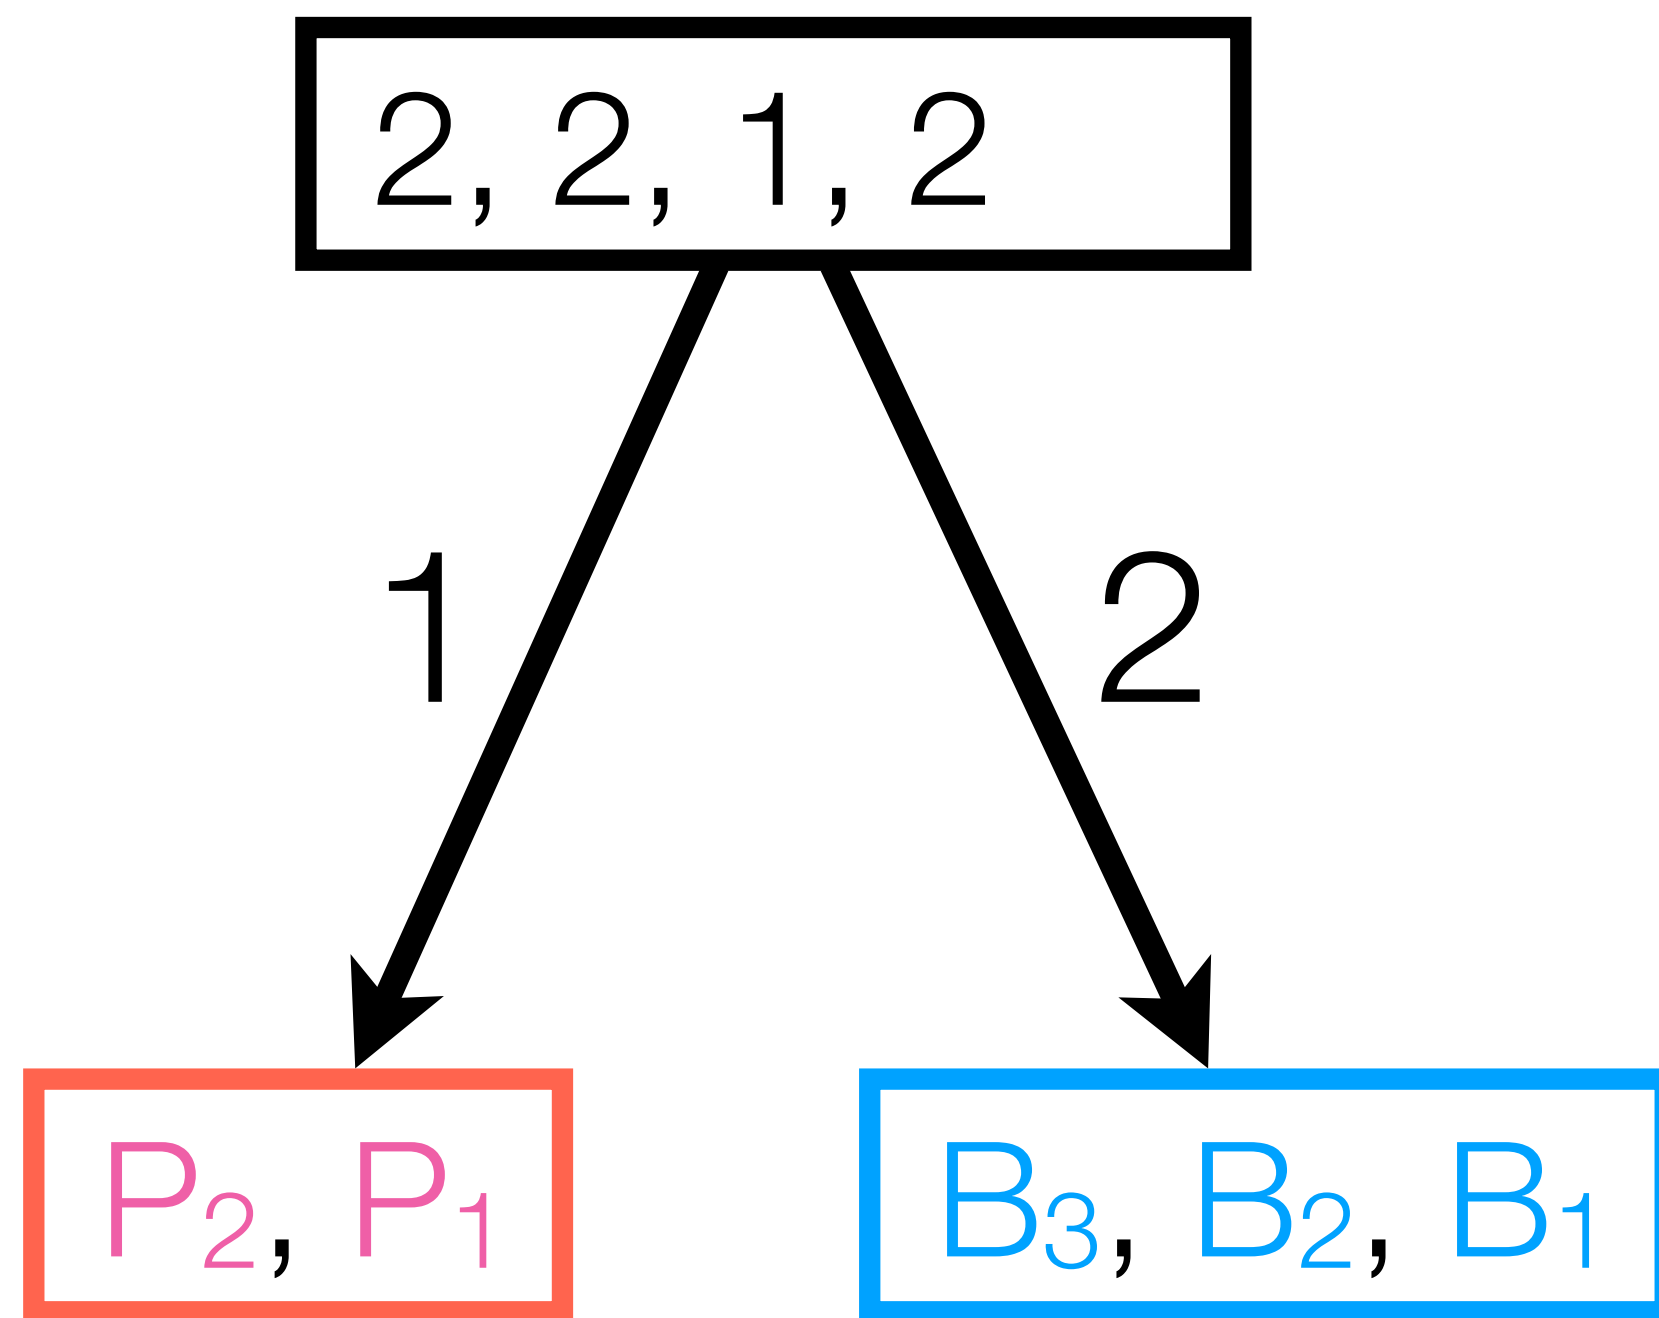
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

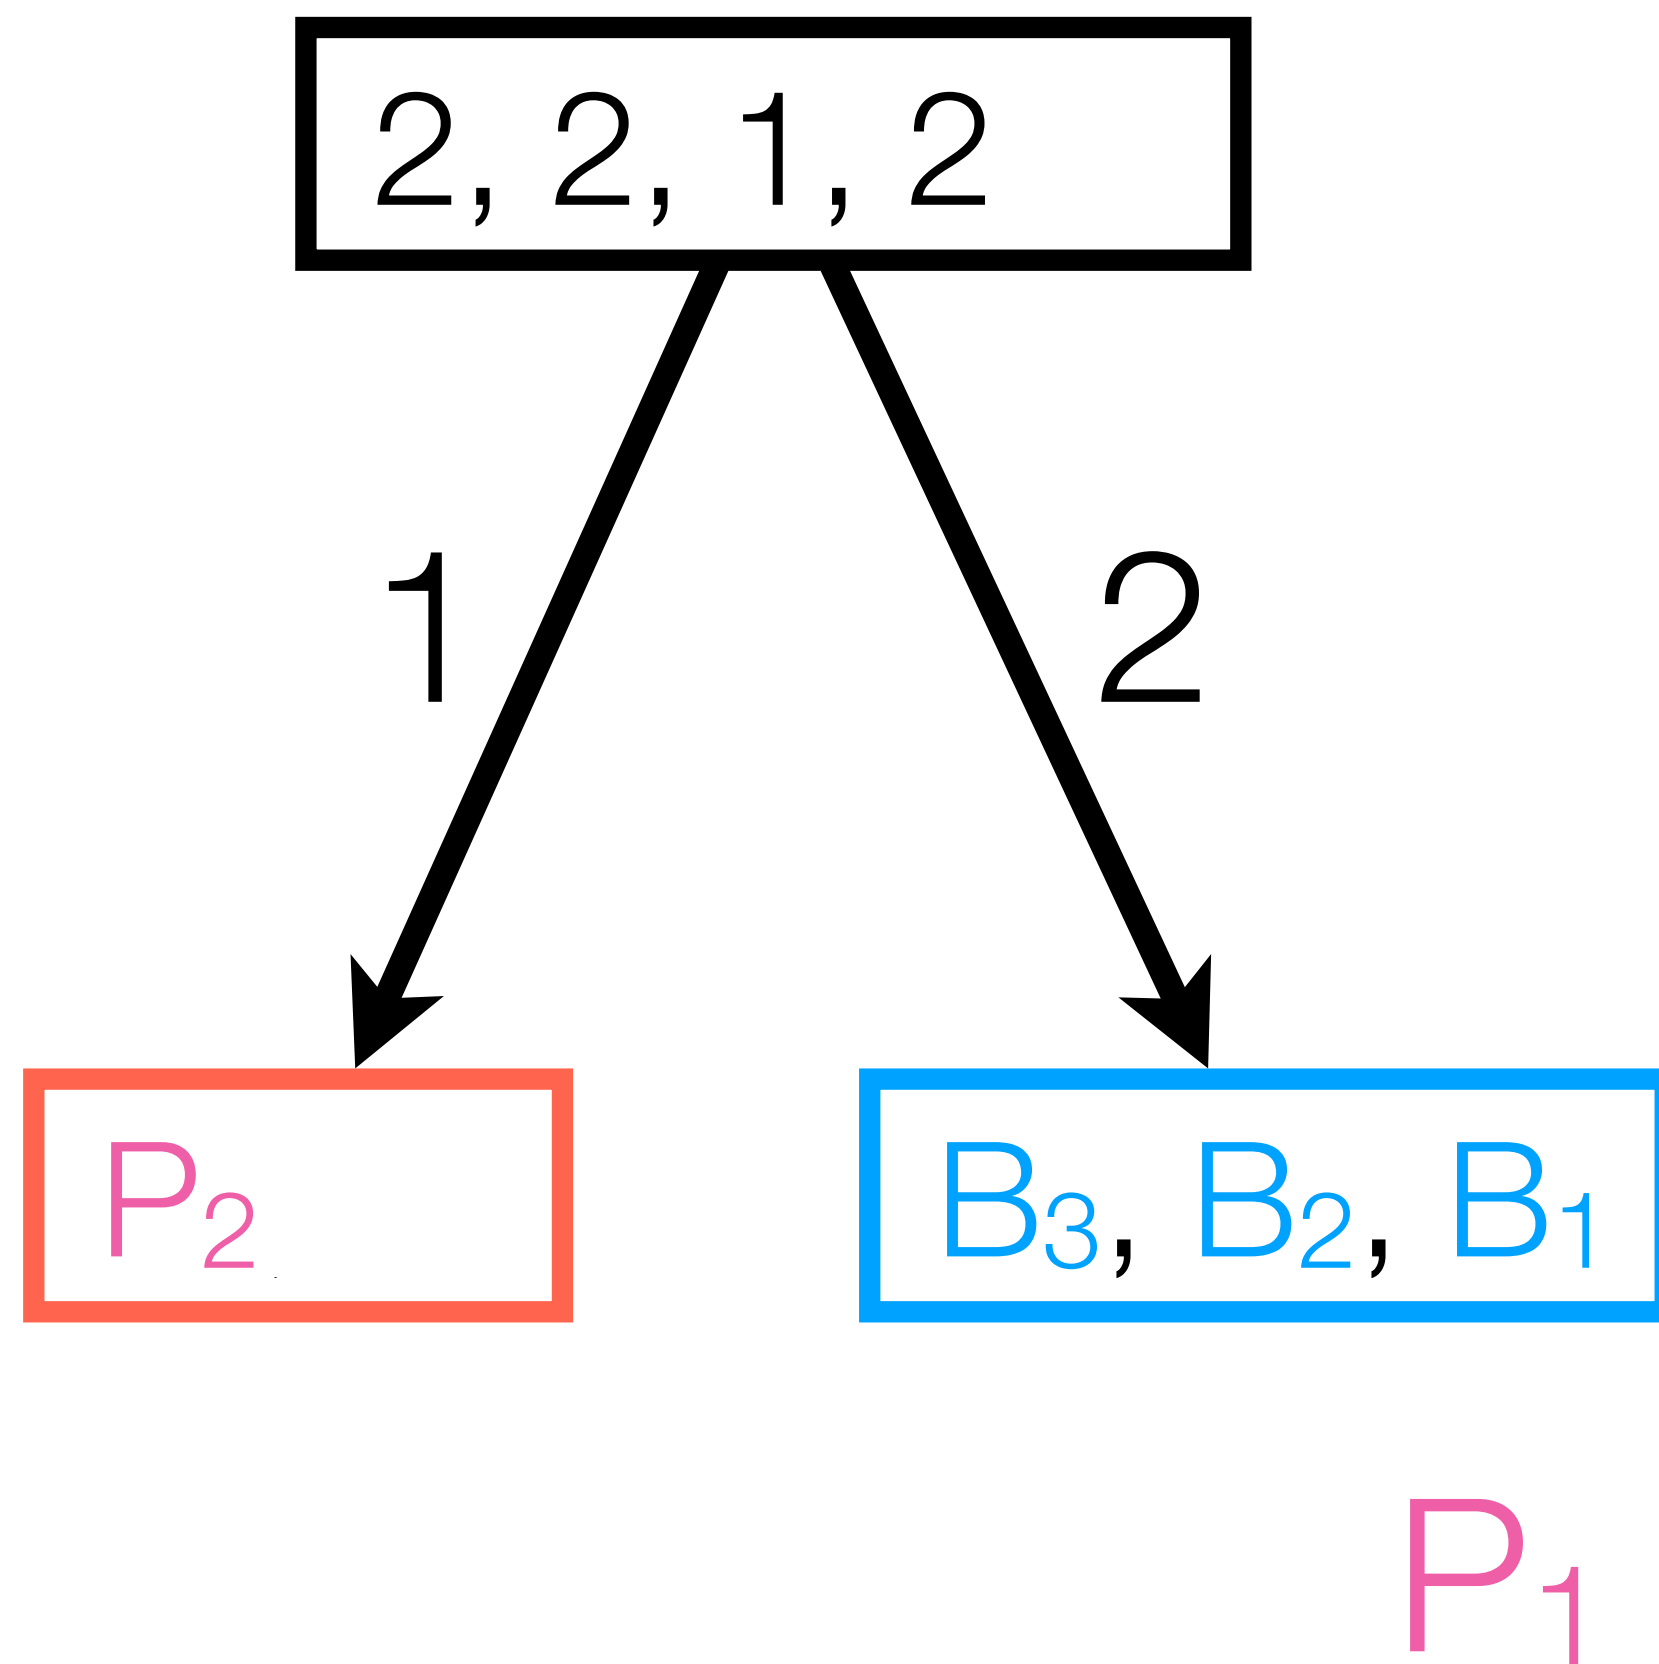
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

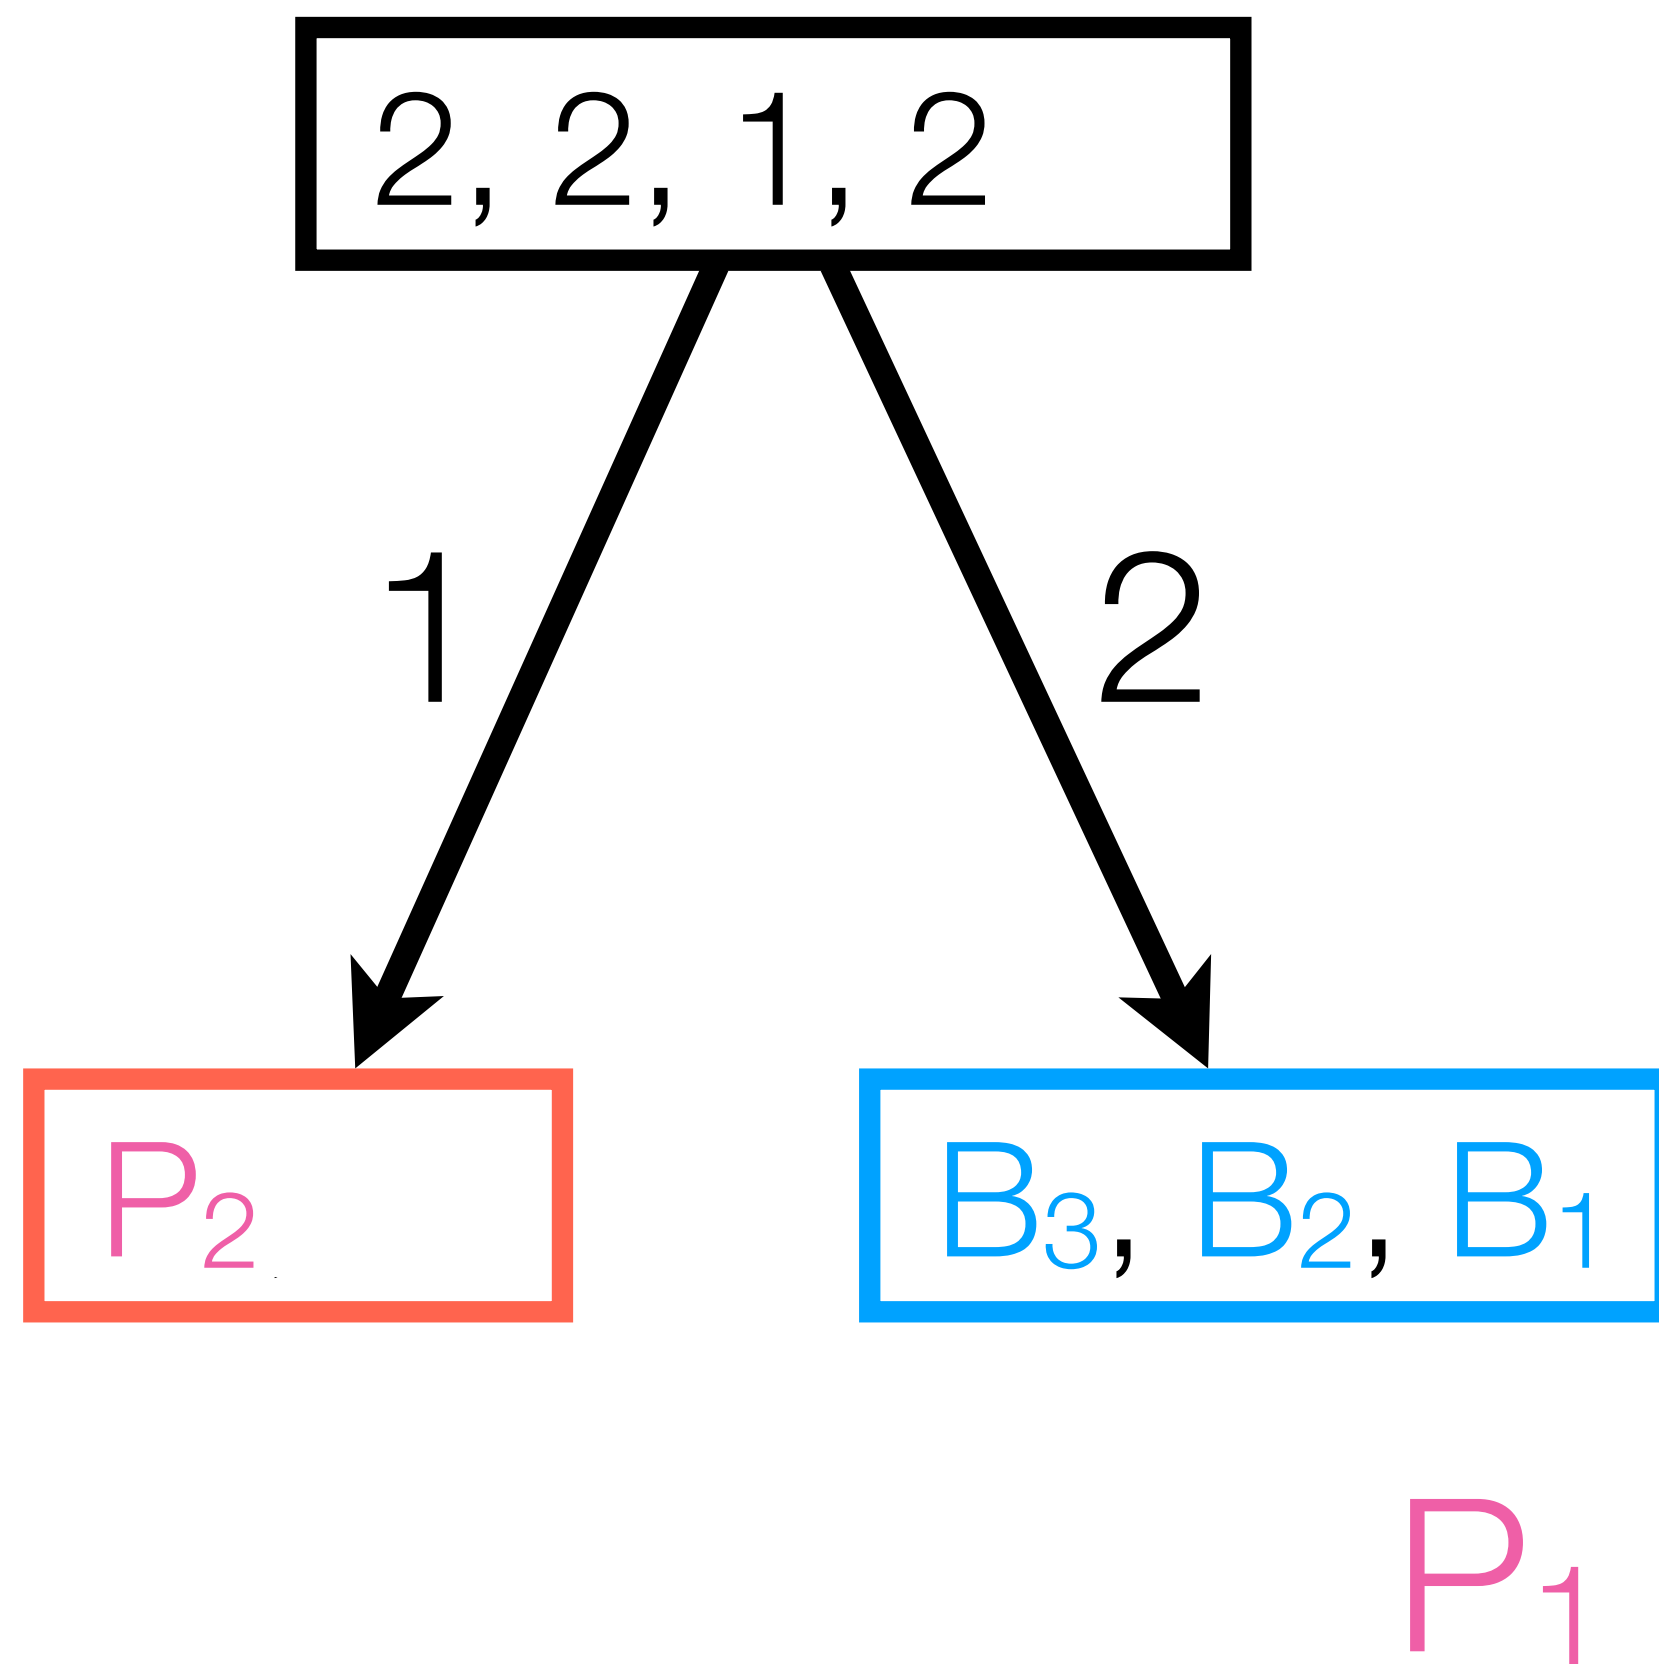
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

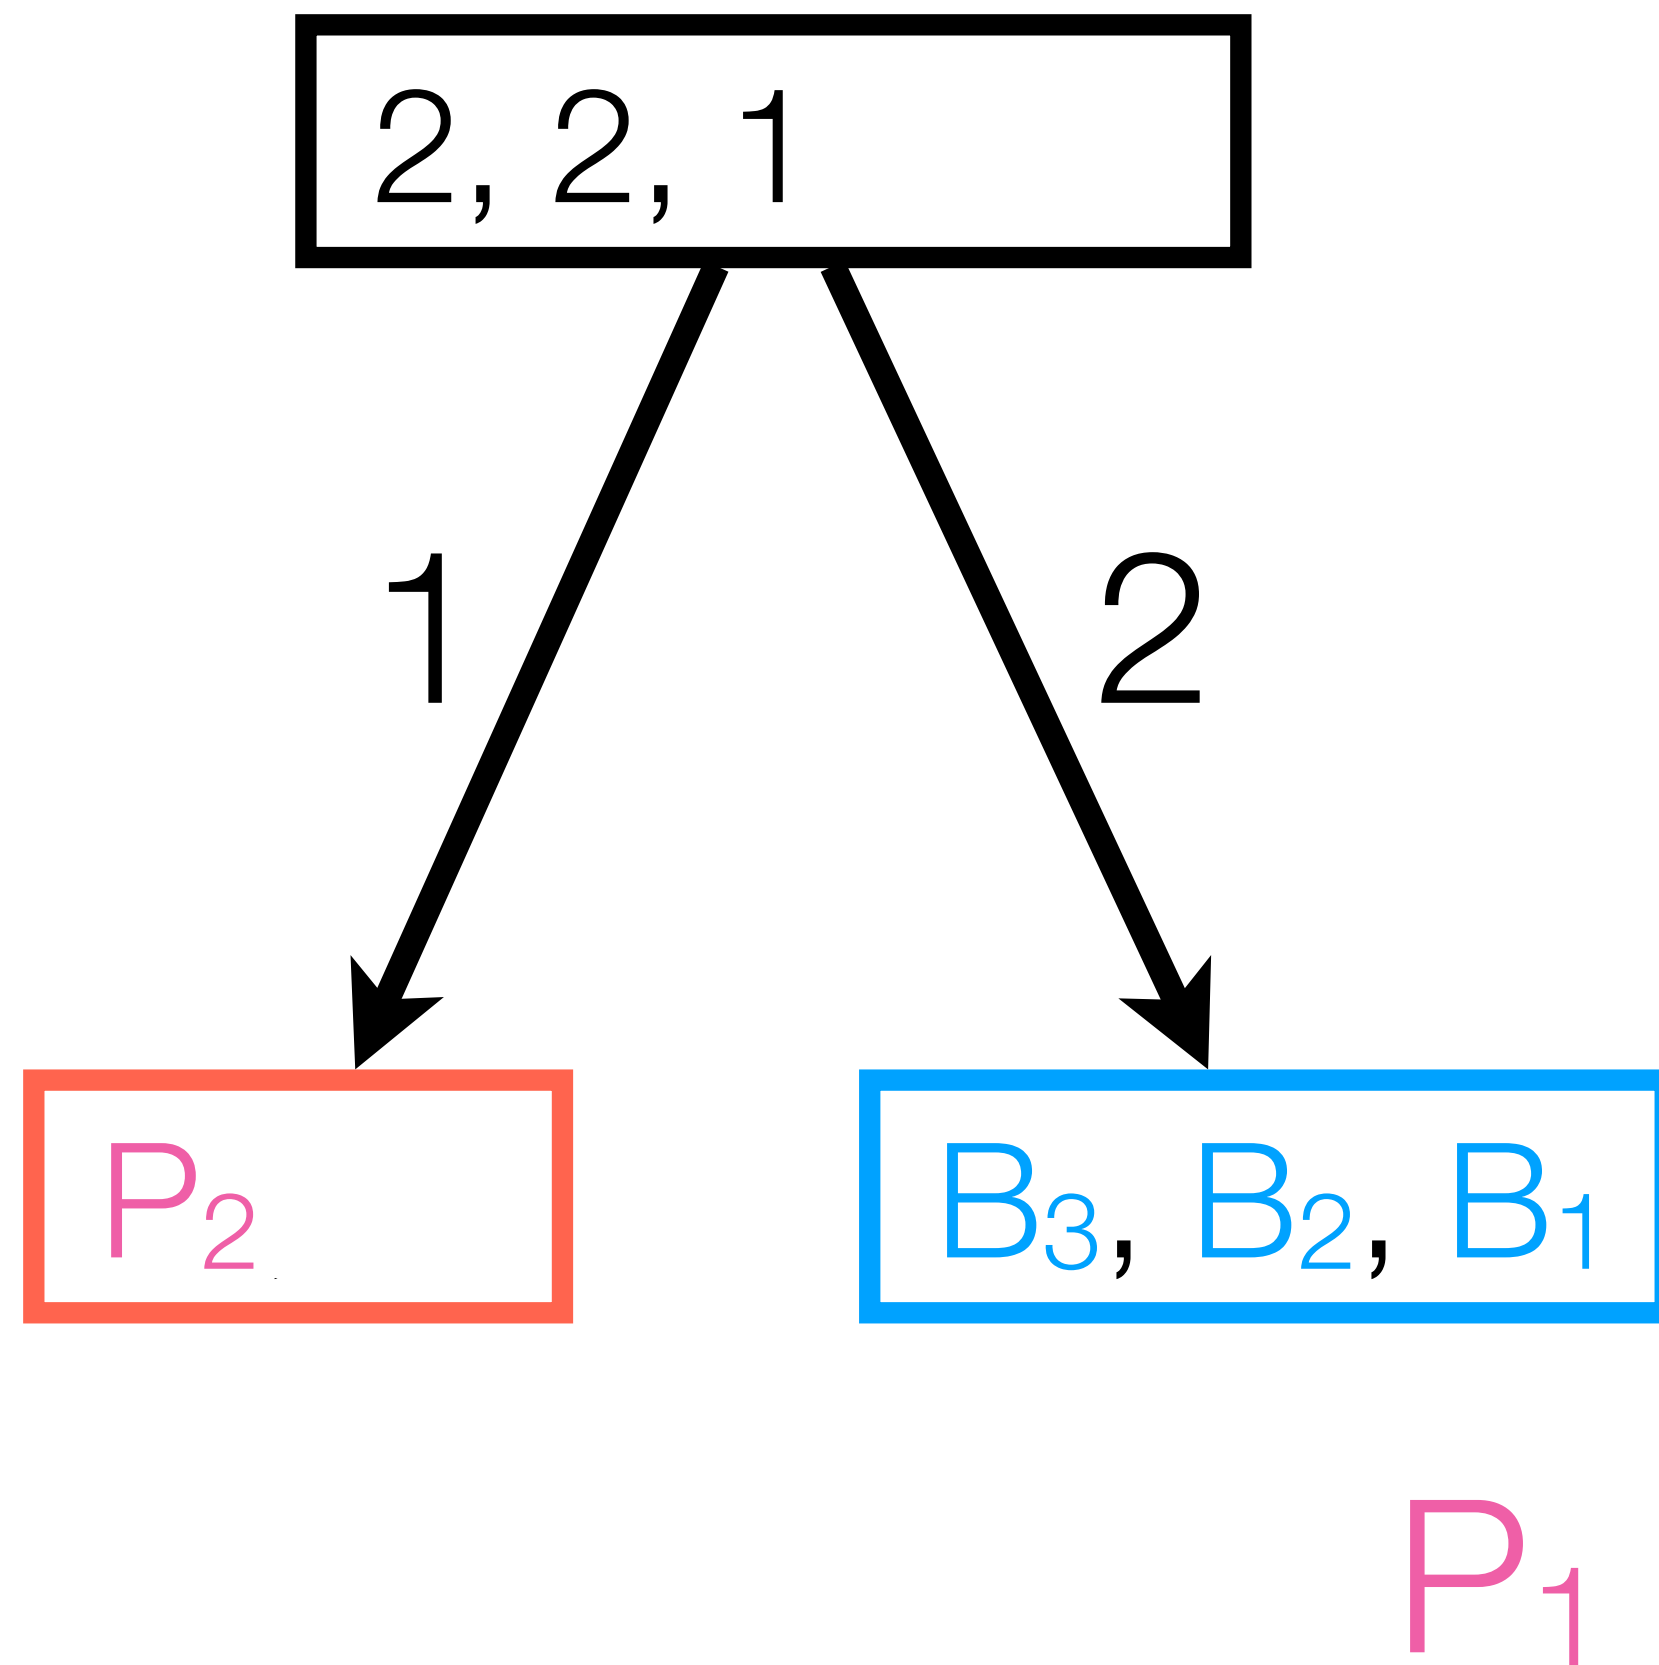
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

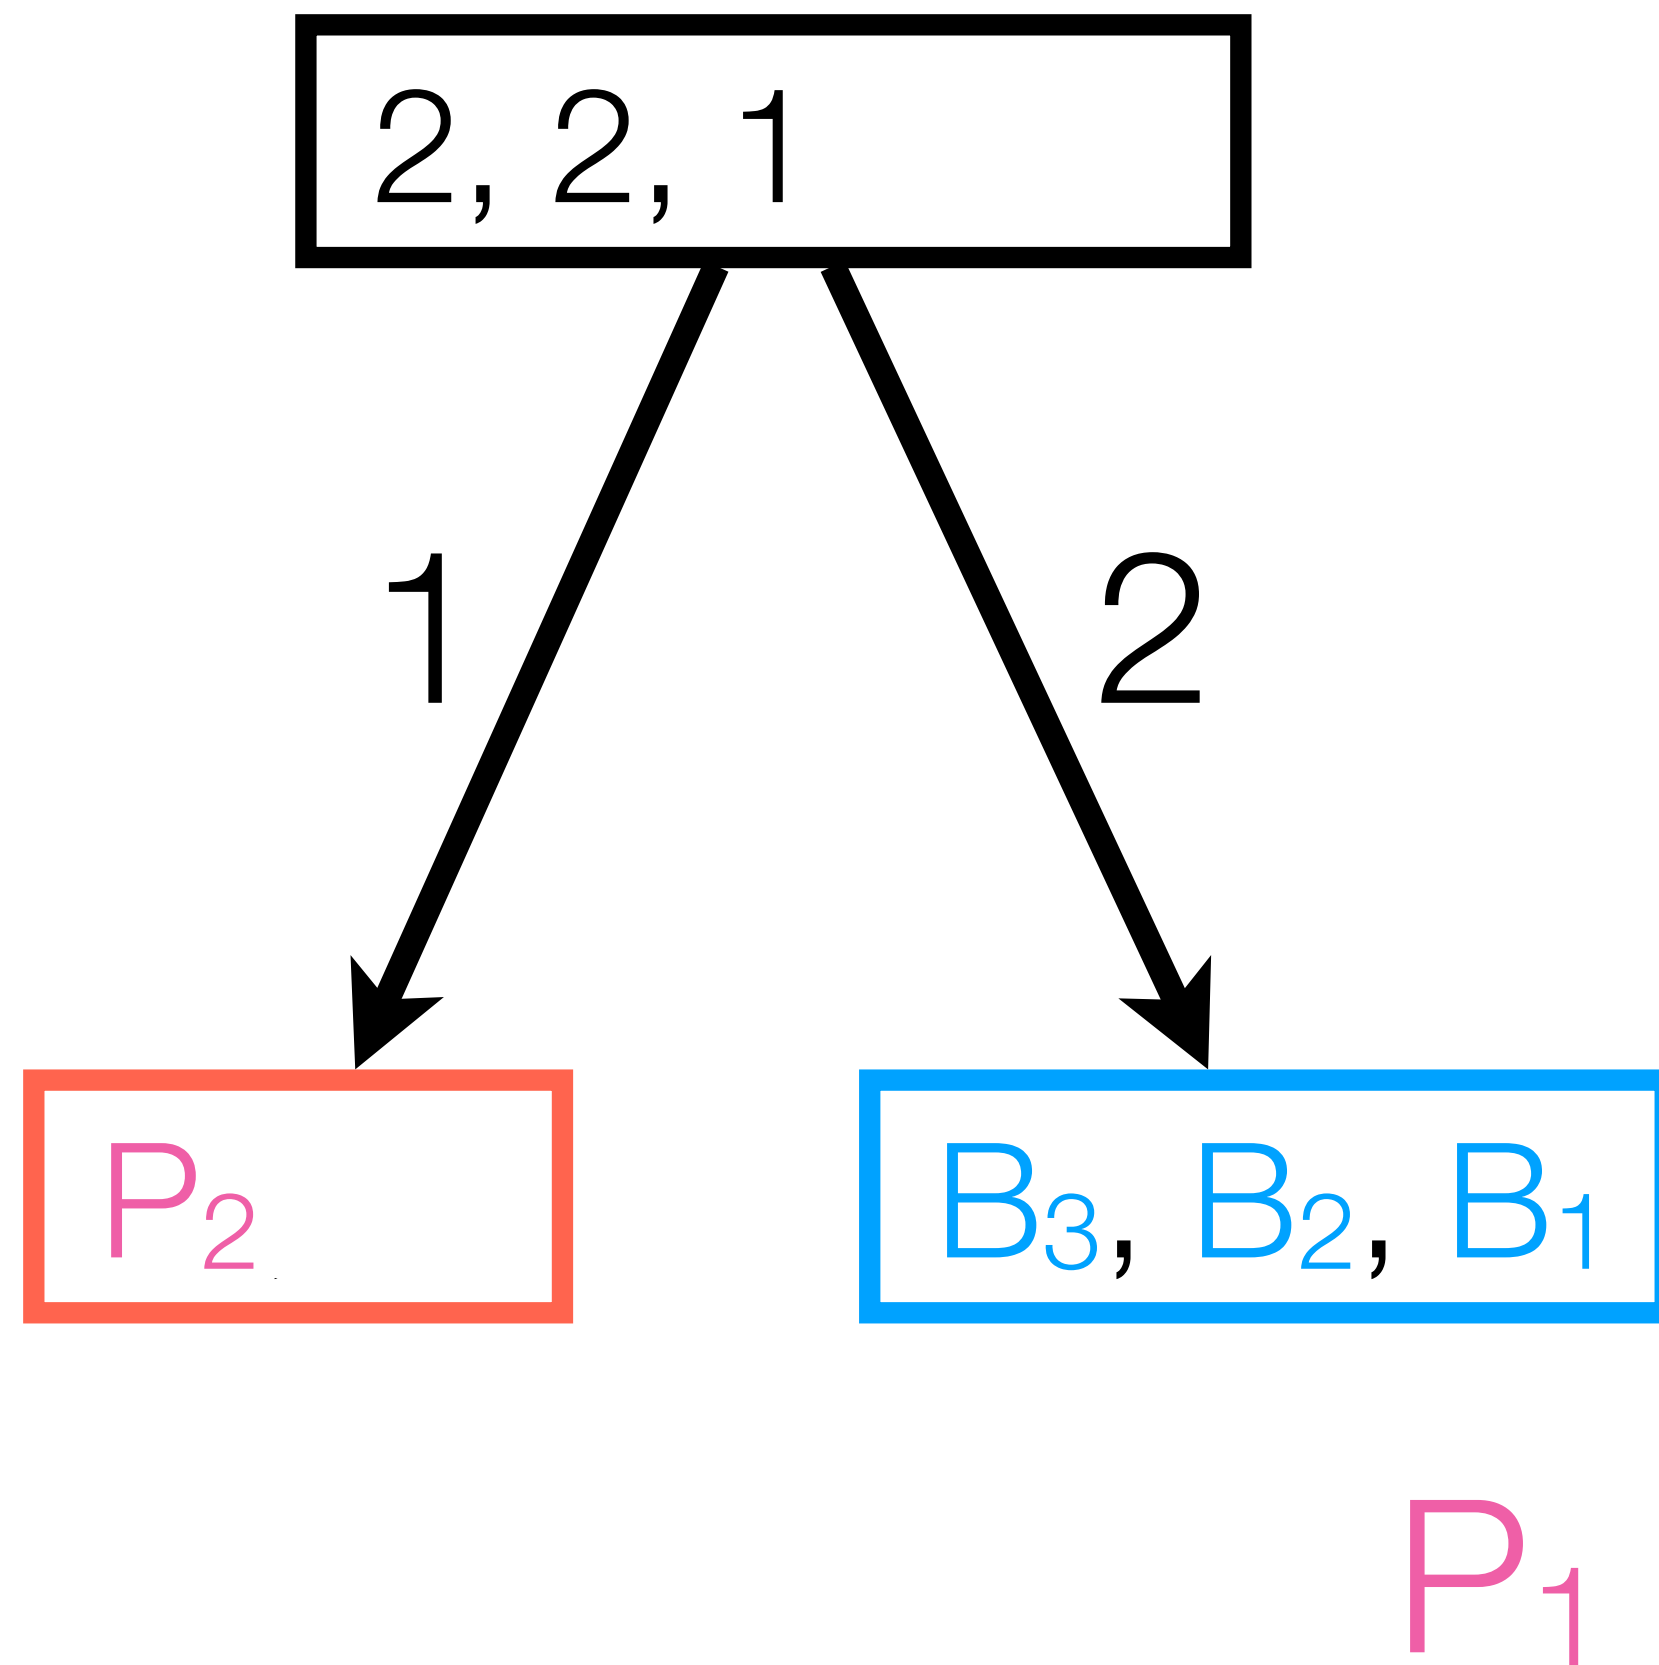
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

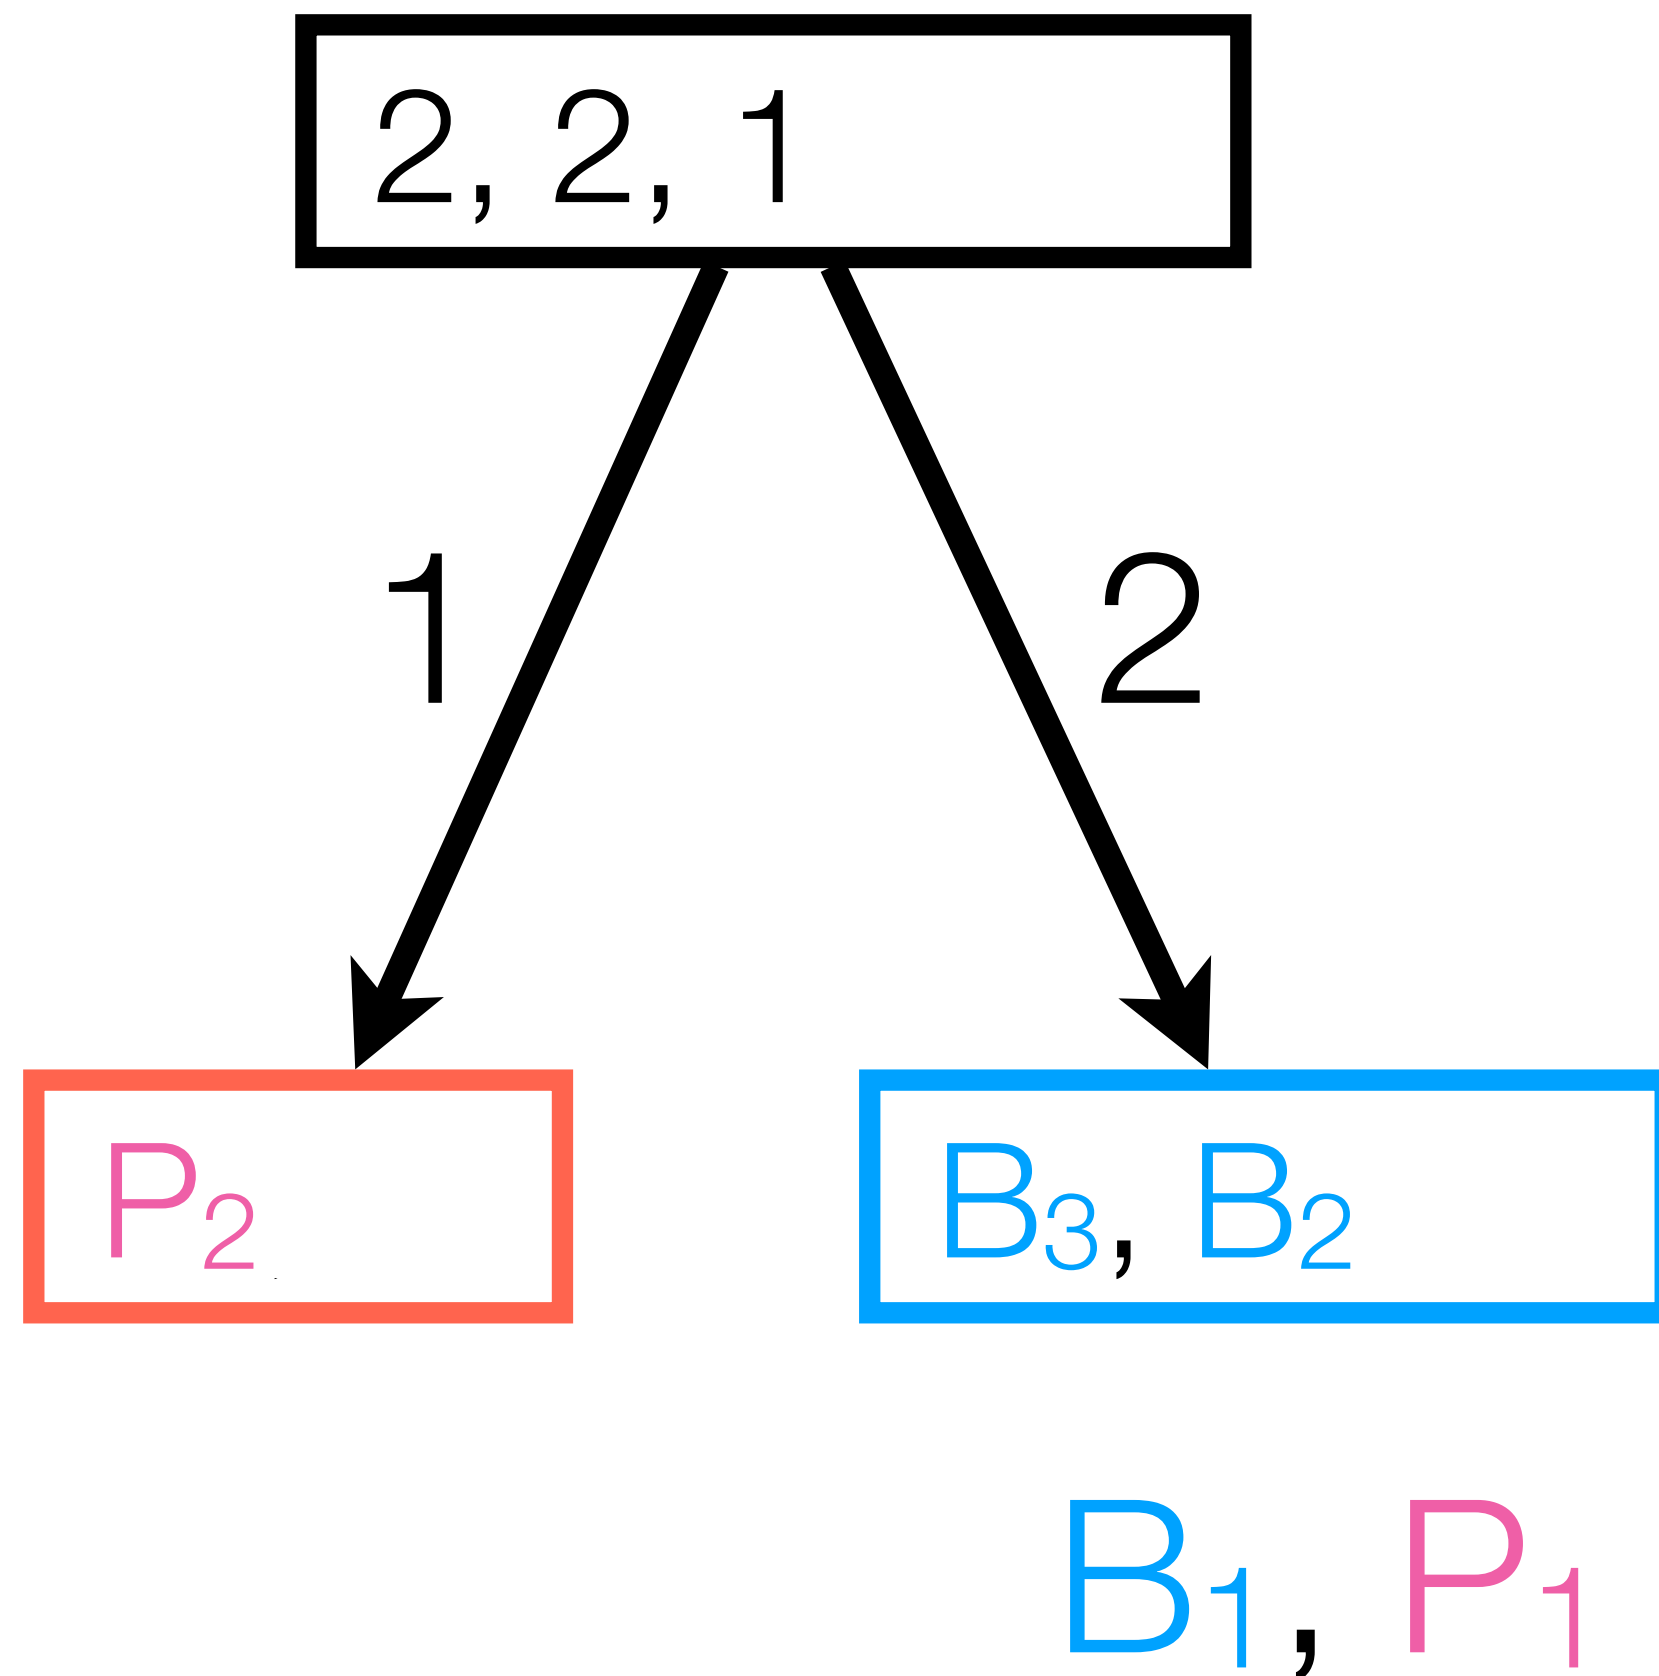
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

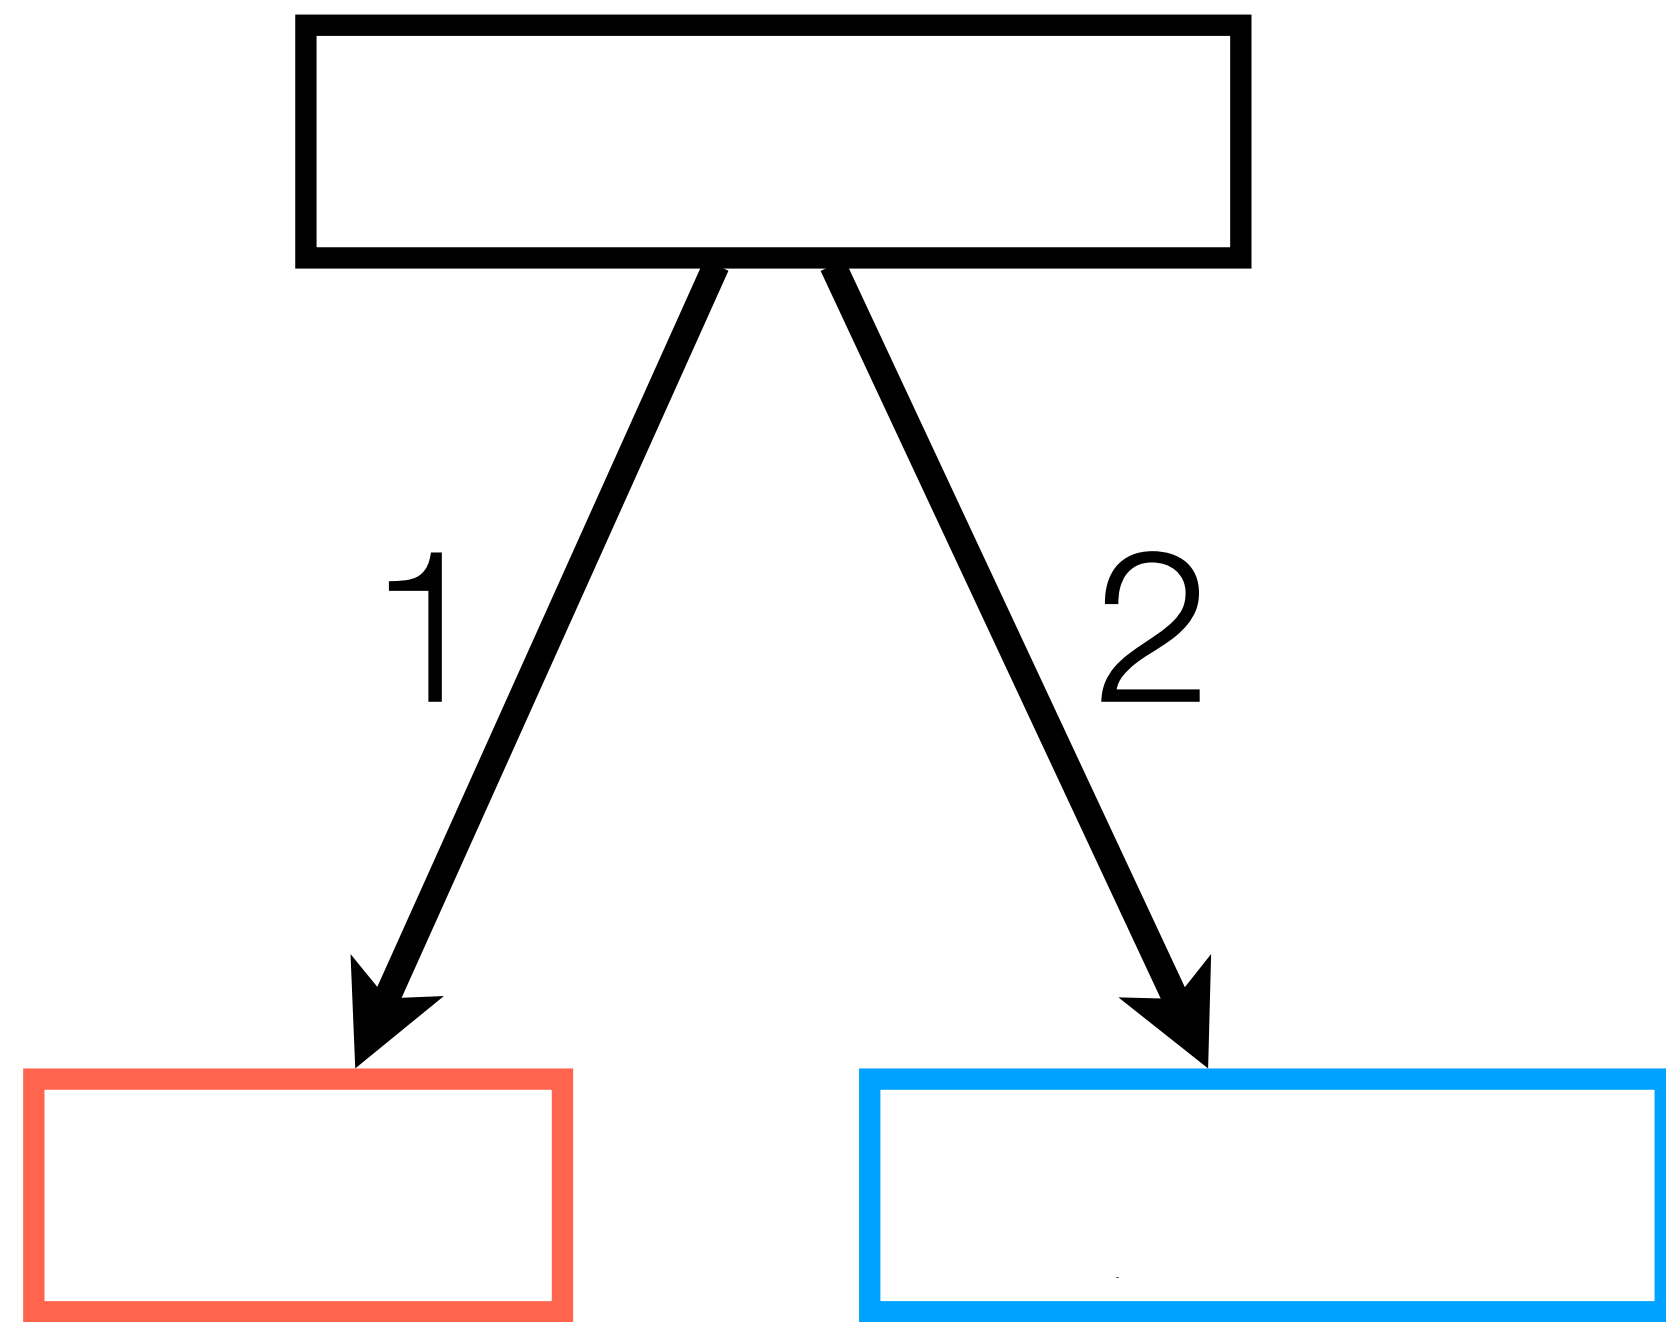
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.

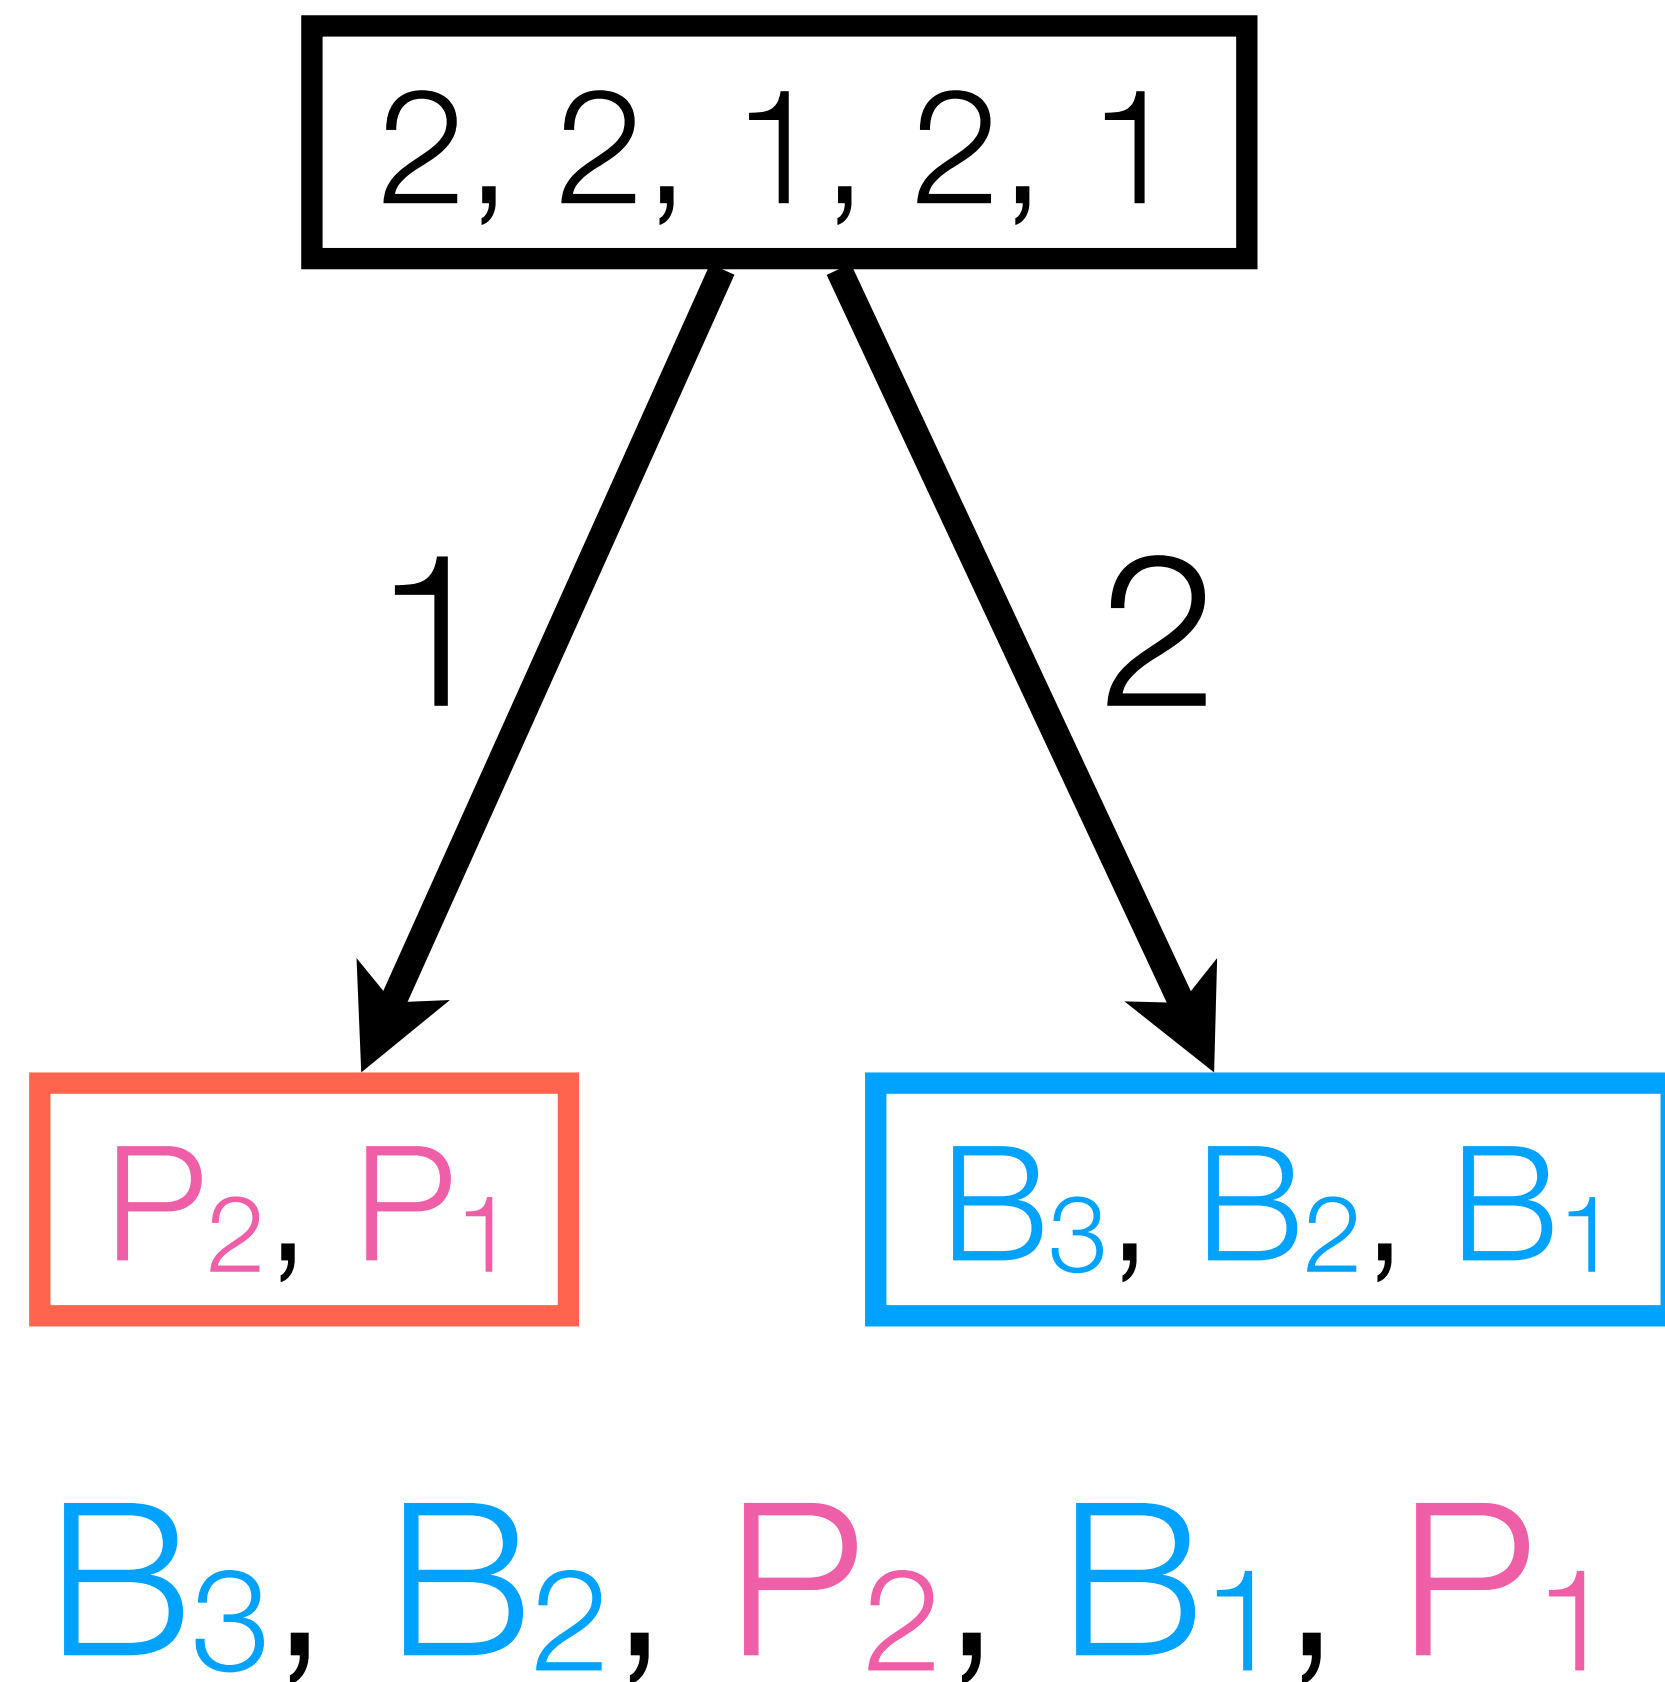


B_3, B_2, P_2, B_1, P_1

This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

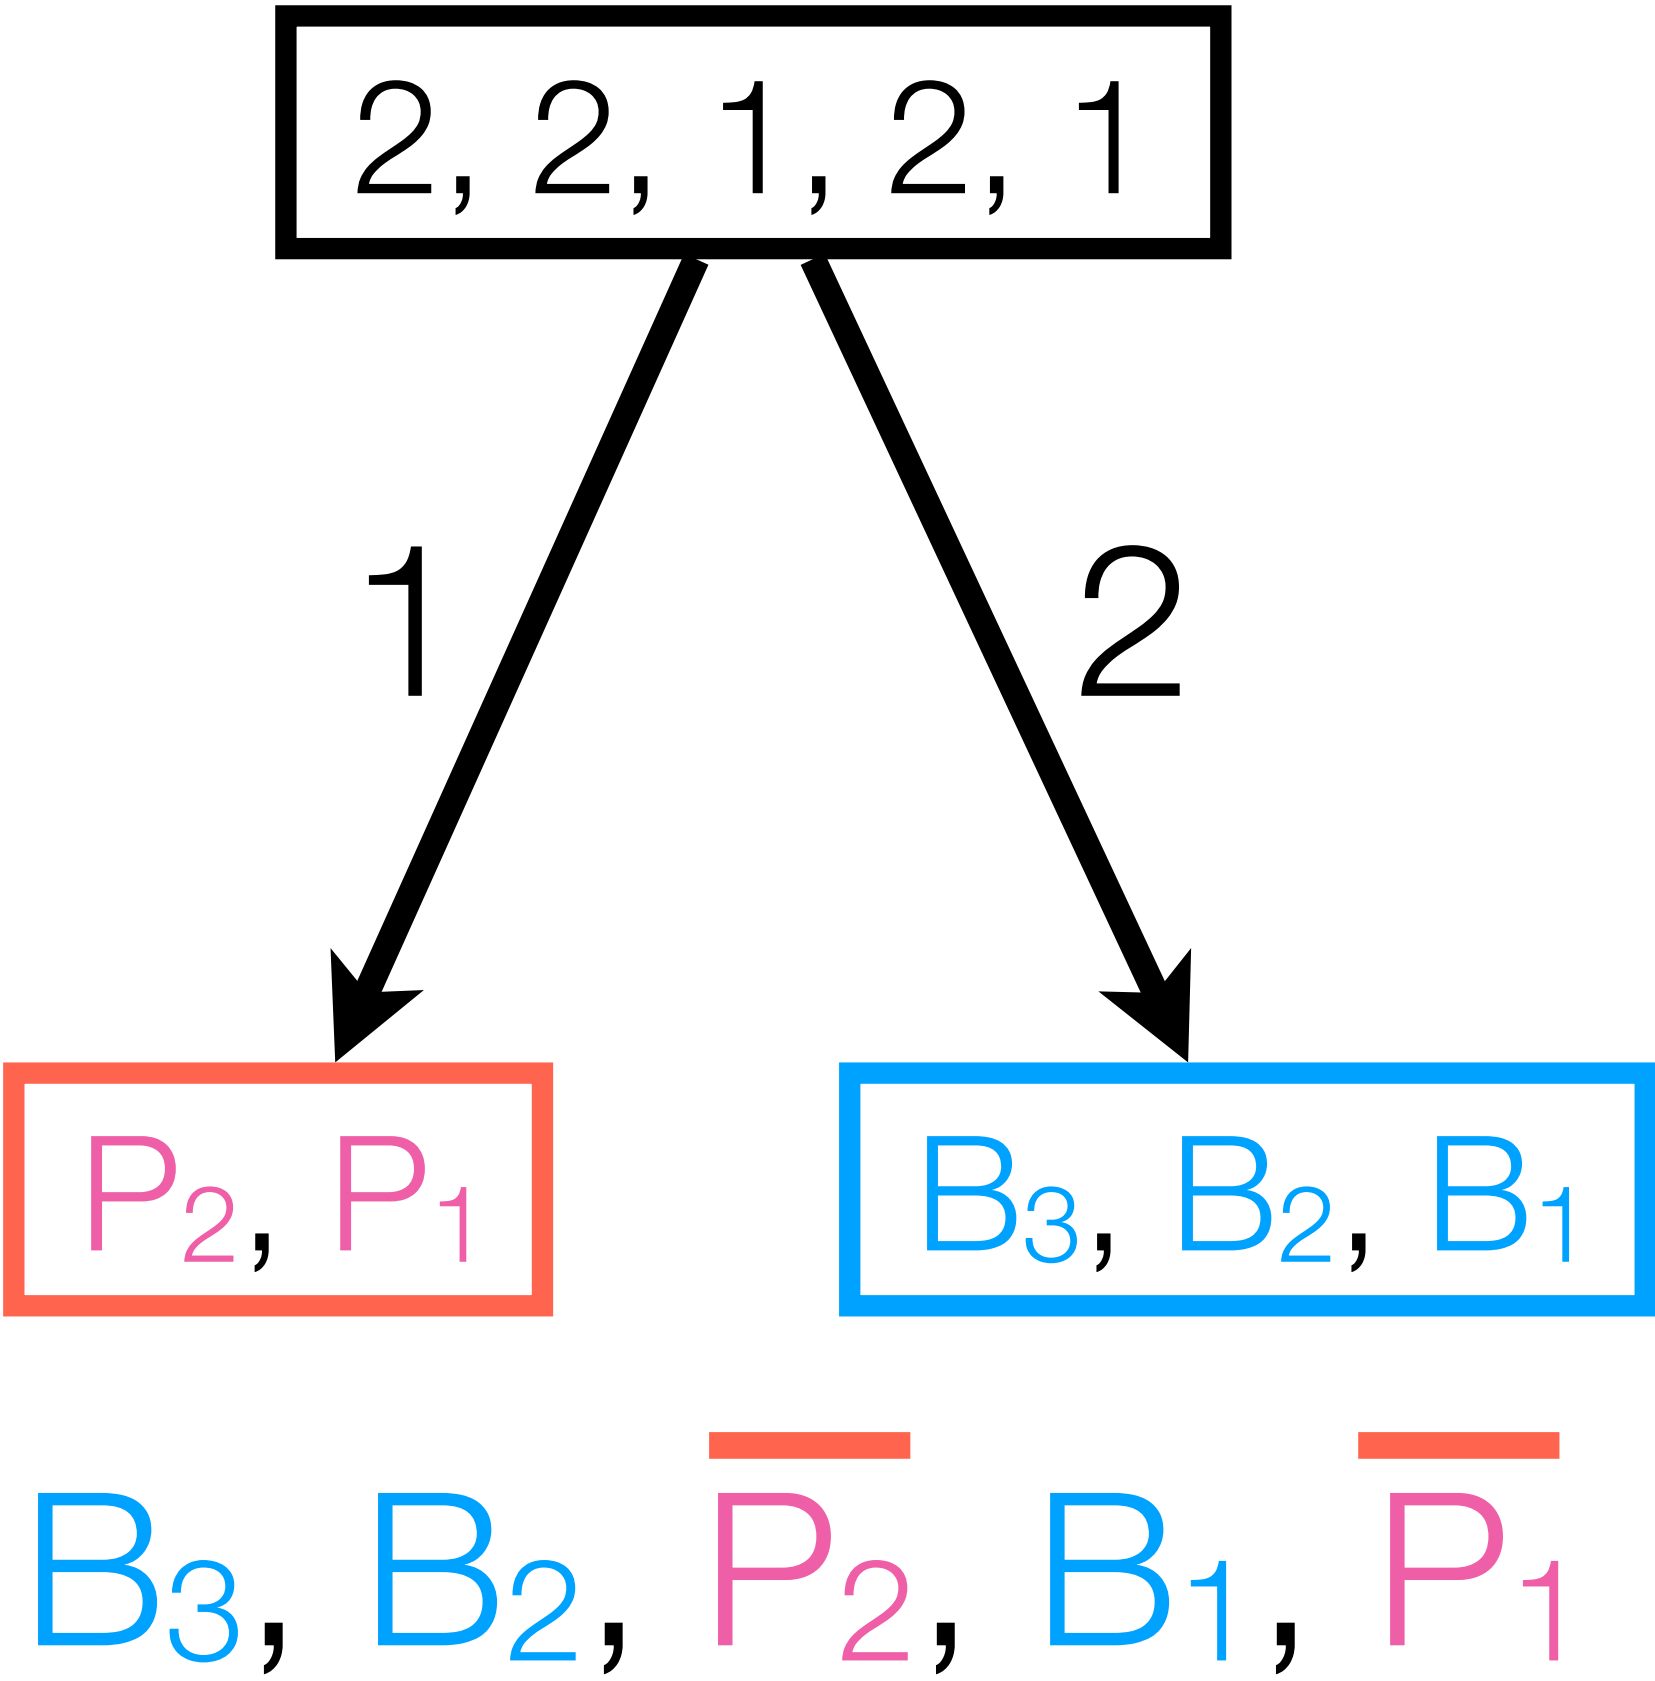
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

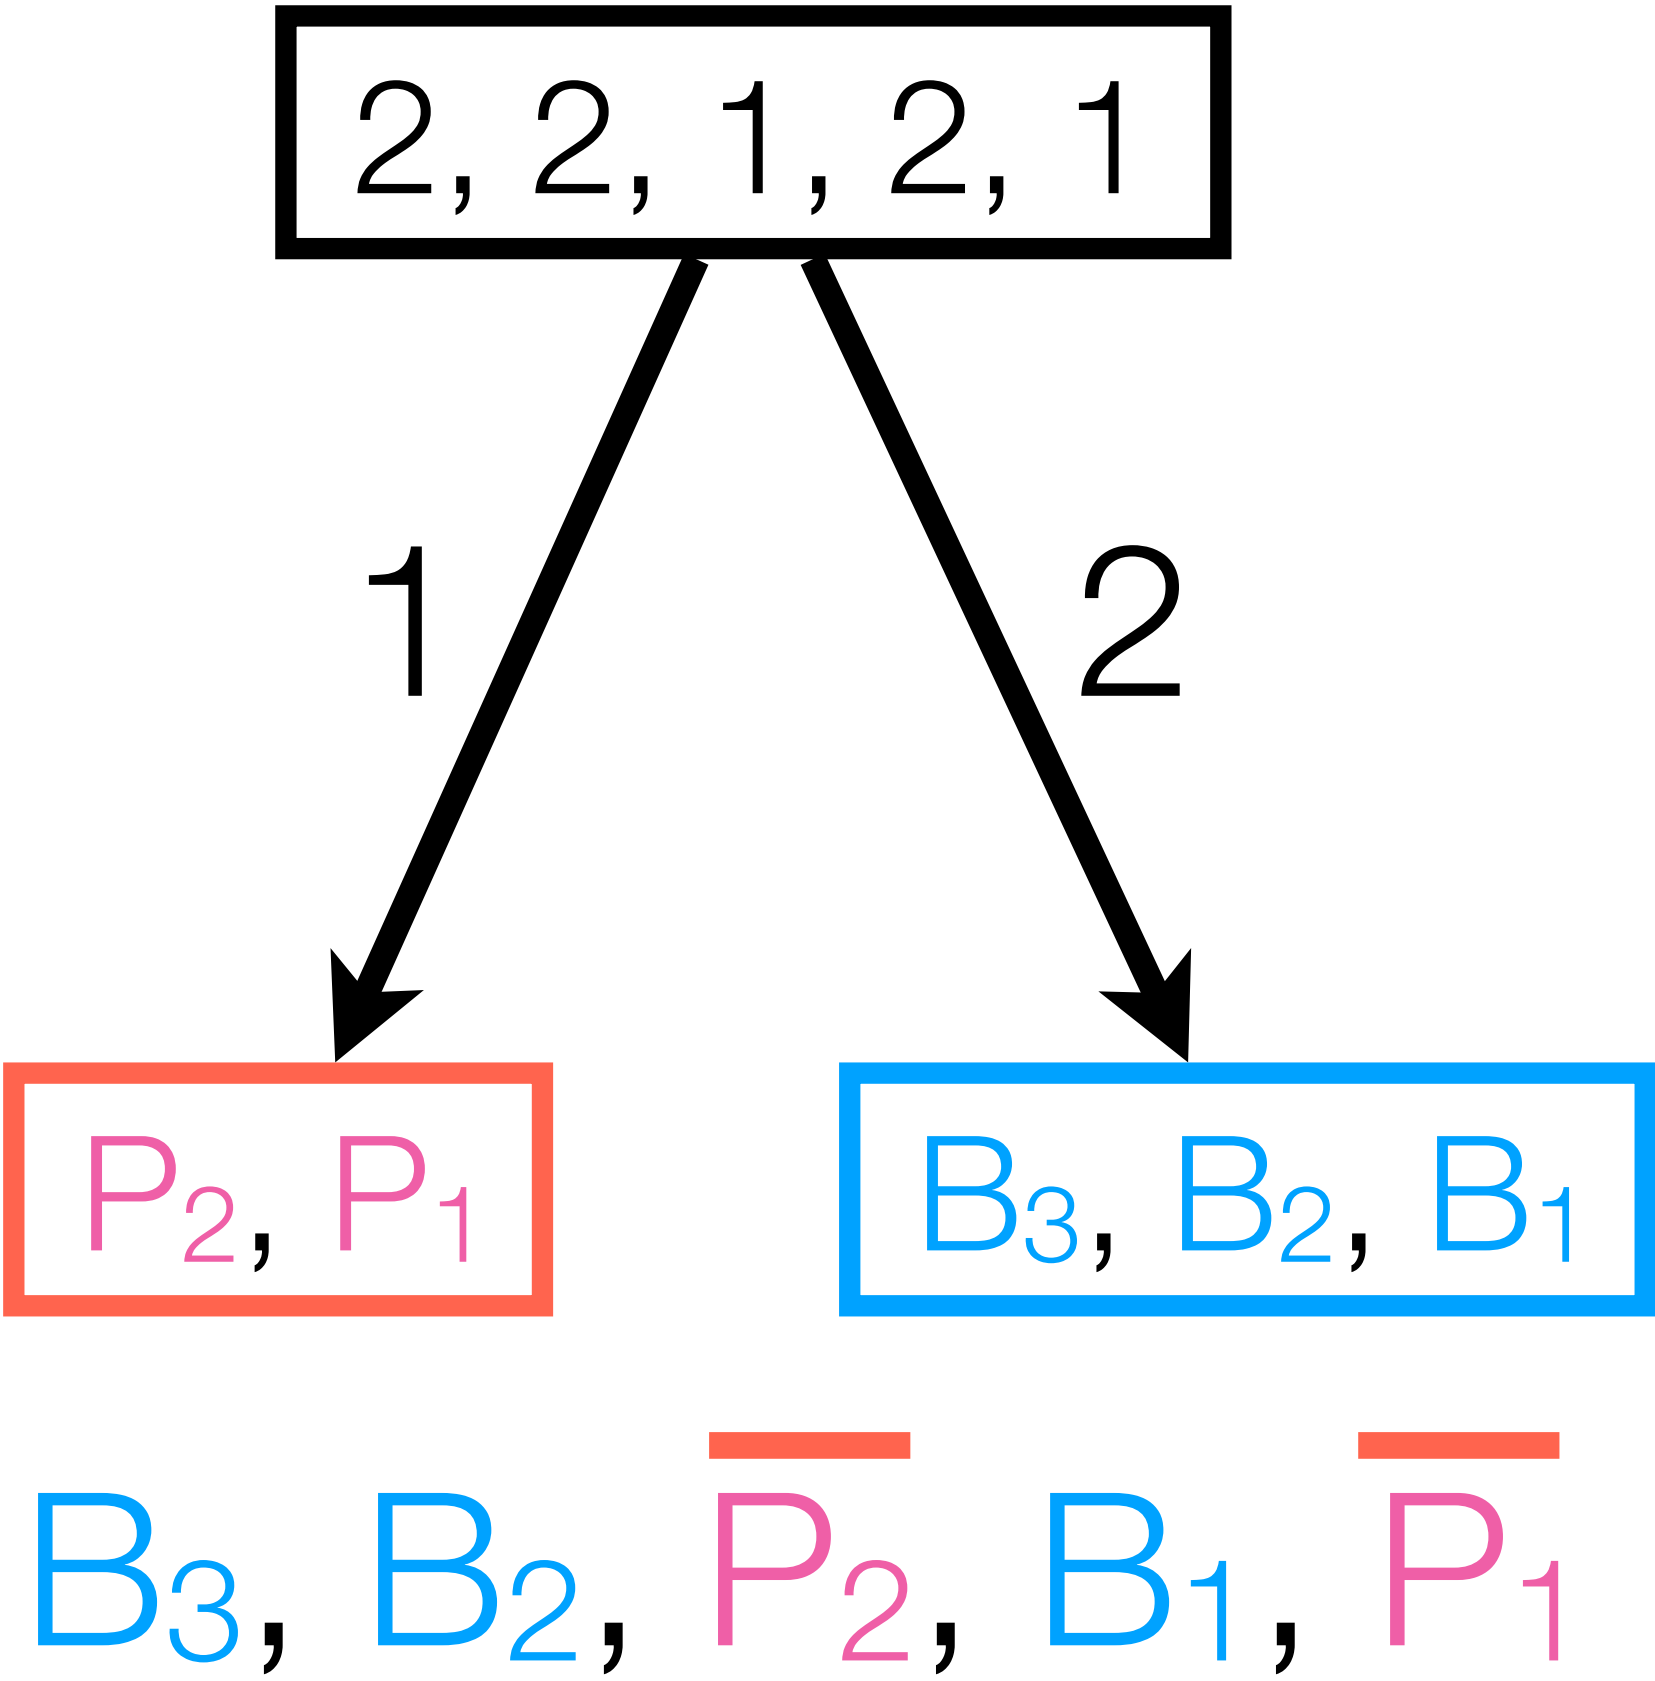
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?

Introducing: PIFO trees

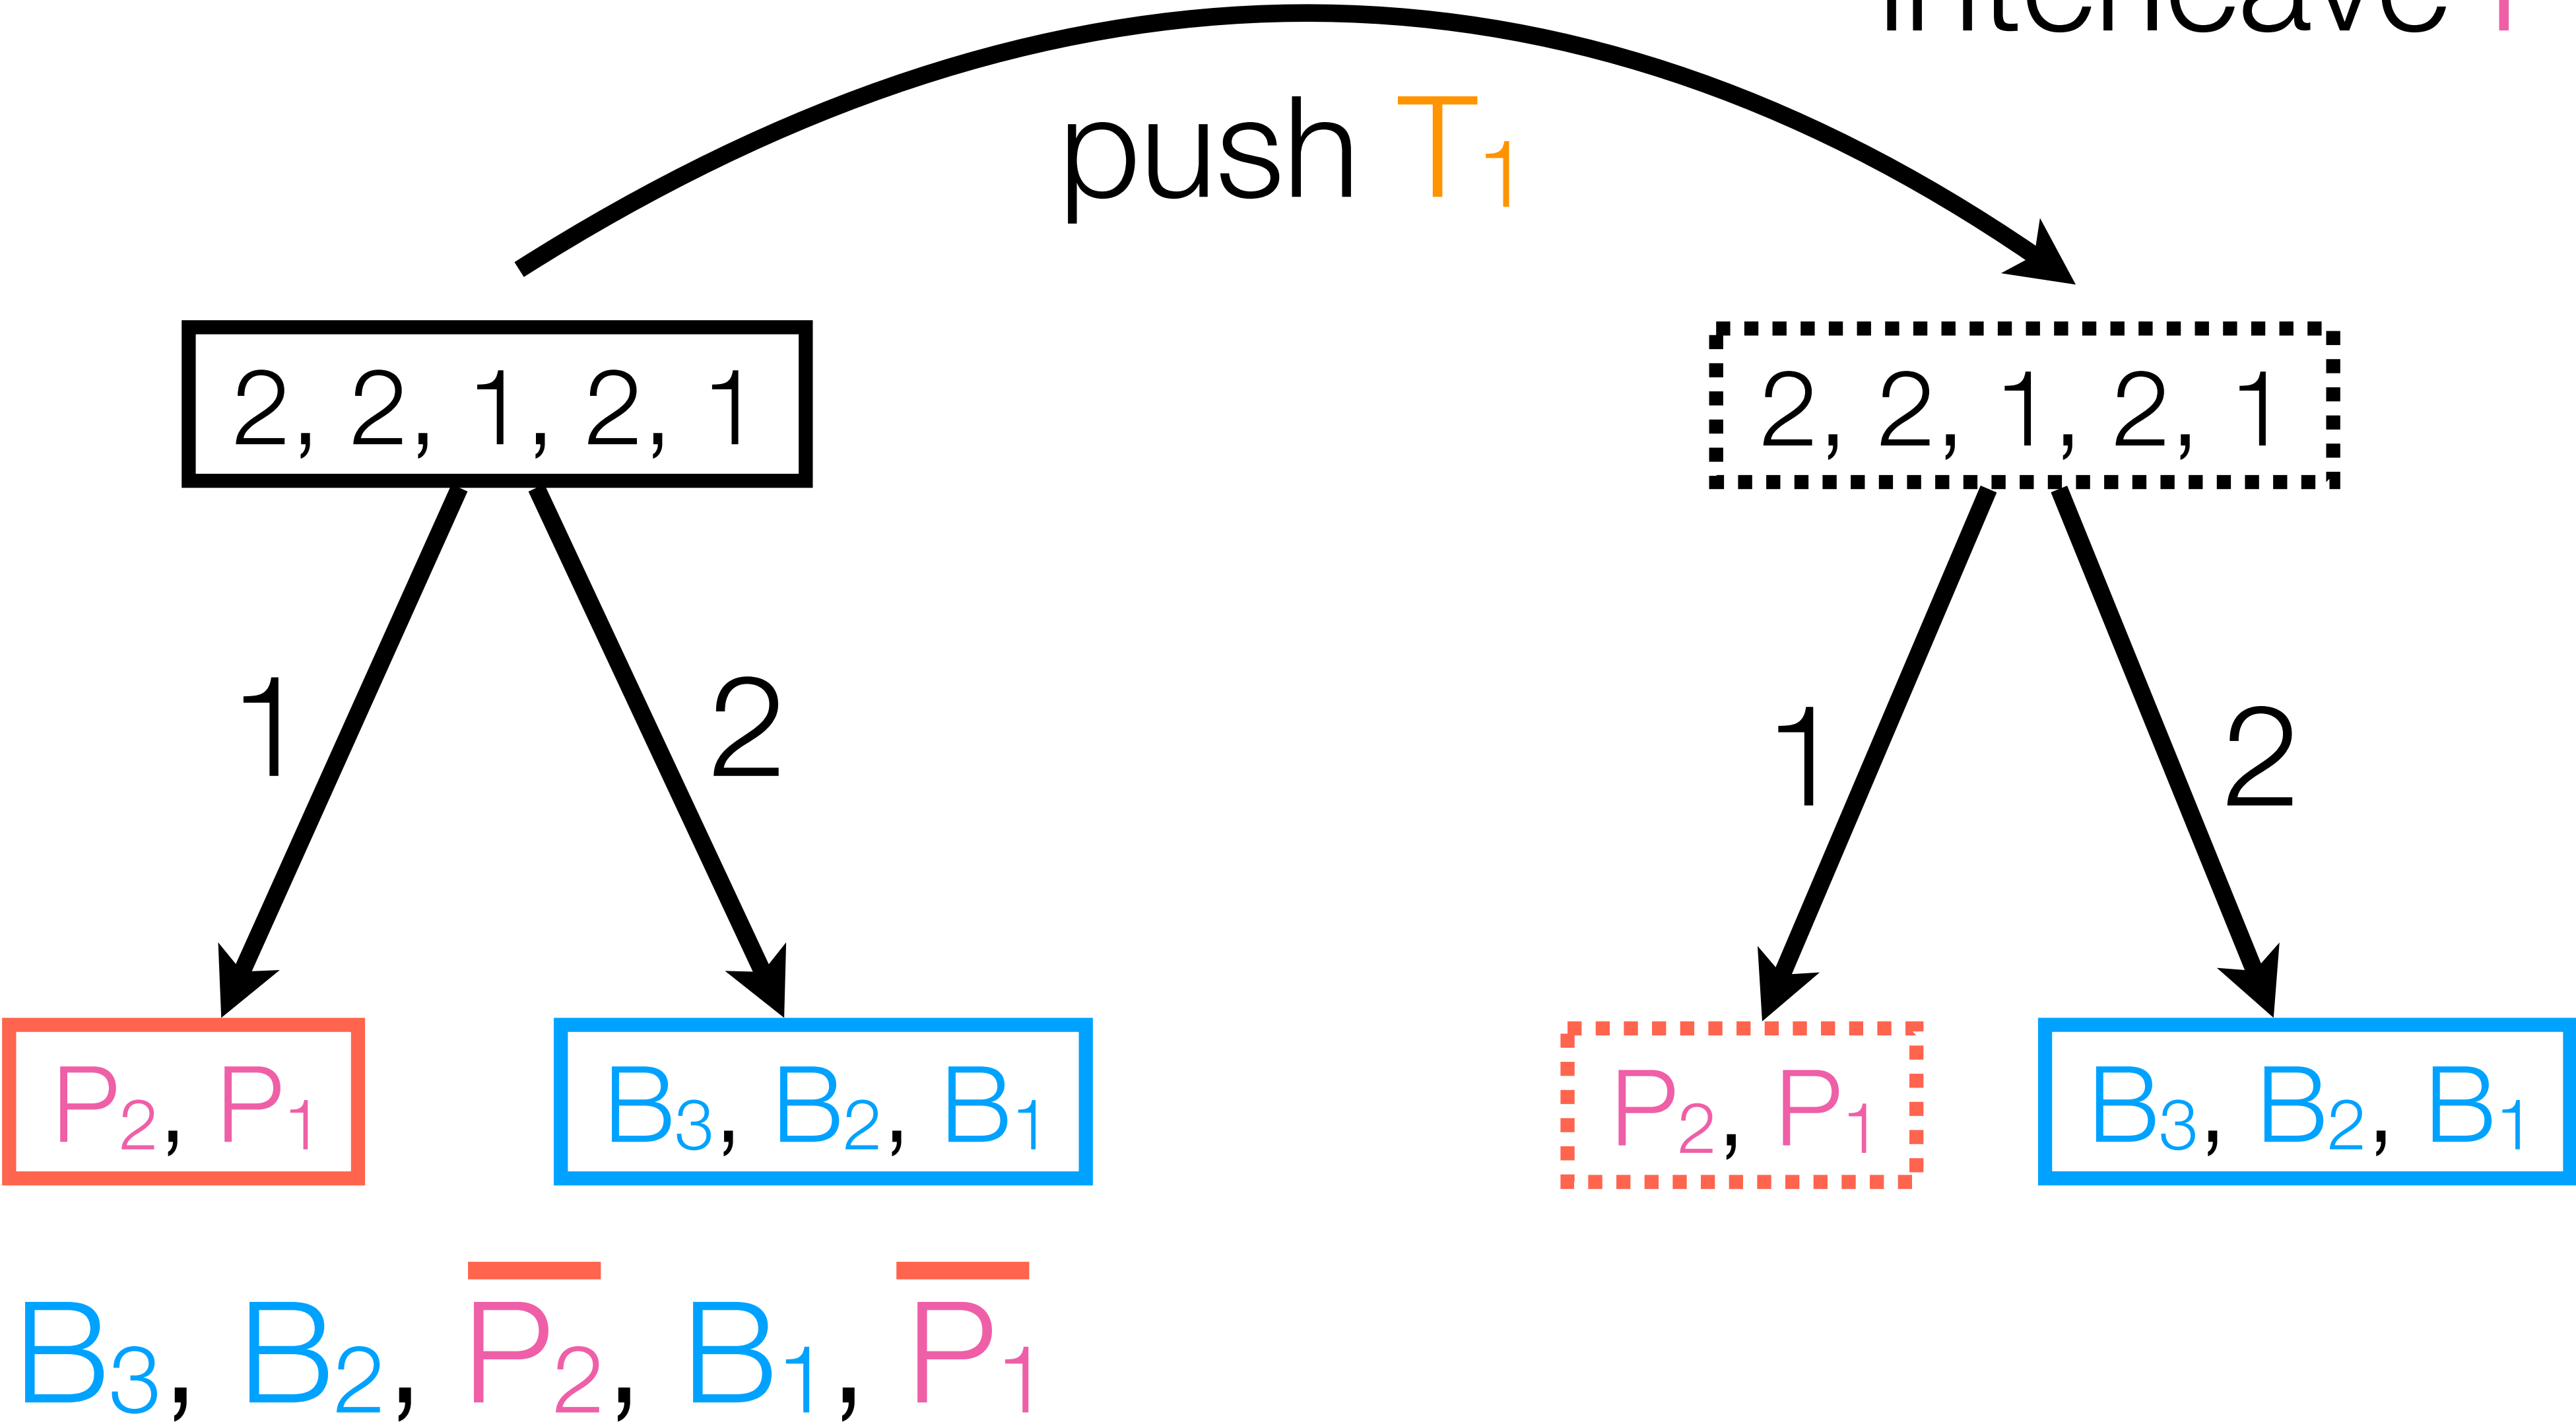
Interleave **R** and **B**;
interleave **P** and **T**.



This behaves like a queue!
How do we pop it?
How do we push into it?

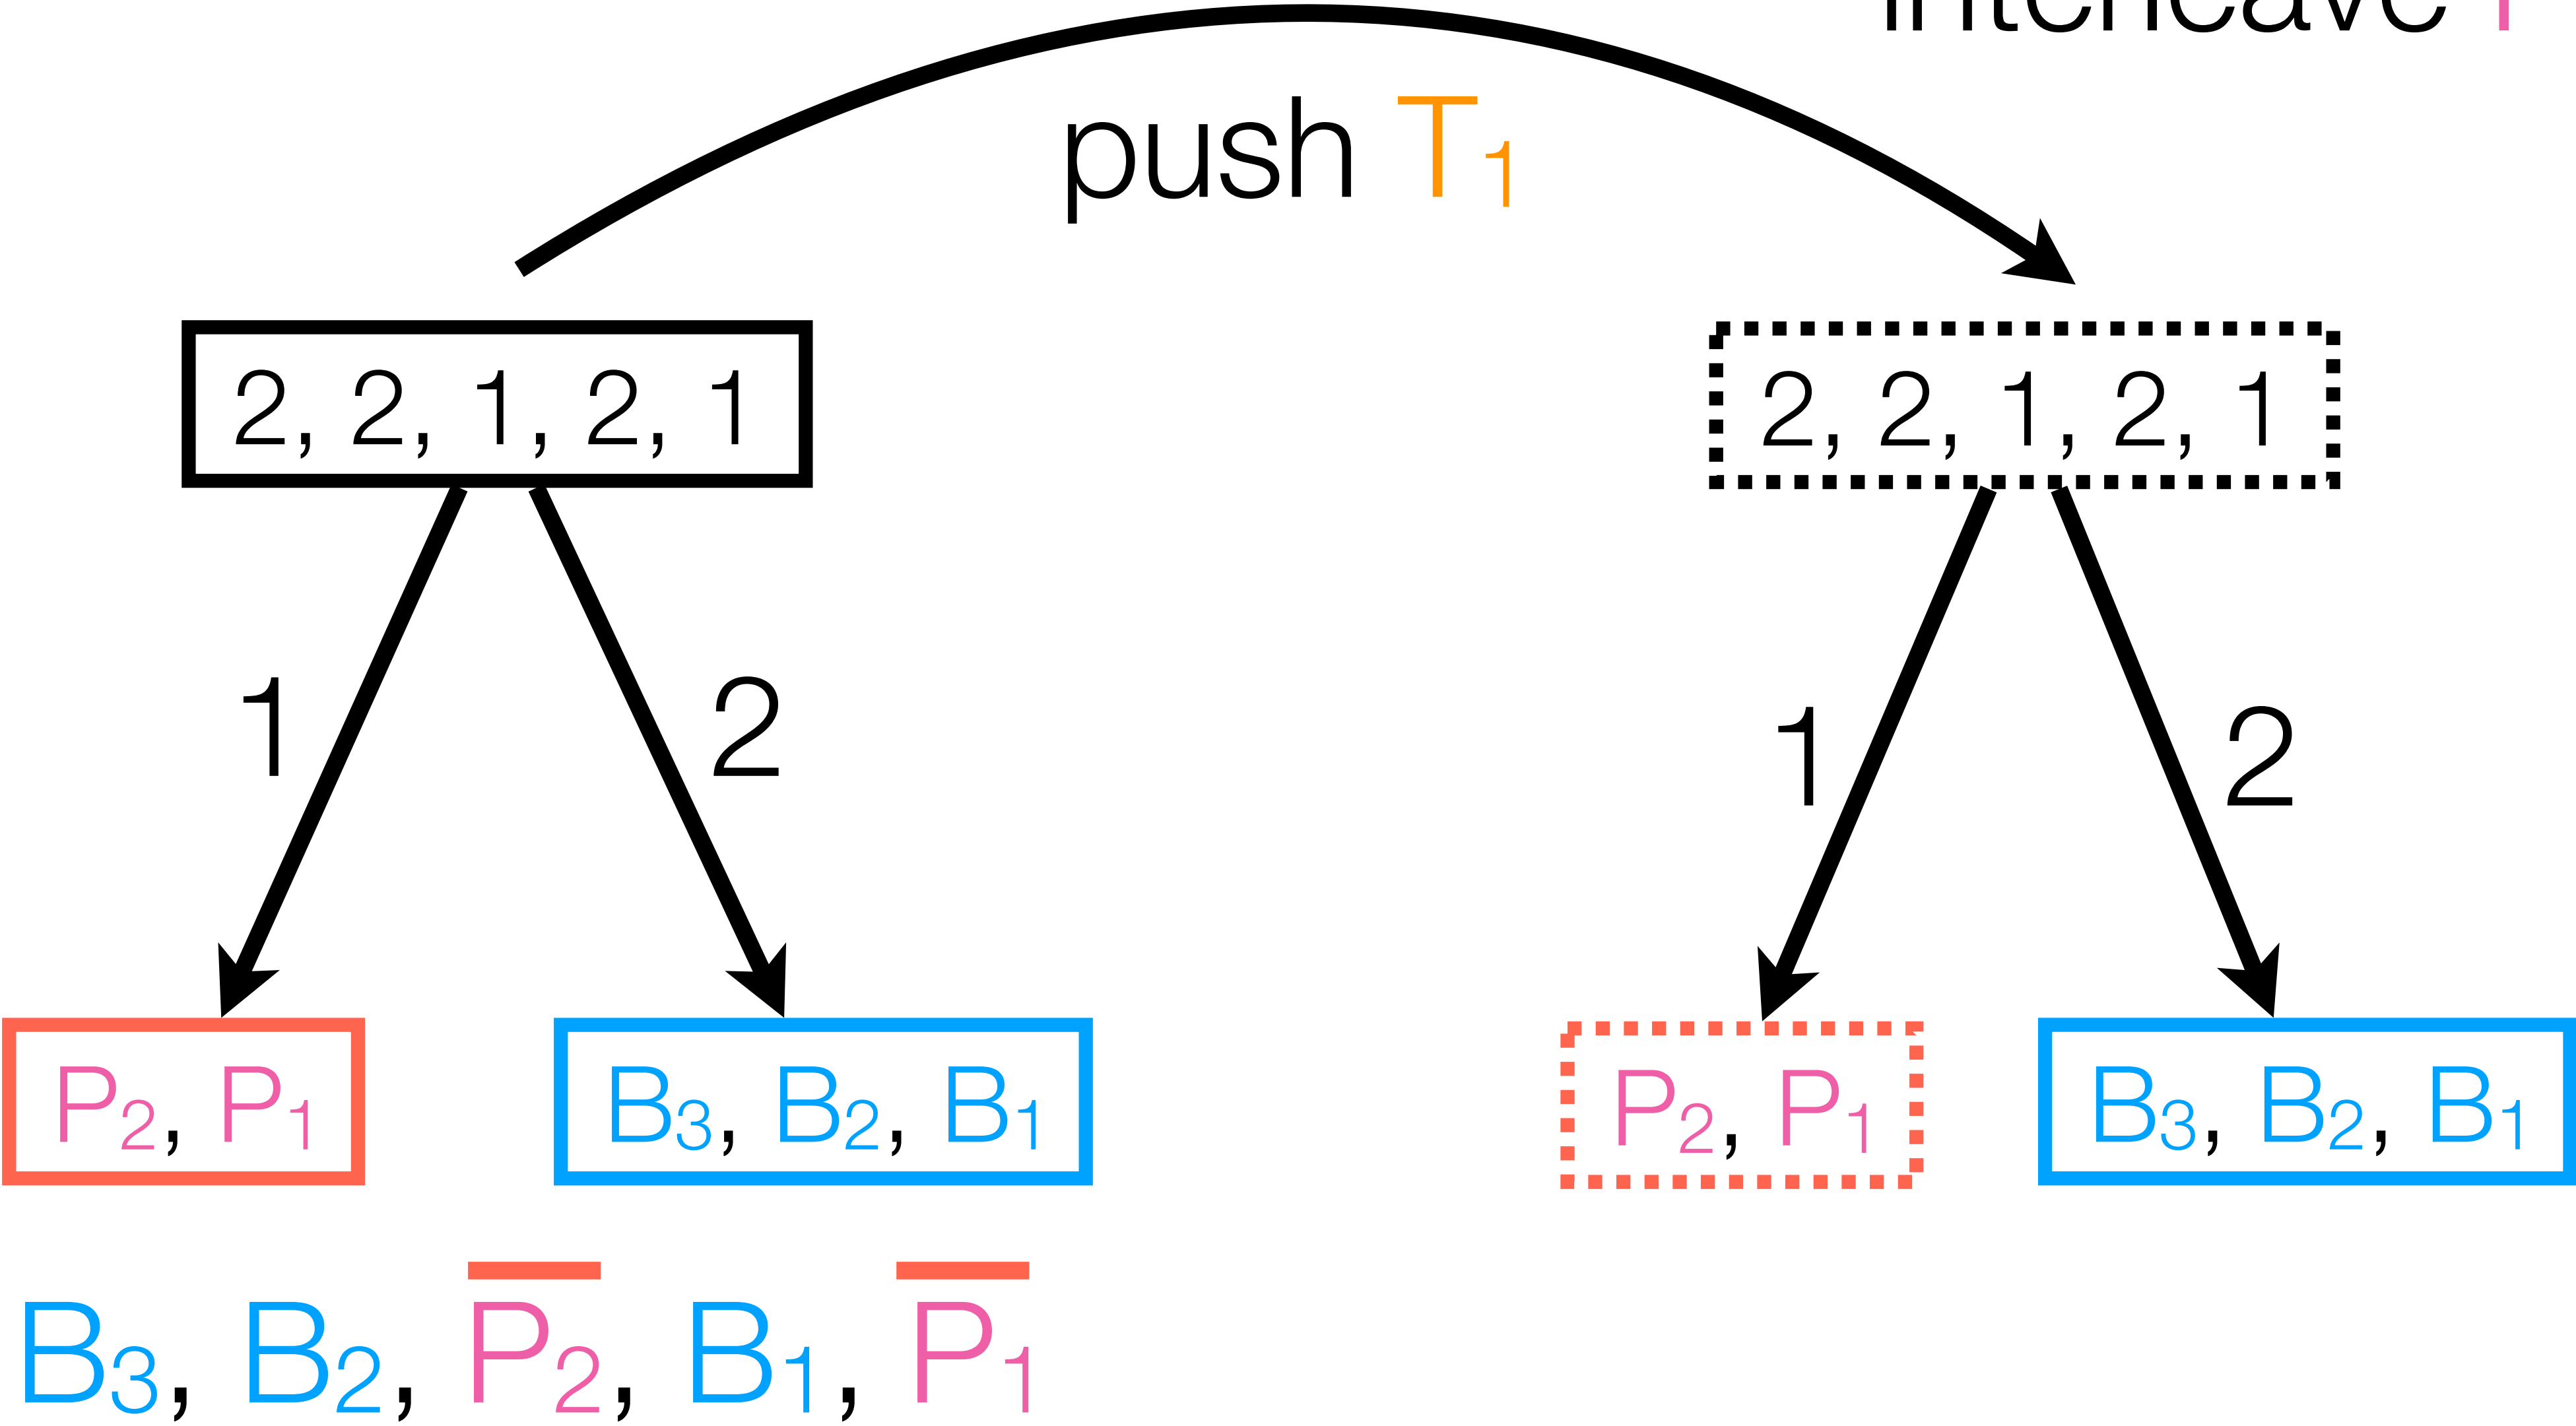
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



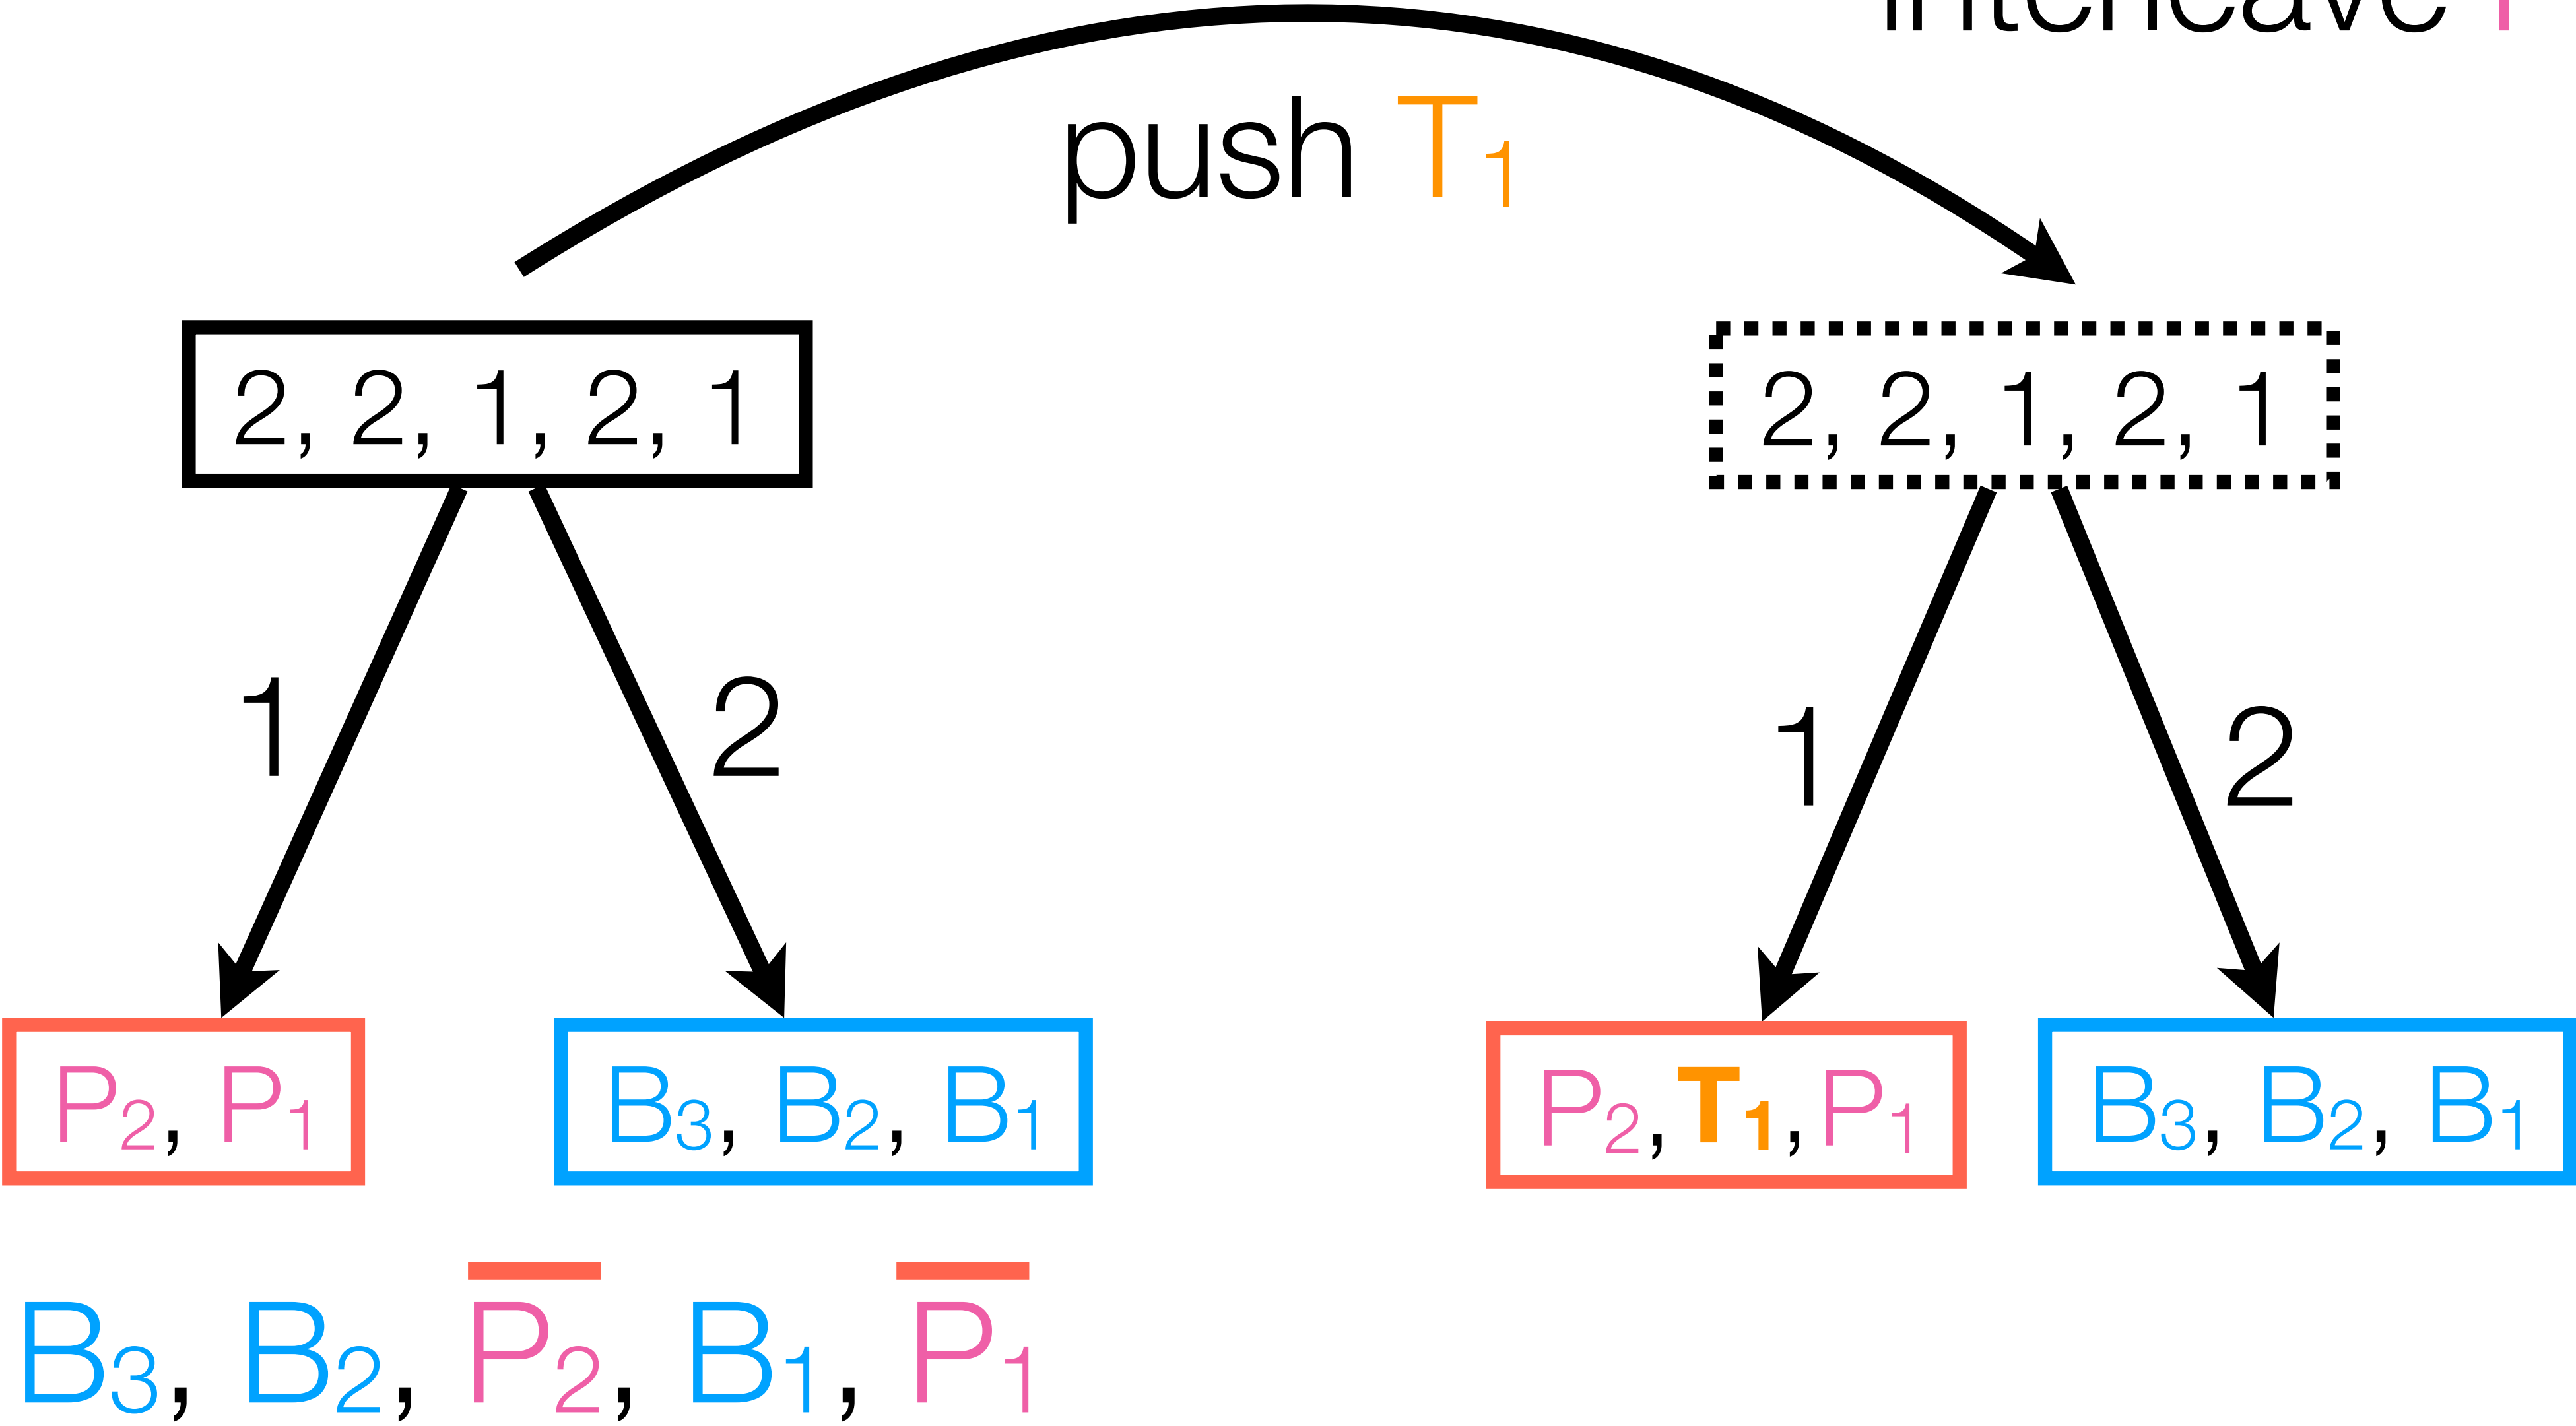
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



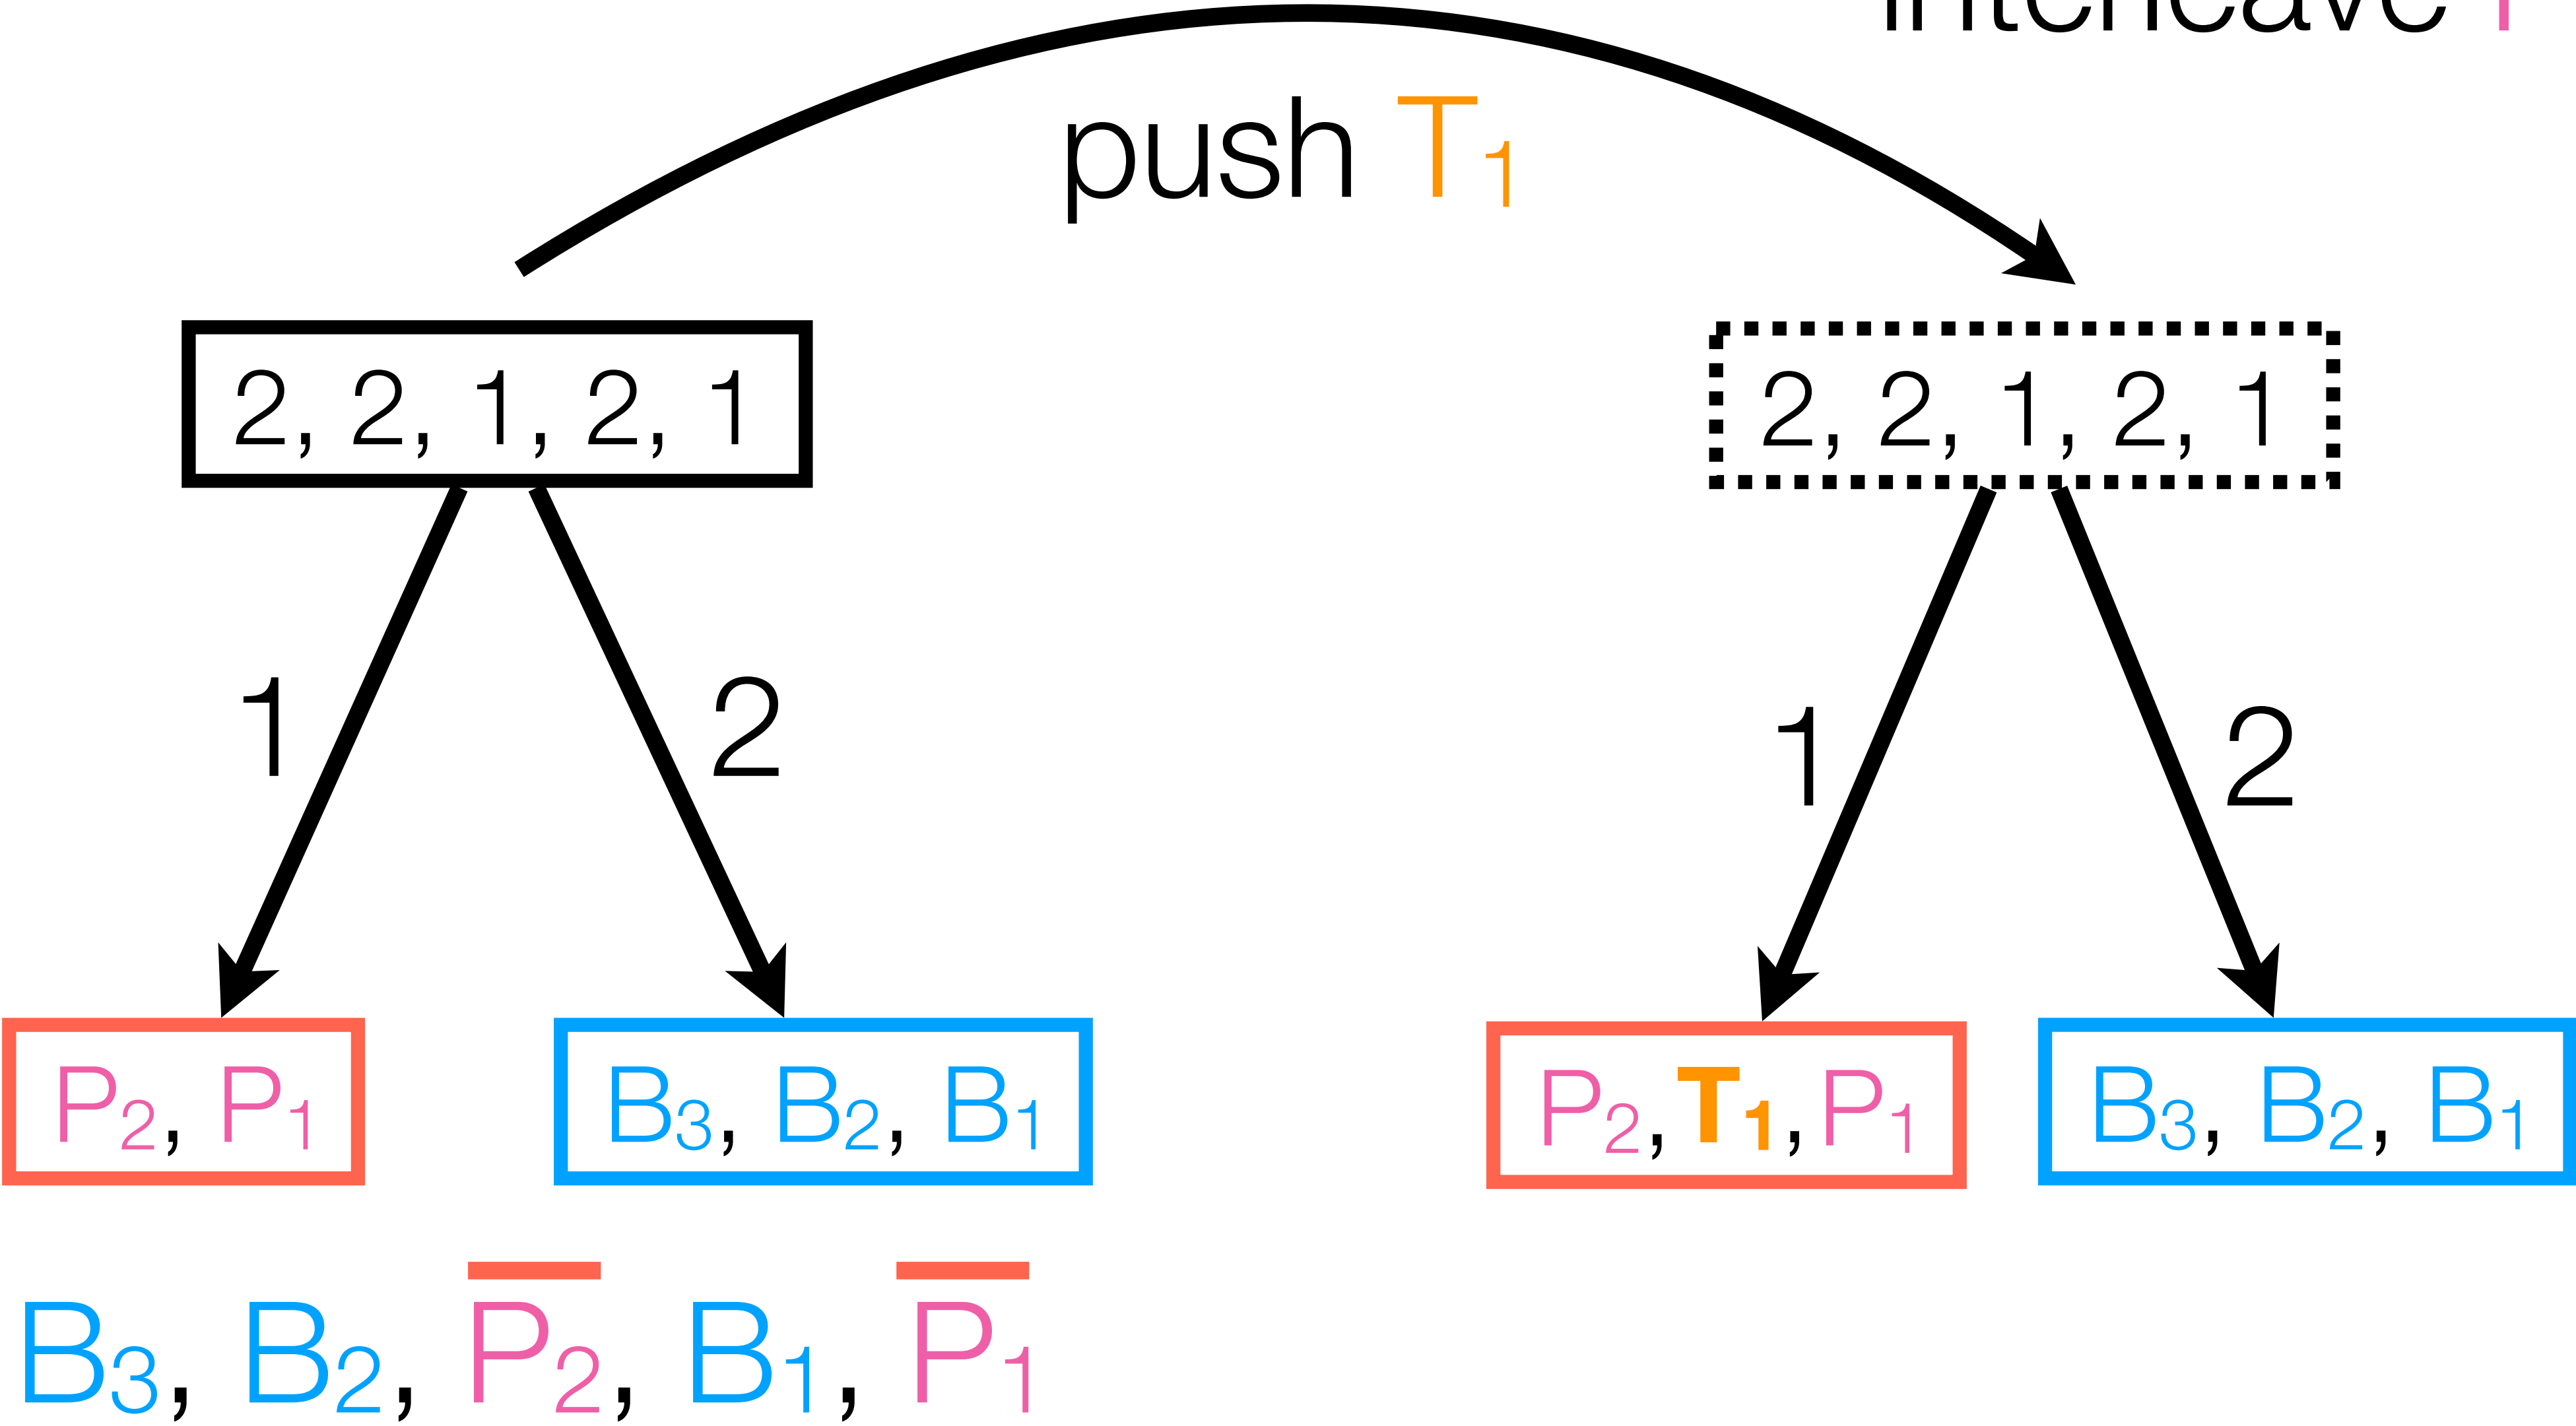
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



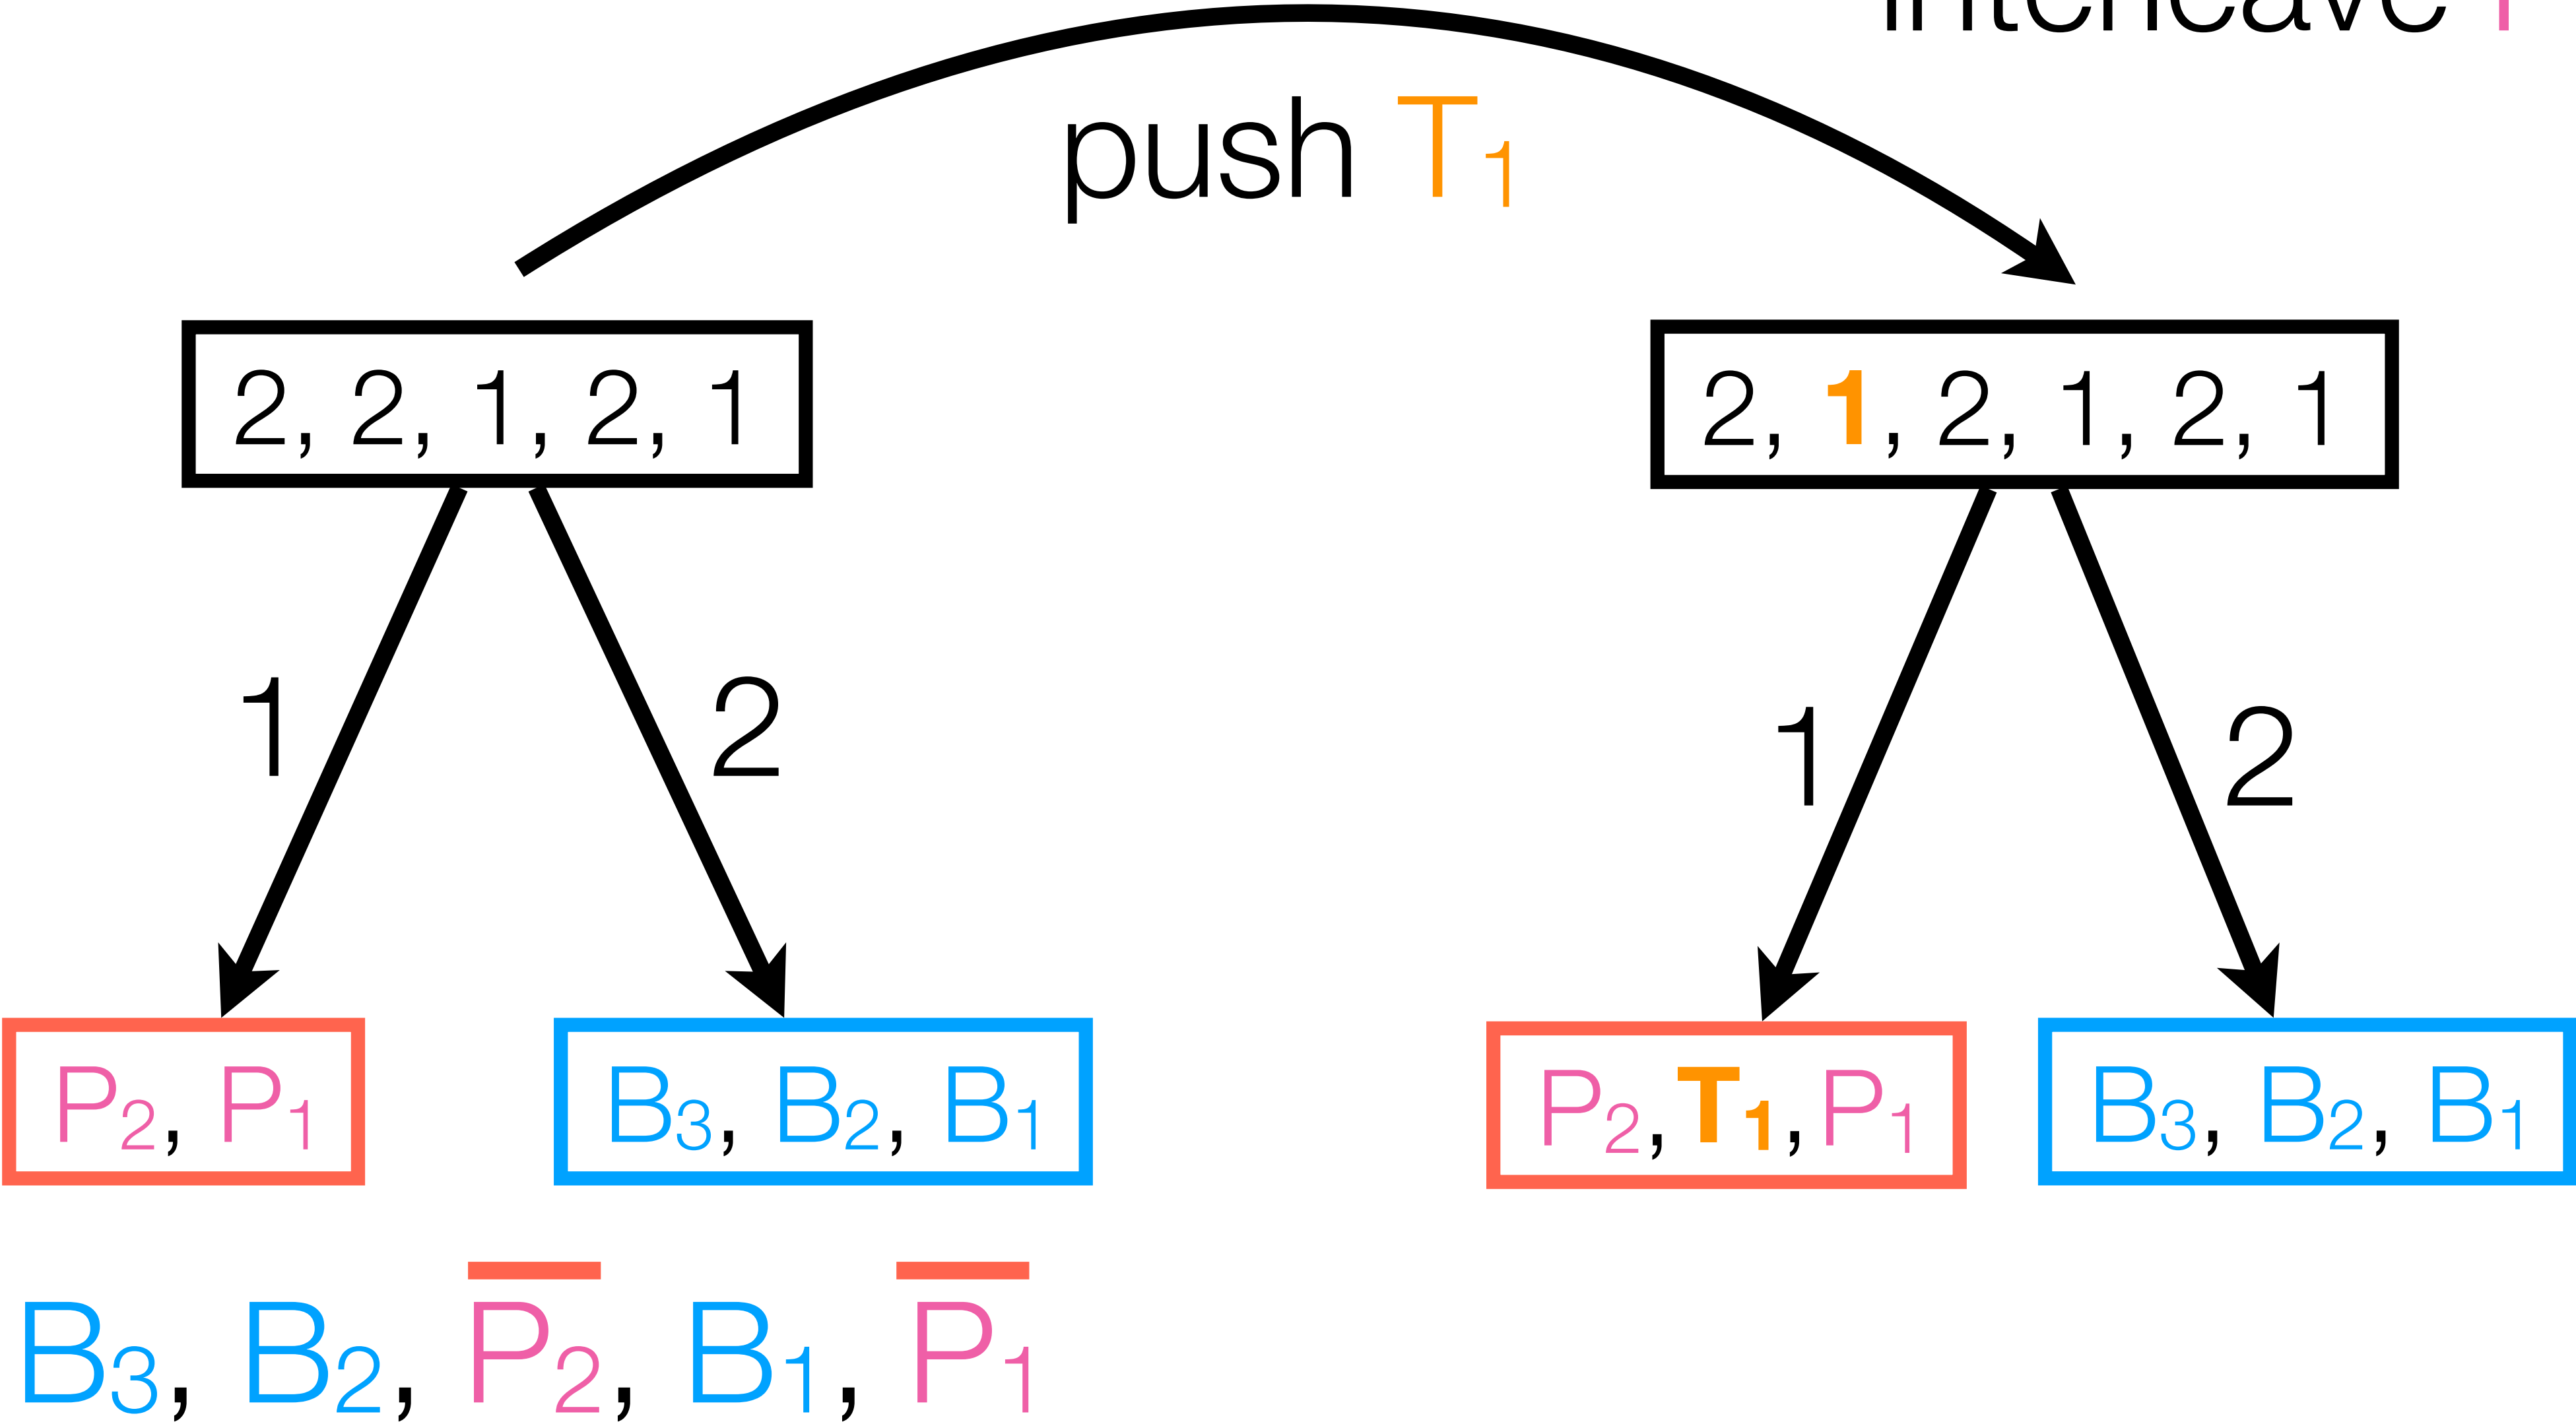
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



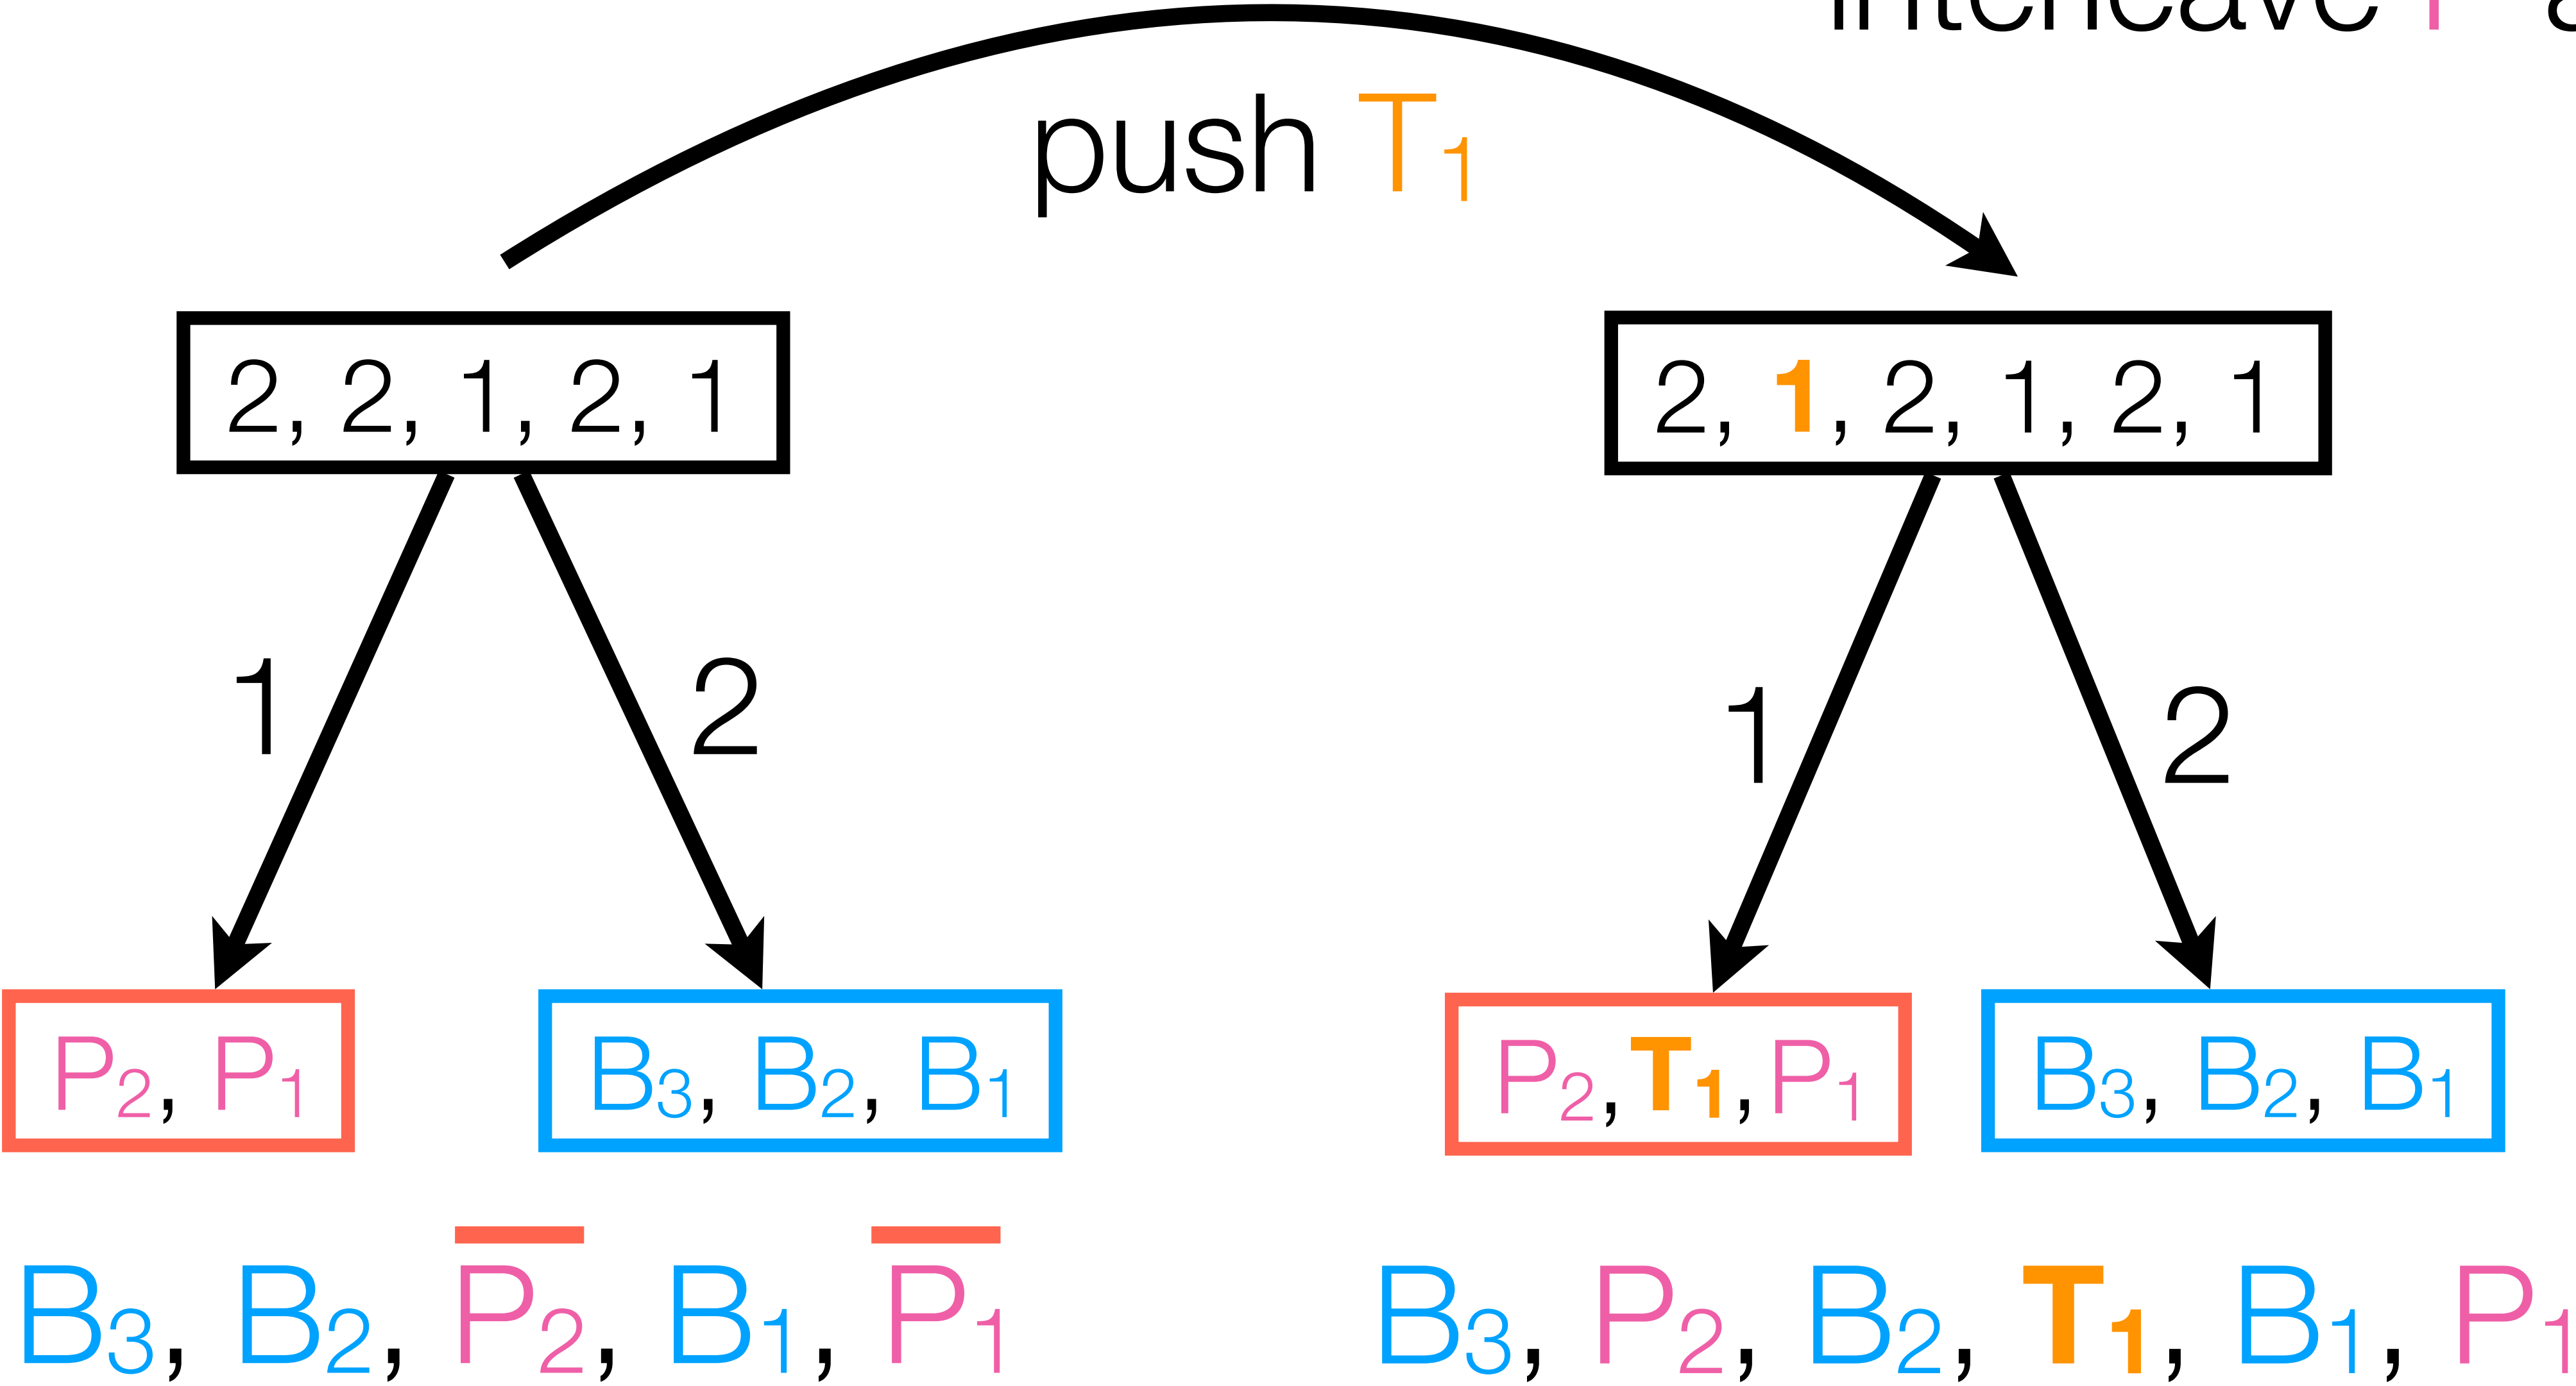
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



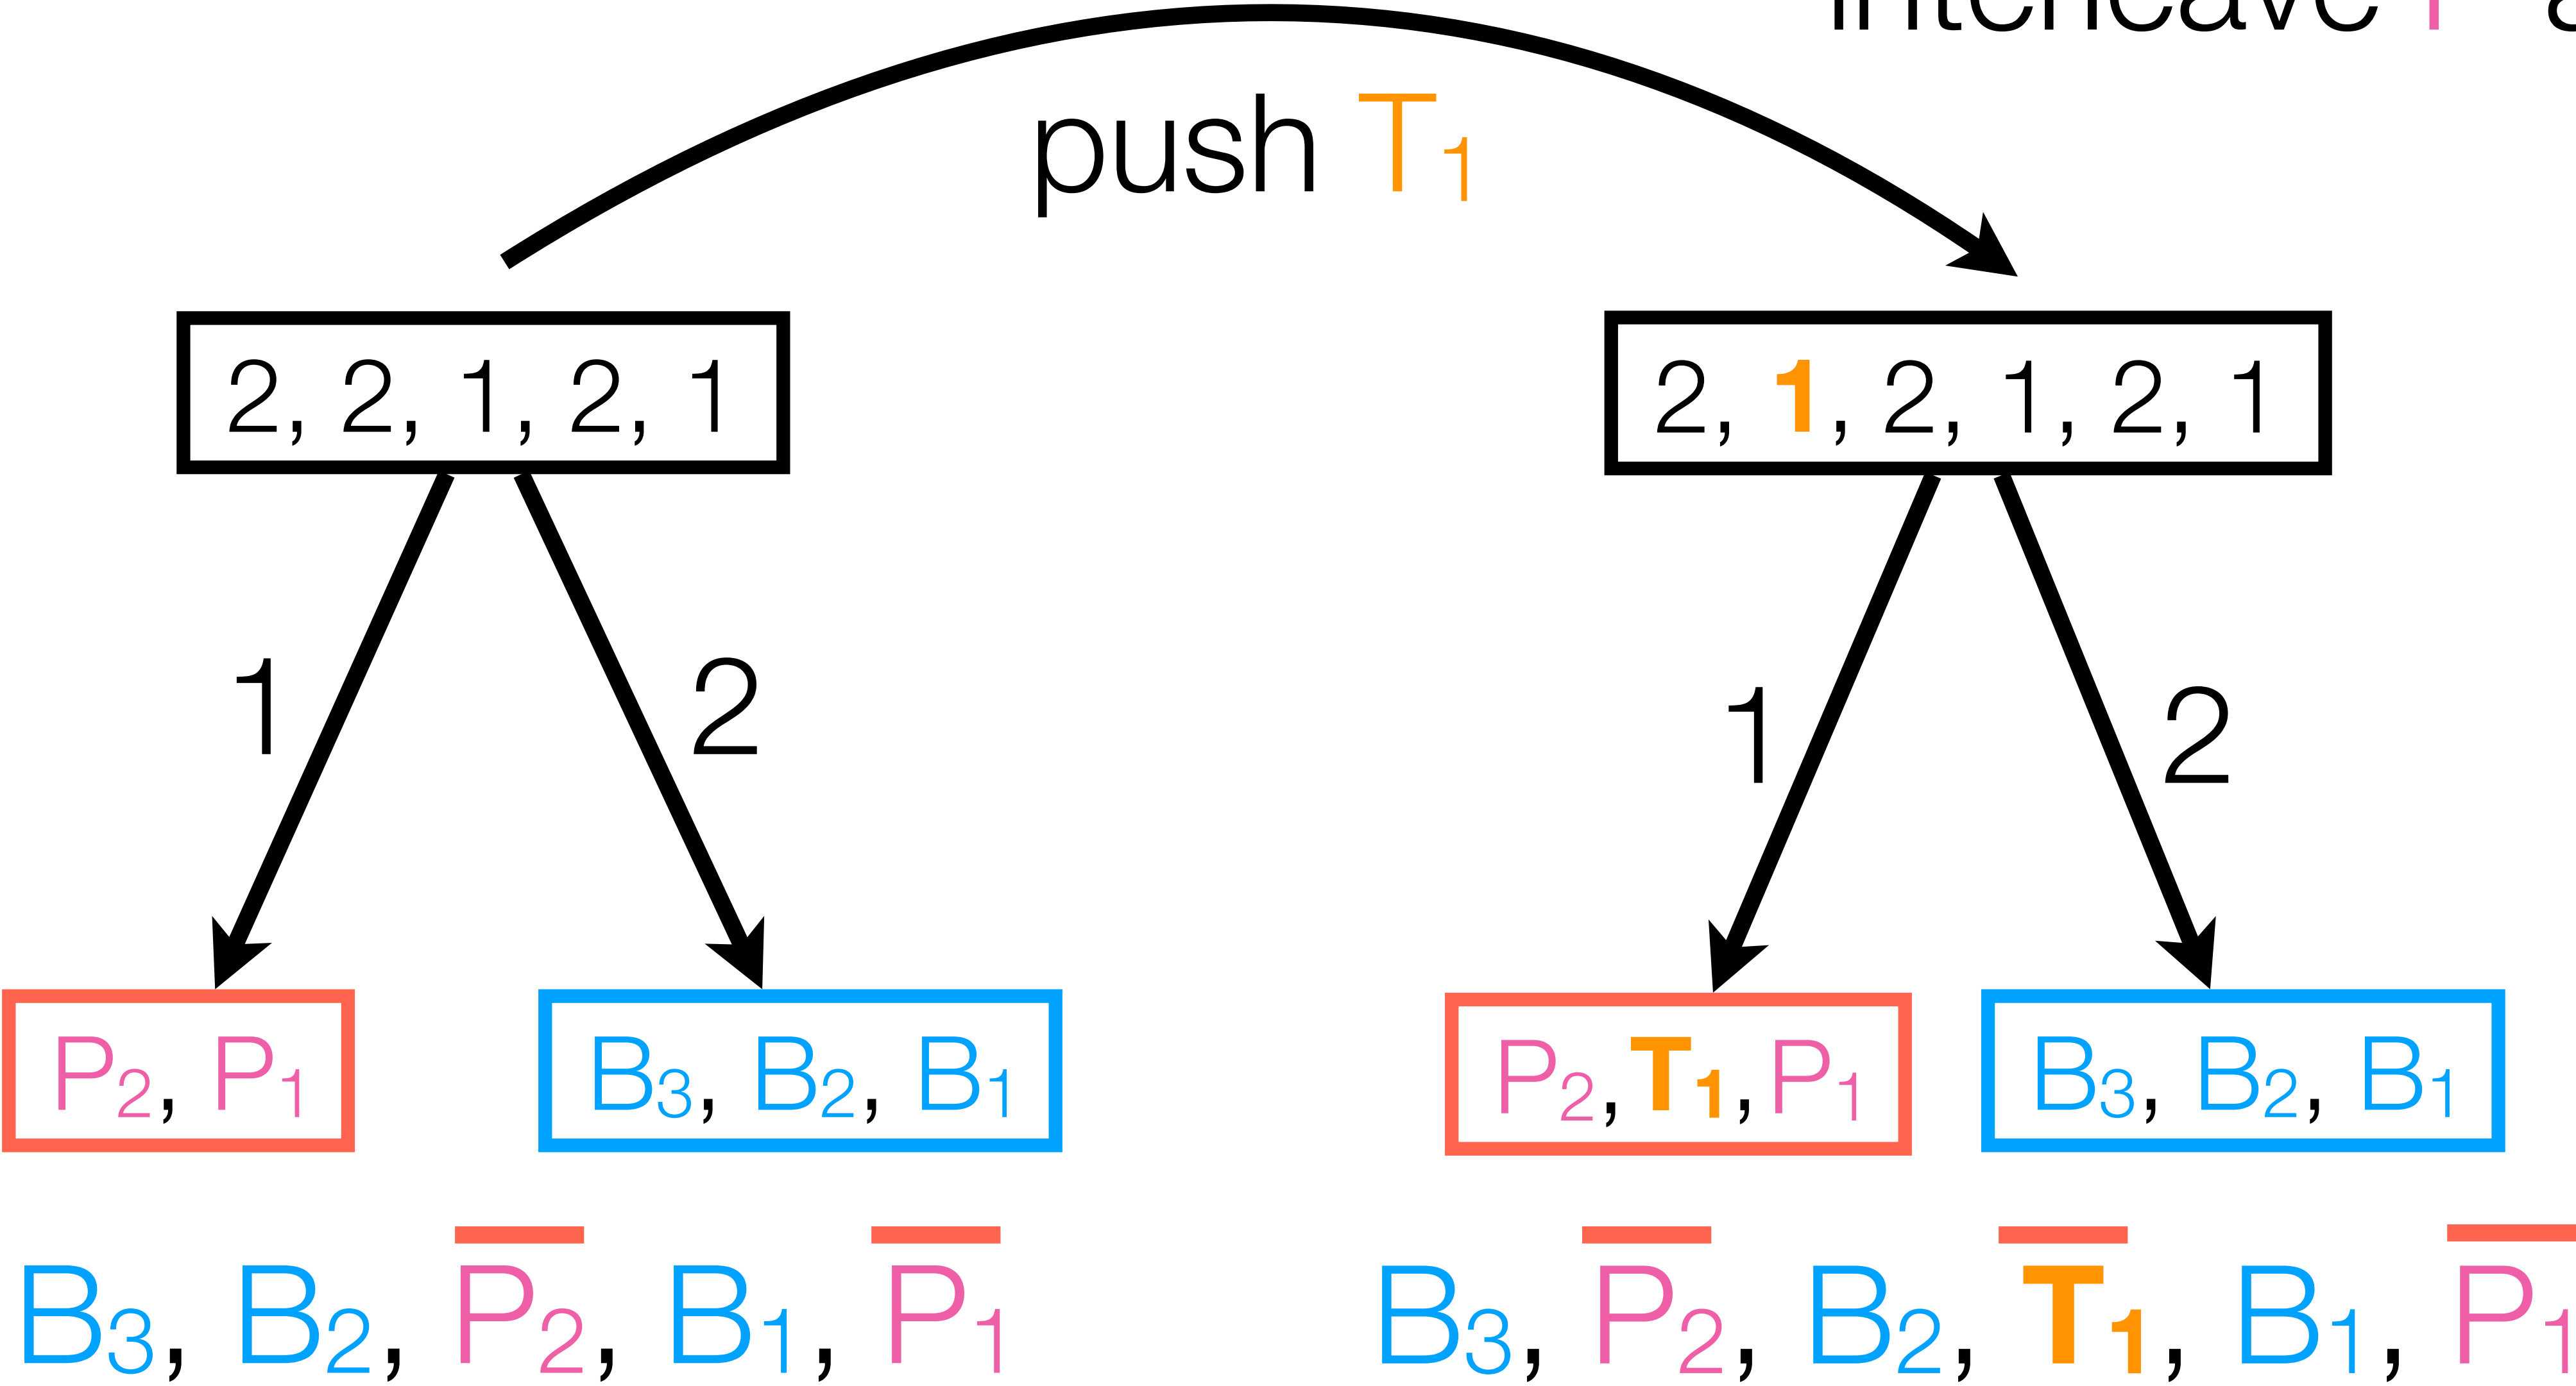
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



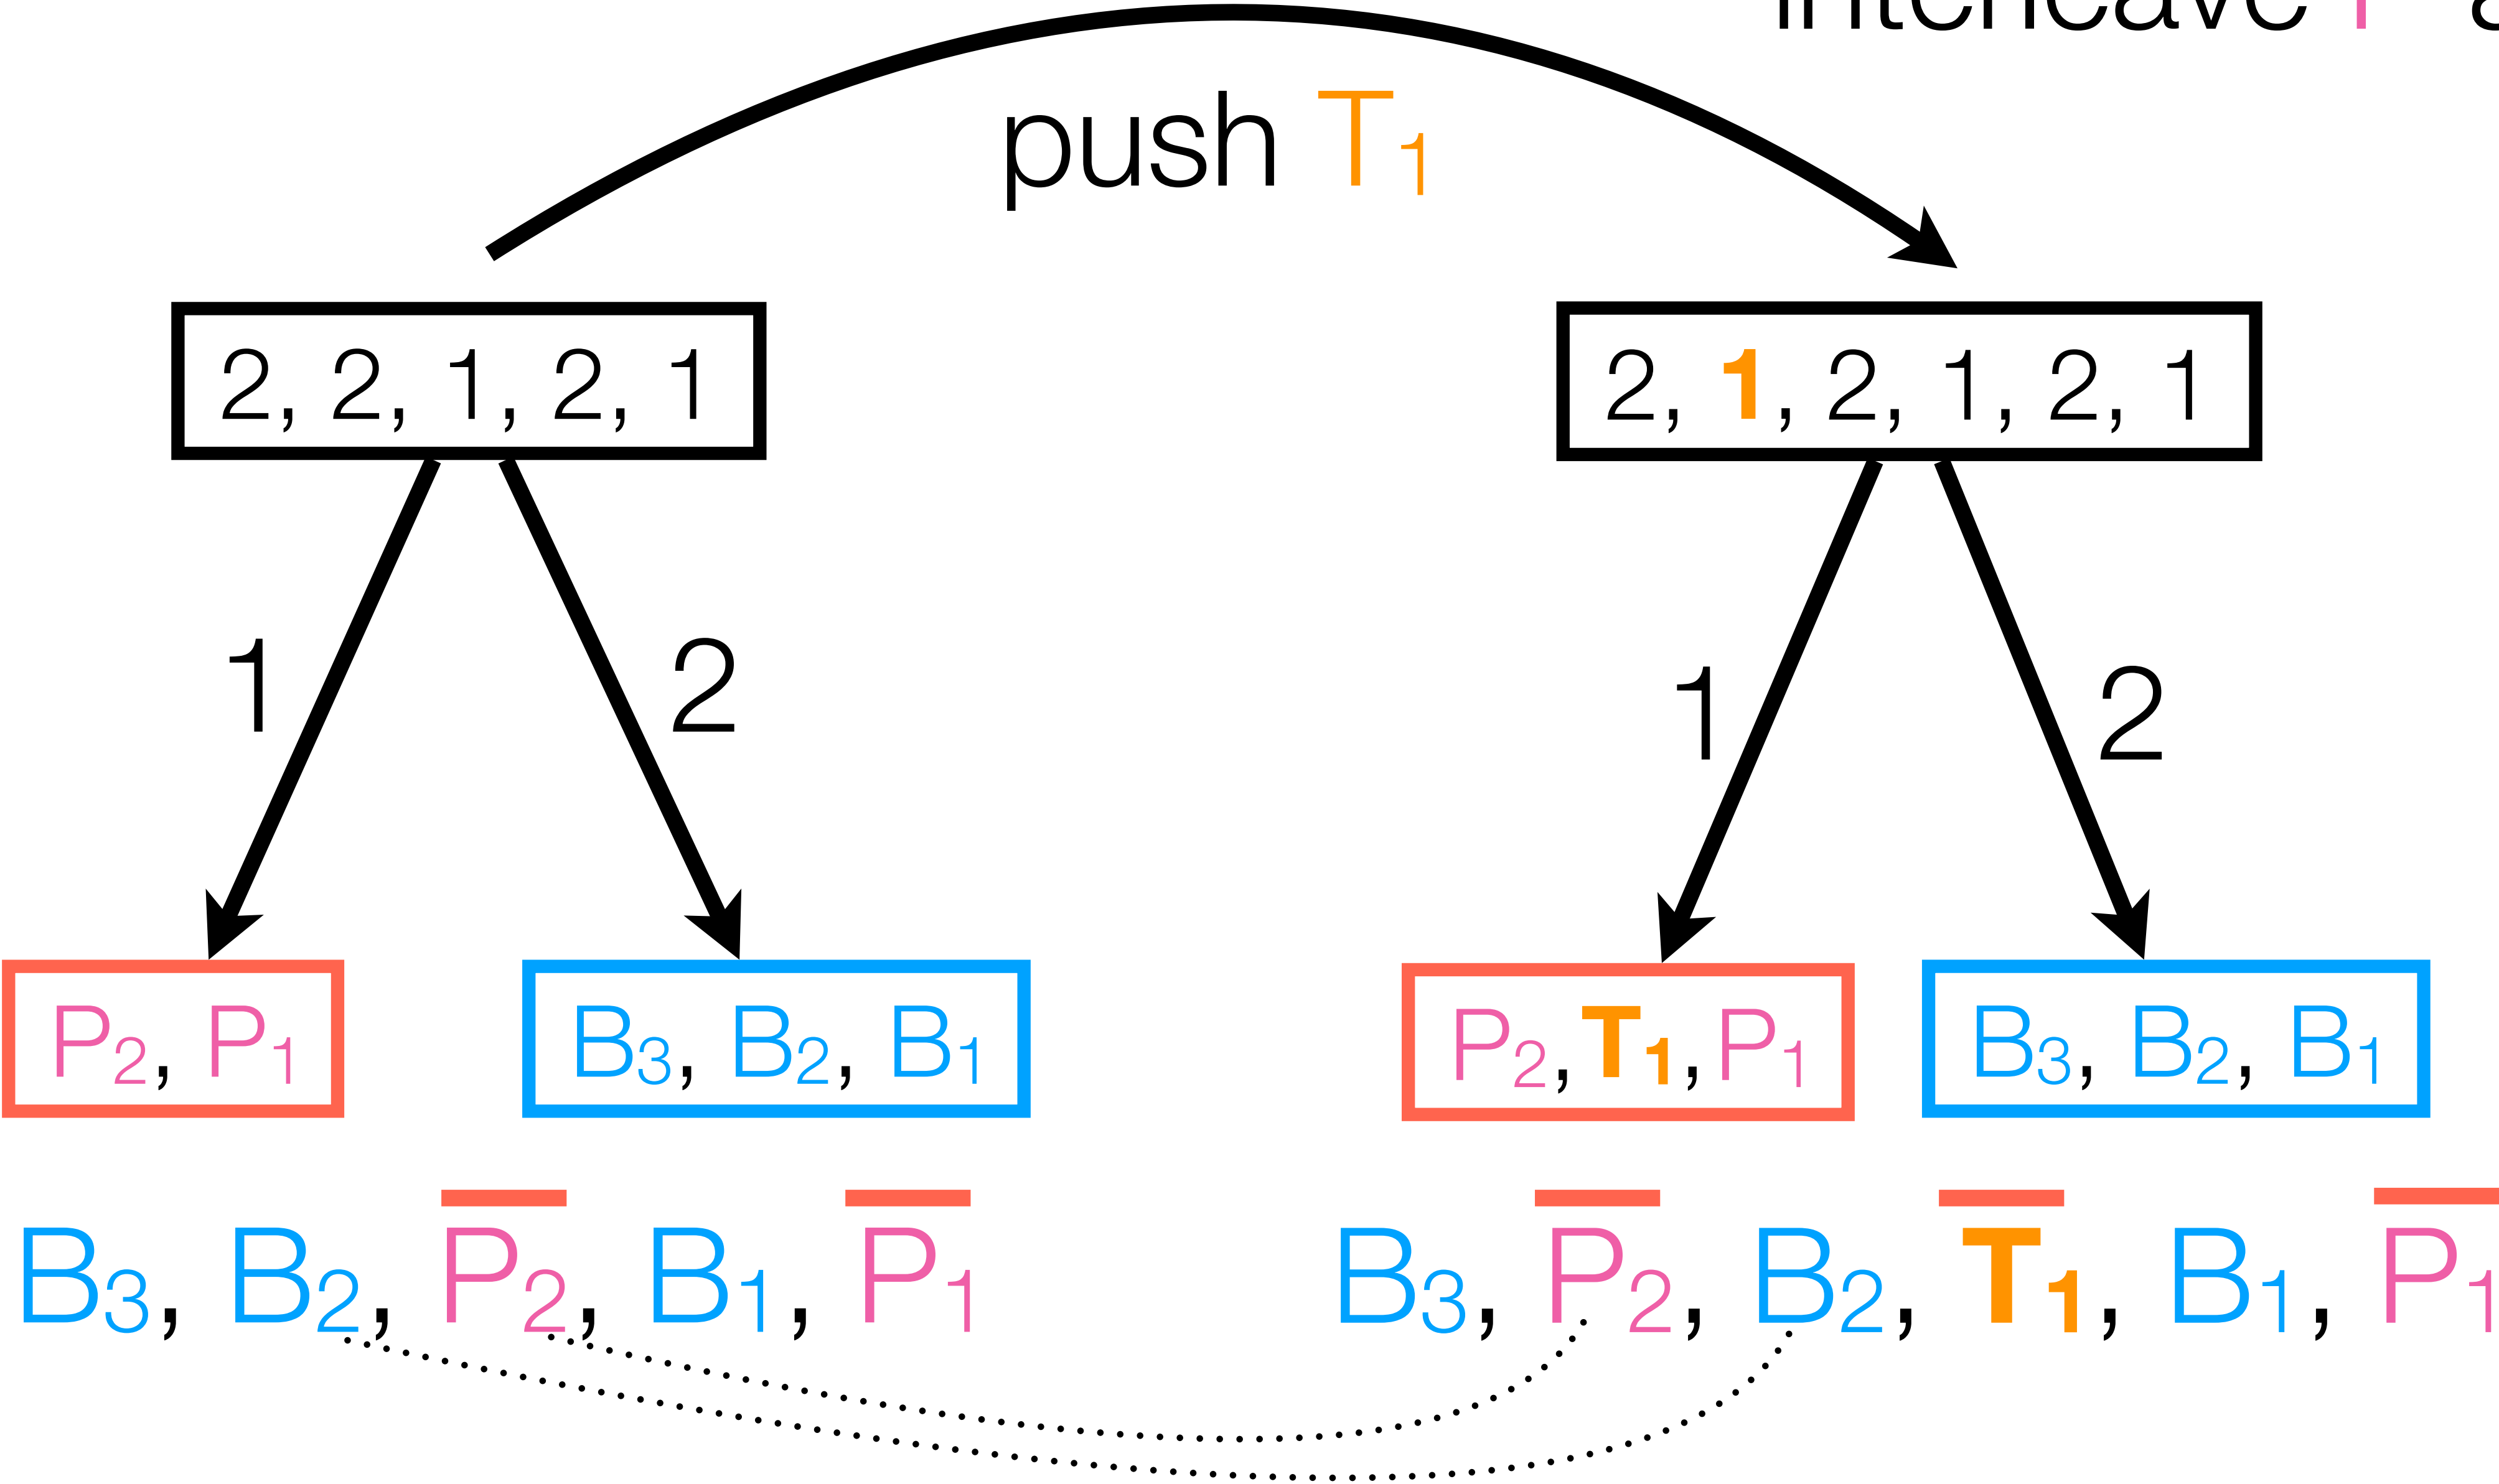
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



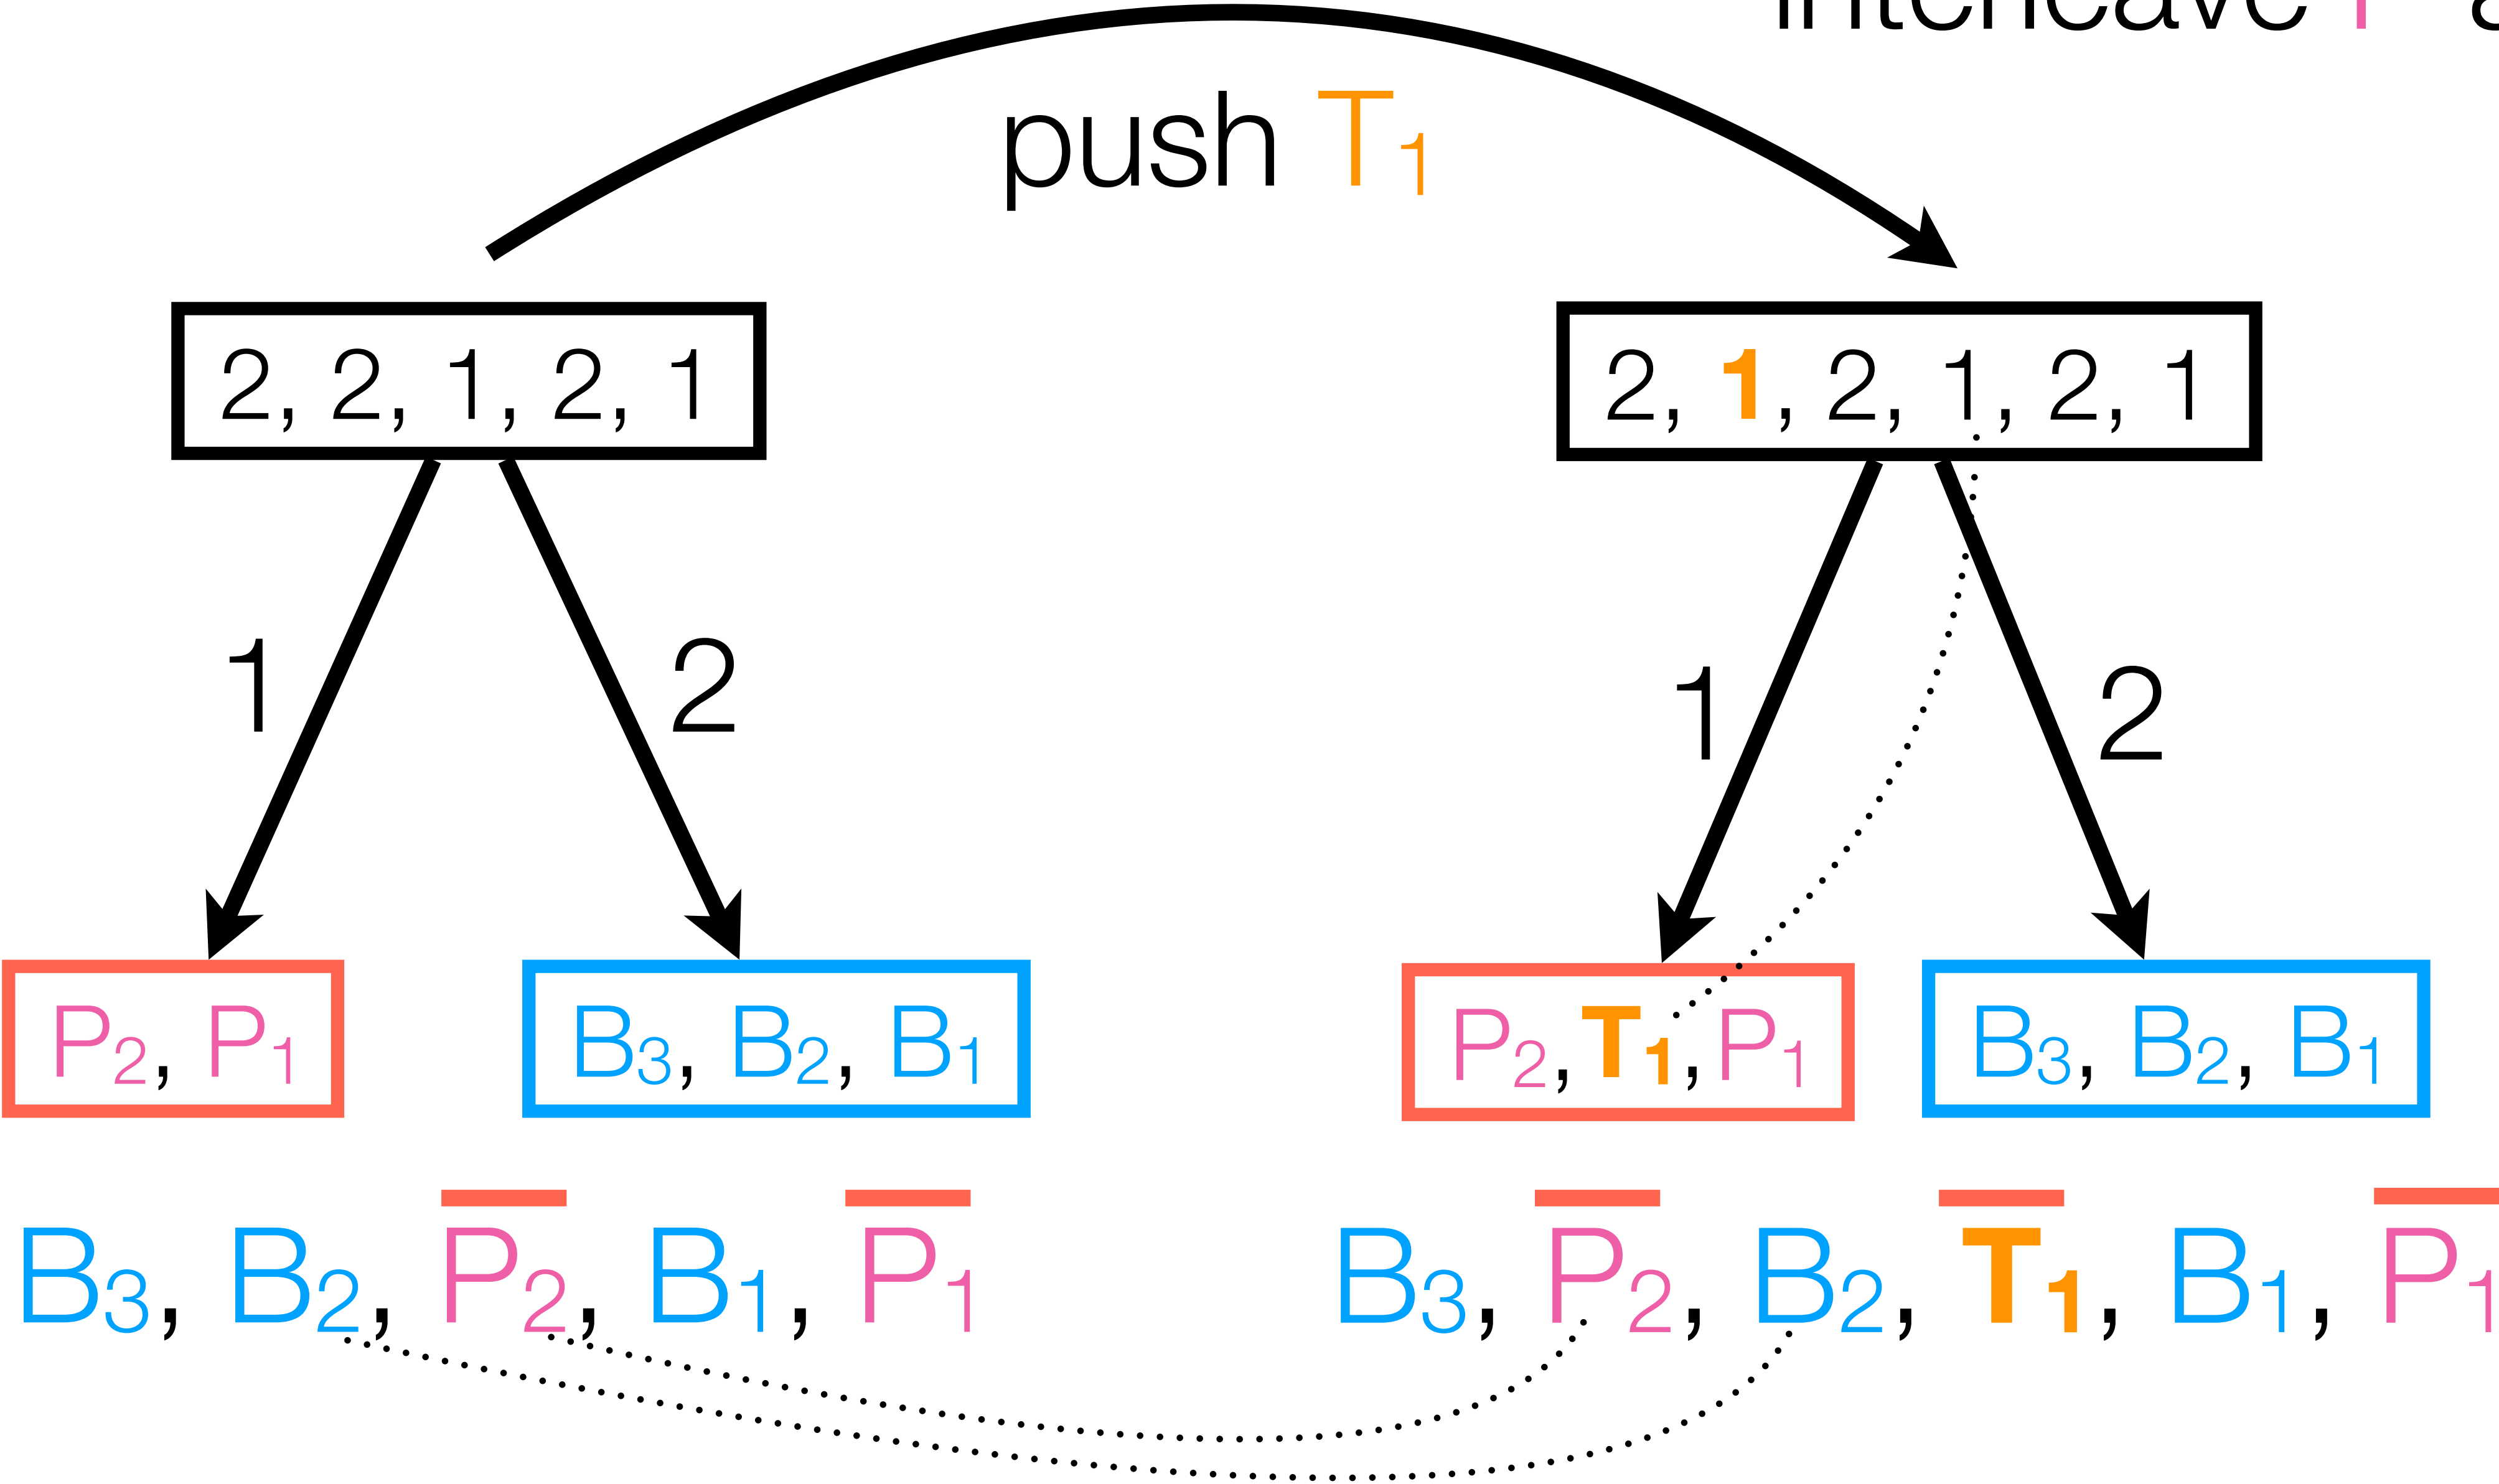
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



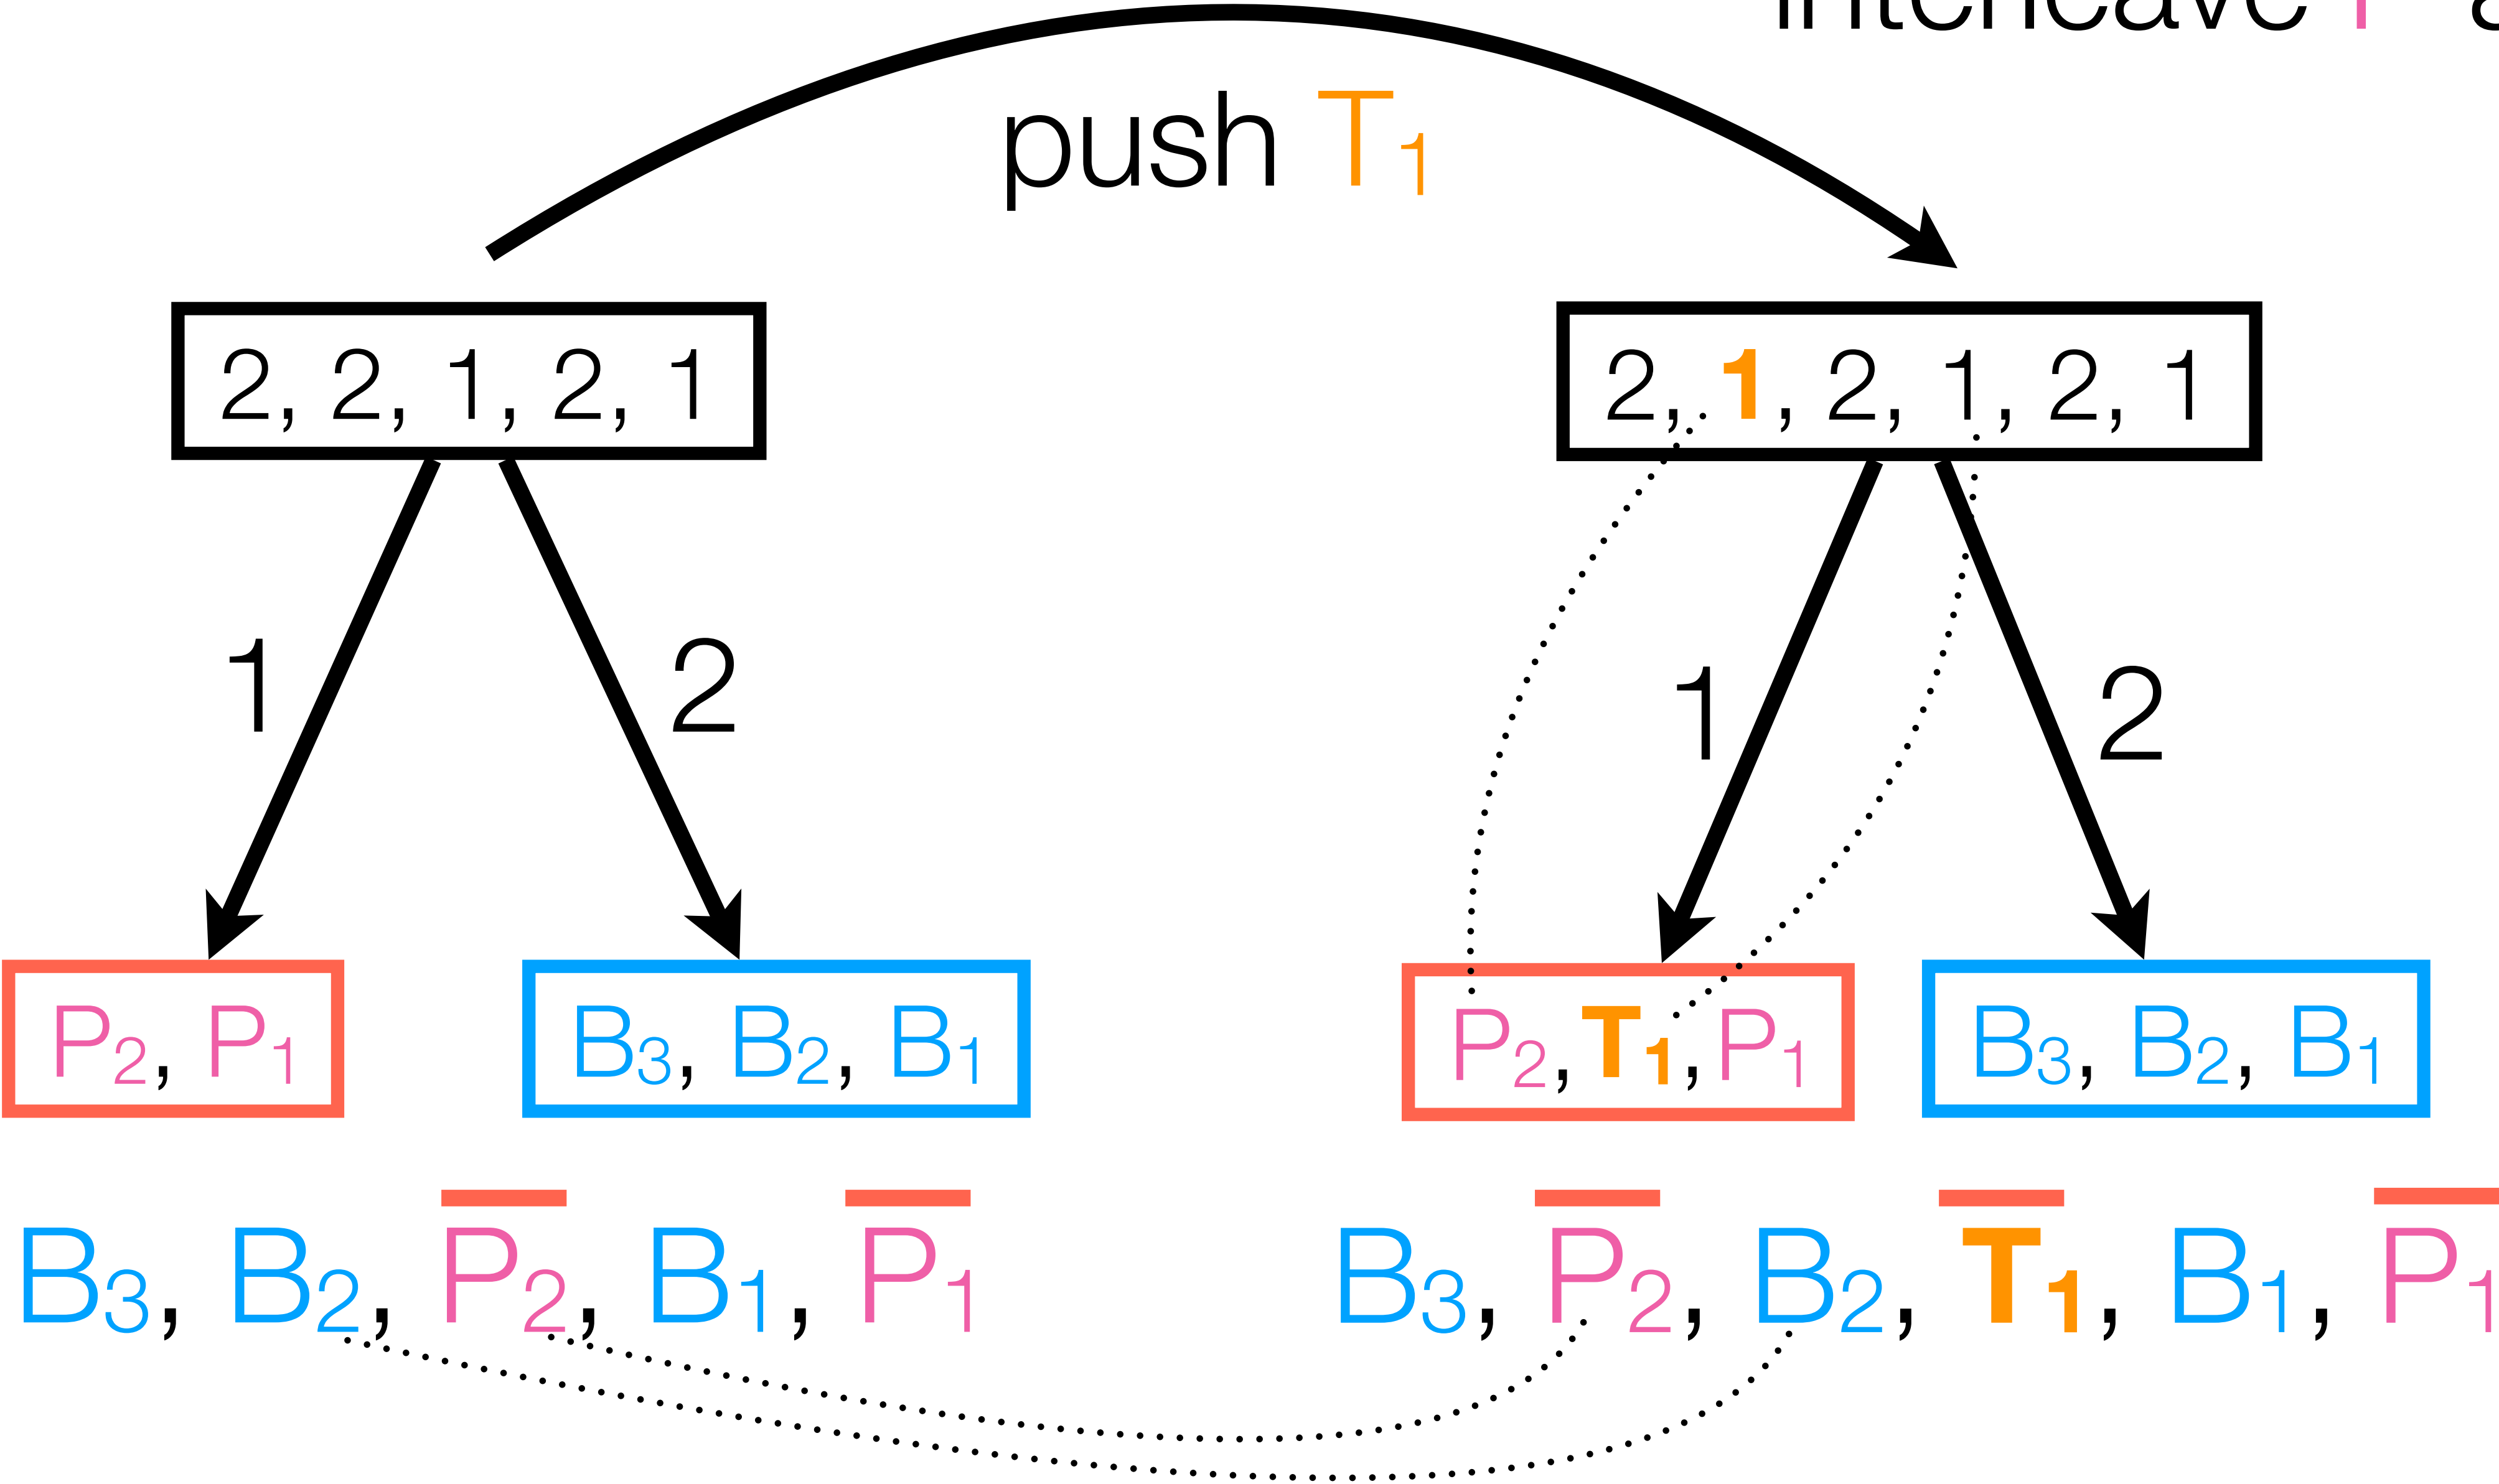
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



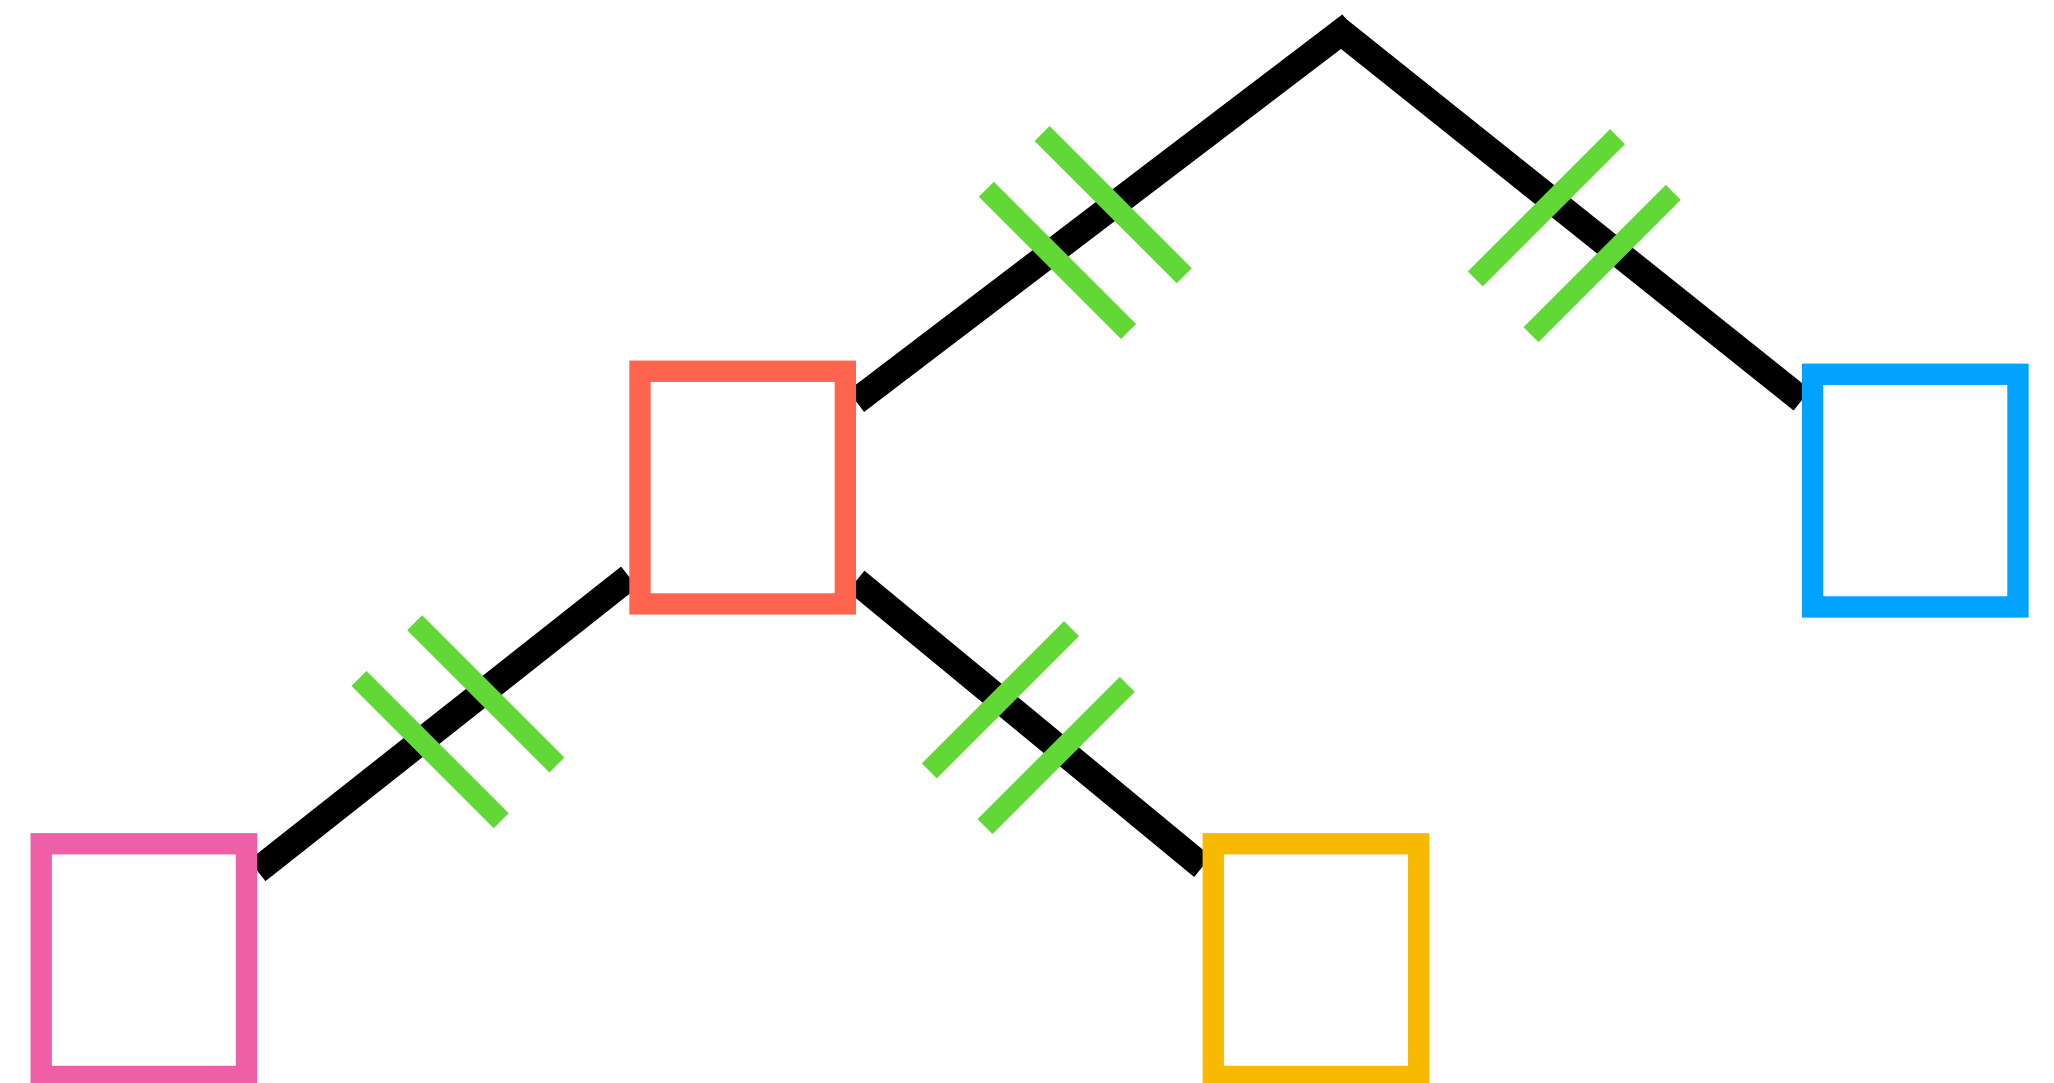
Introducing: PIFO trees

Interleave **R** and **B**;
interleave **P** and **T**.



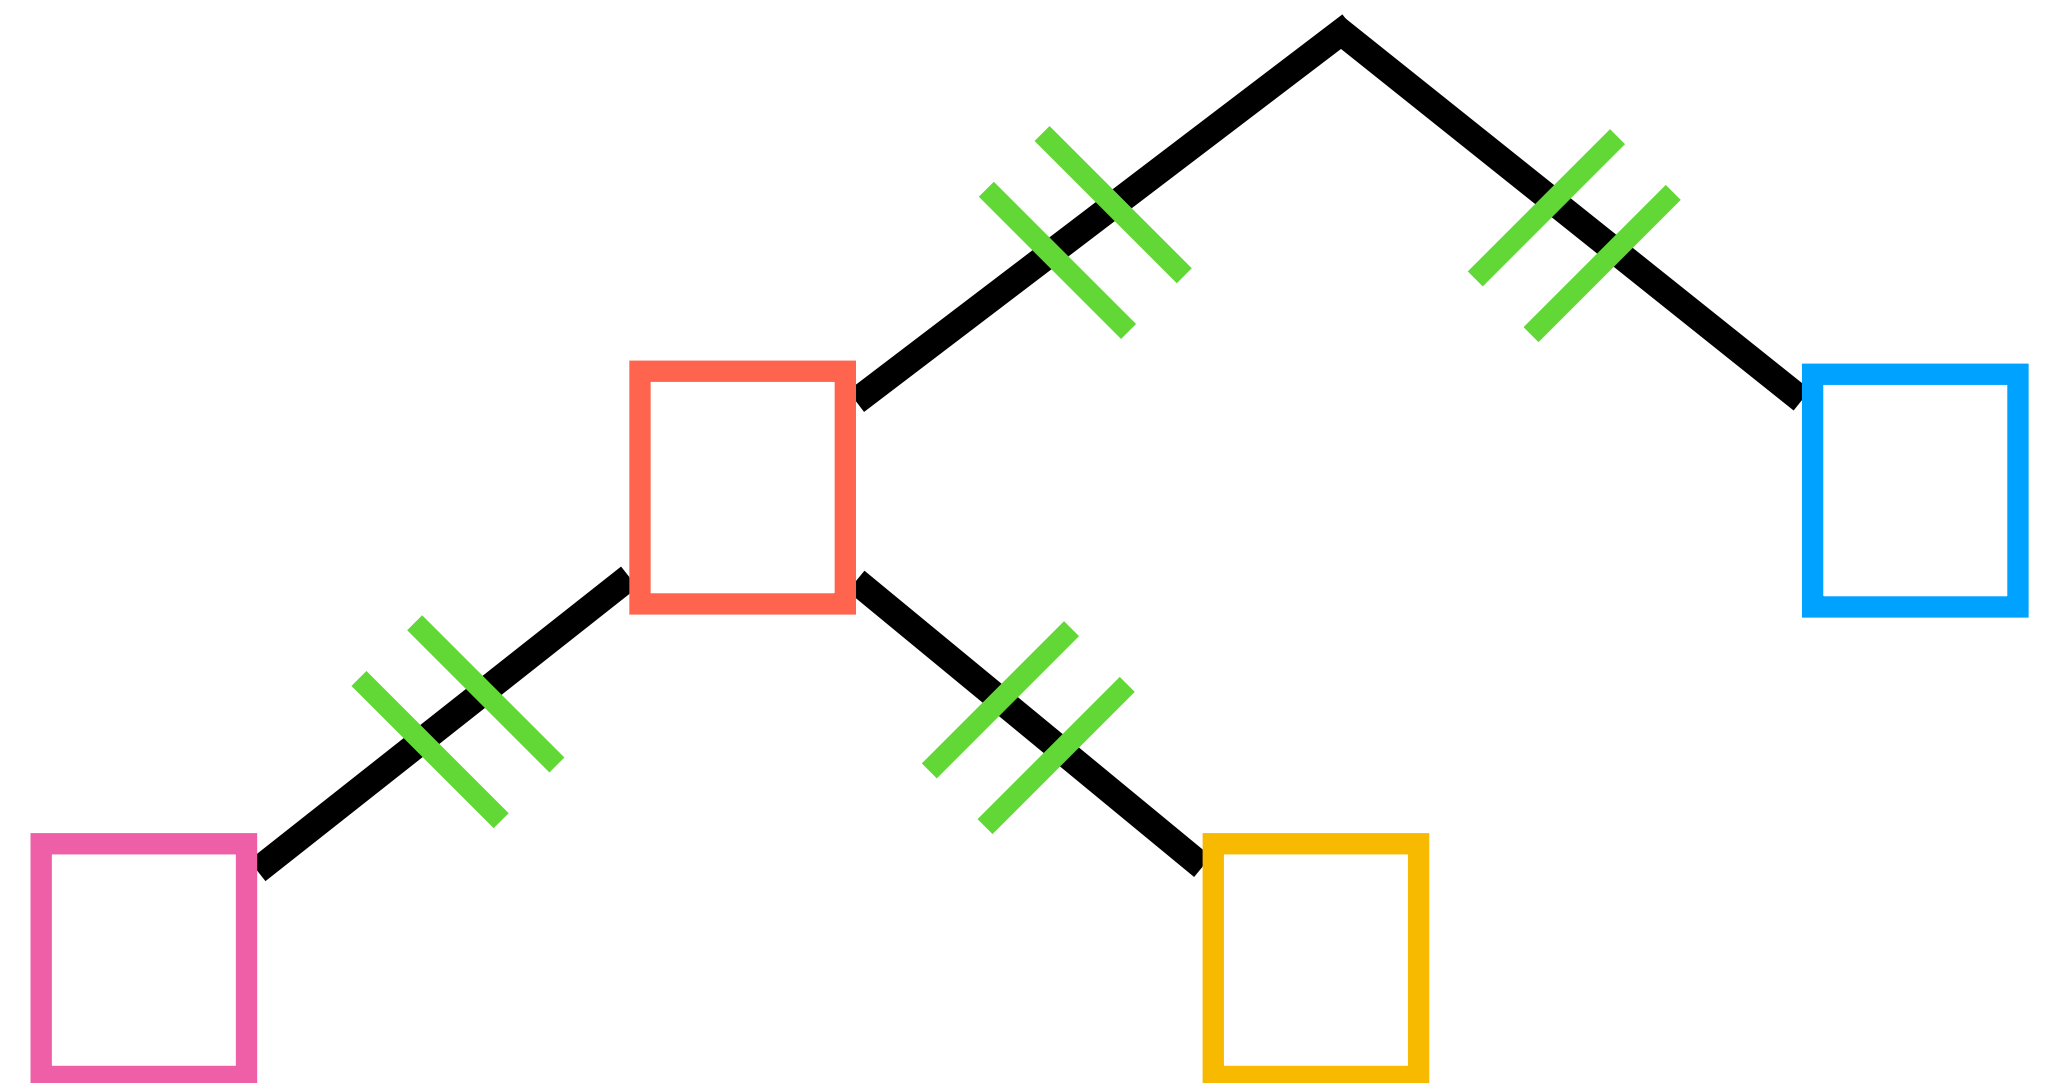


Goal:
interleave **R** and **B**;
interleave **P** and **T**.





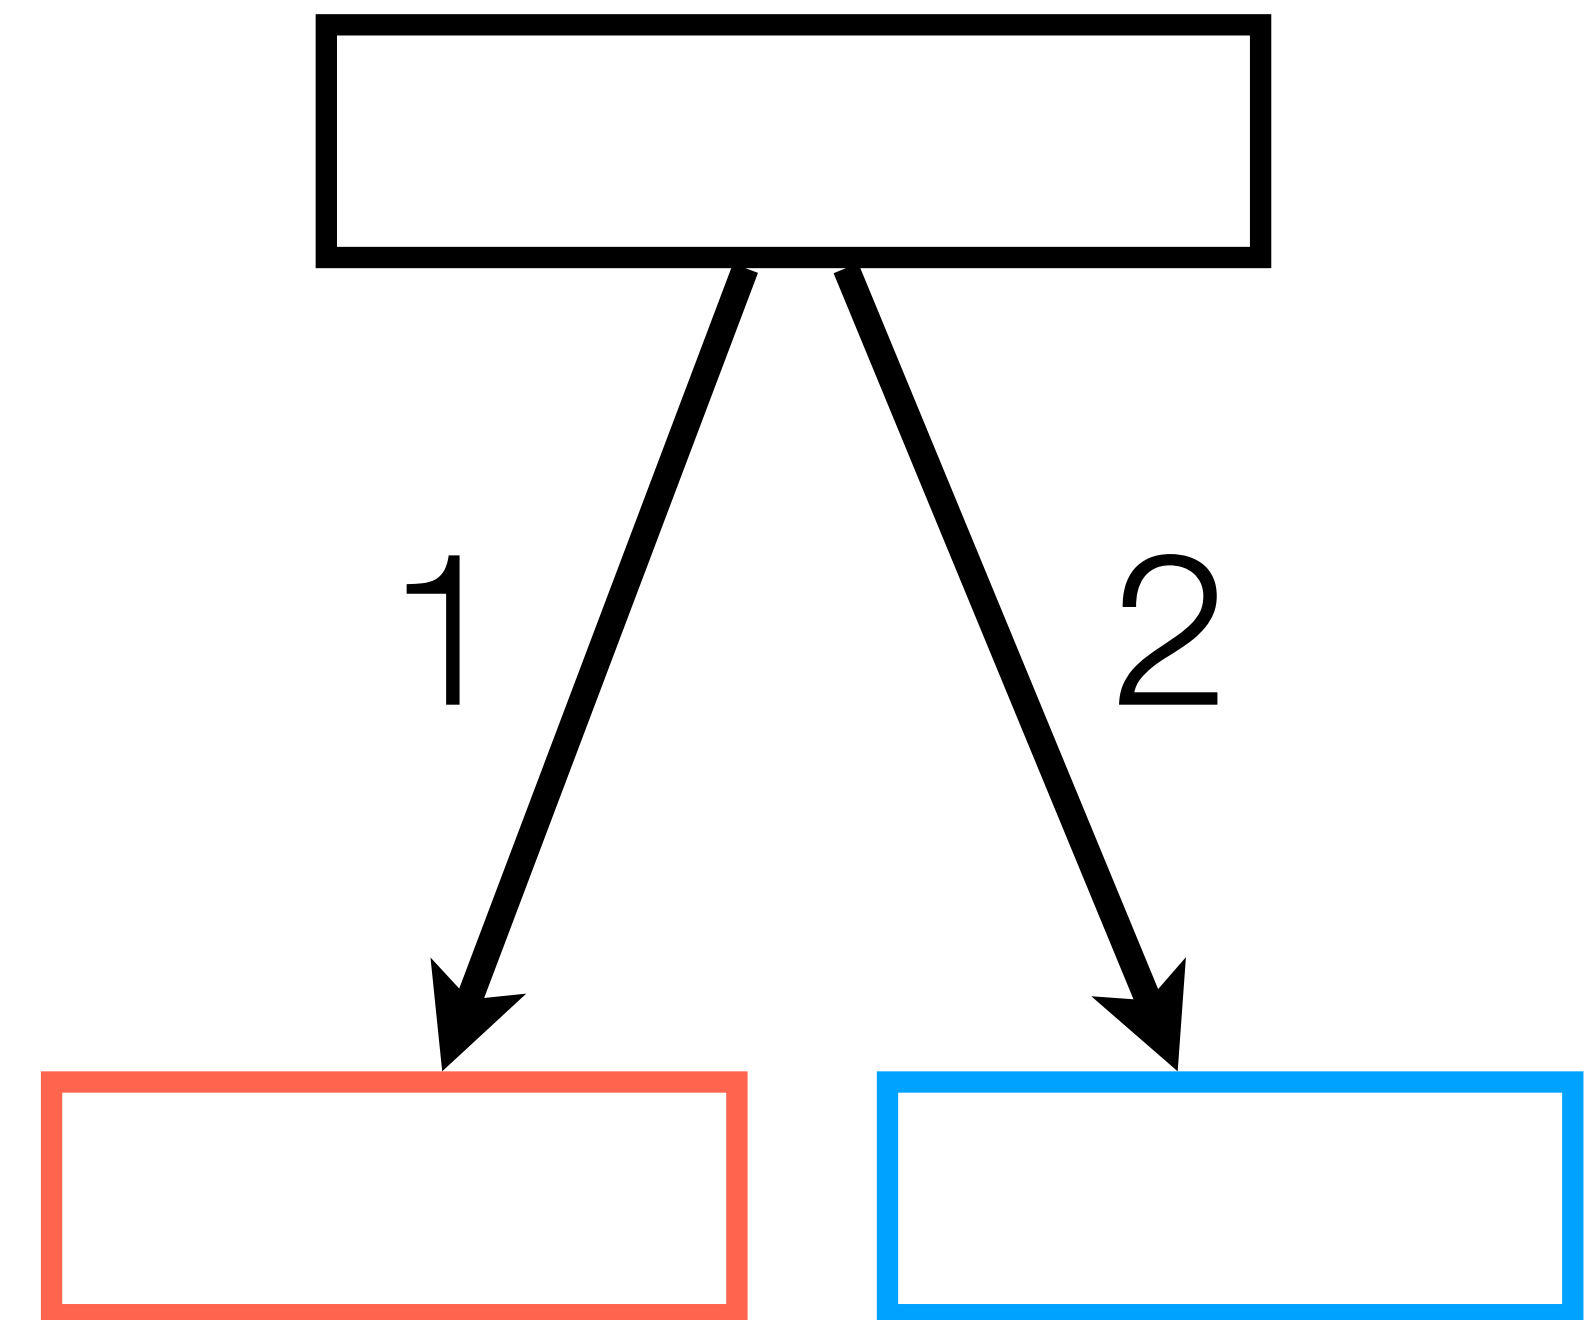
Goal:
interleave **R** and **B**;
interleave **P** and **T**.



Aside: PIFO Trees

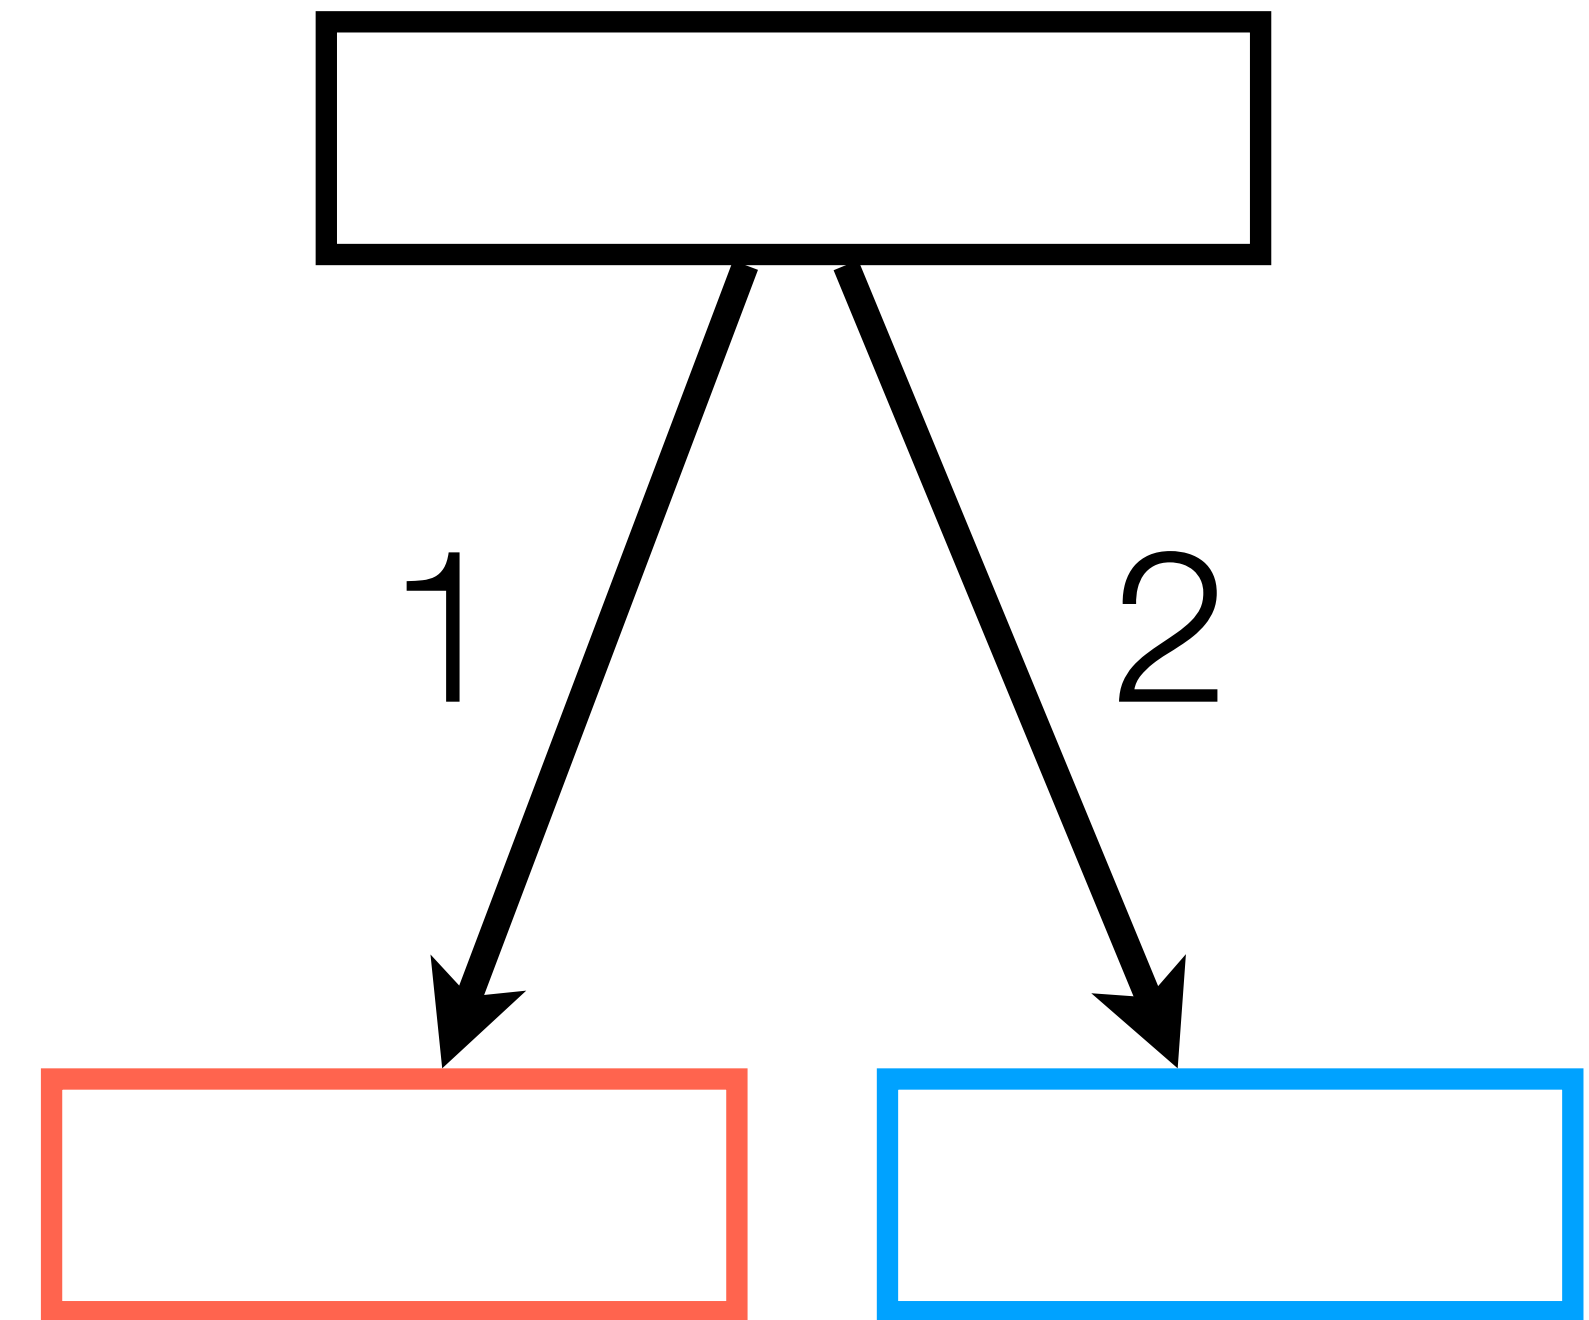
Sivaraman et al. at SIGCOMM '16

Key Insight



Key Insight

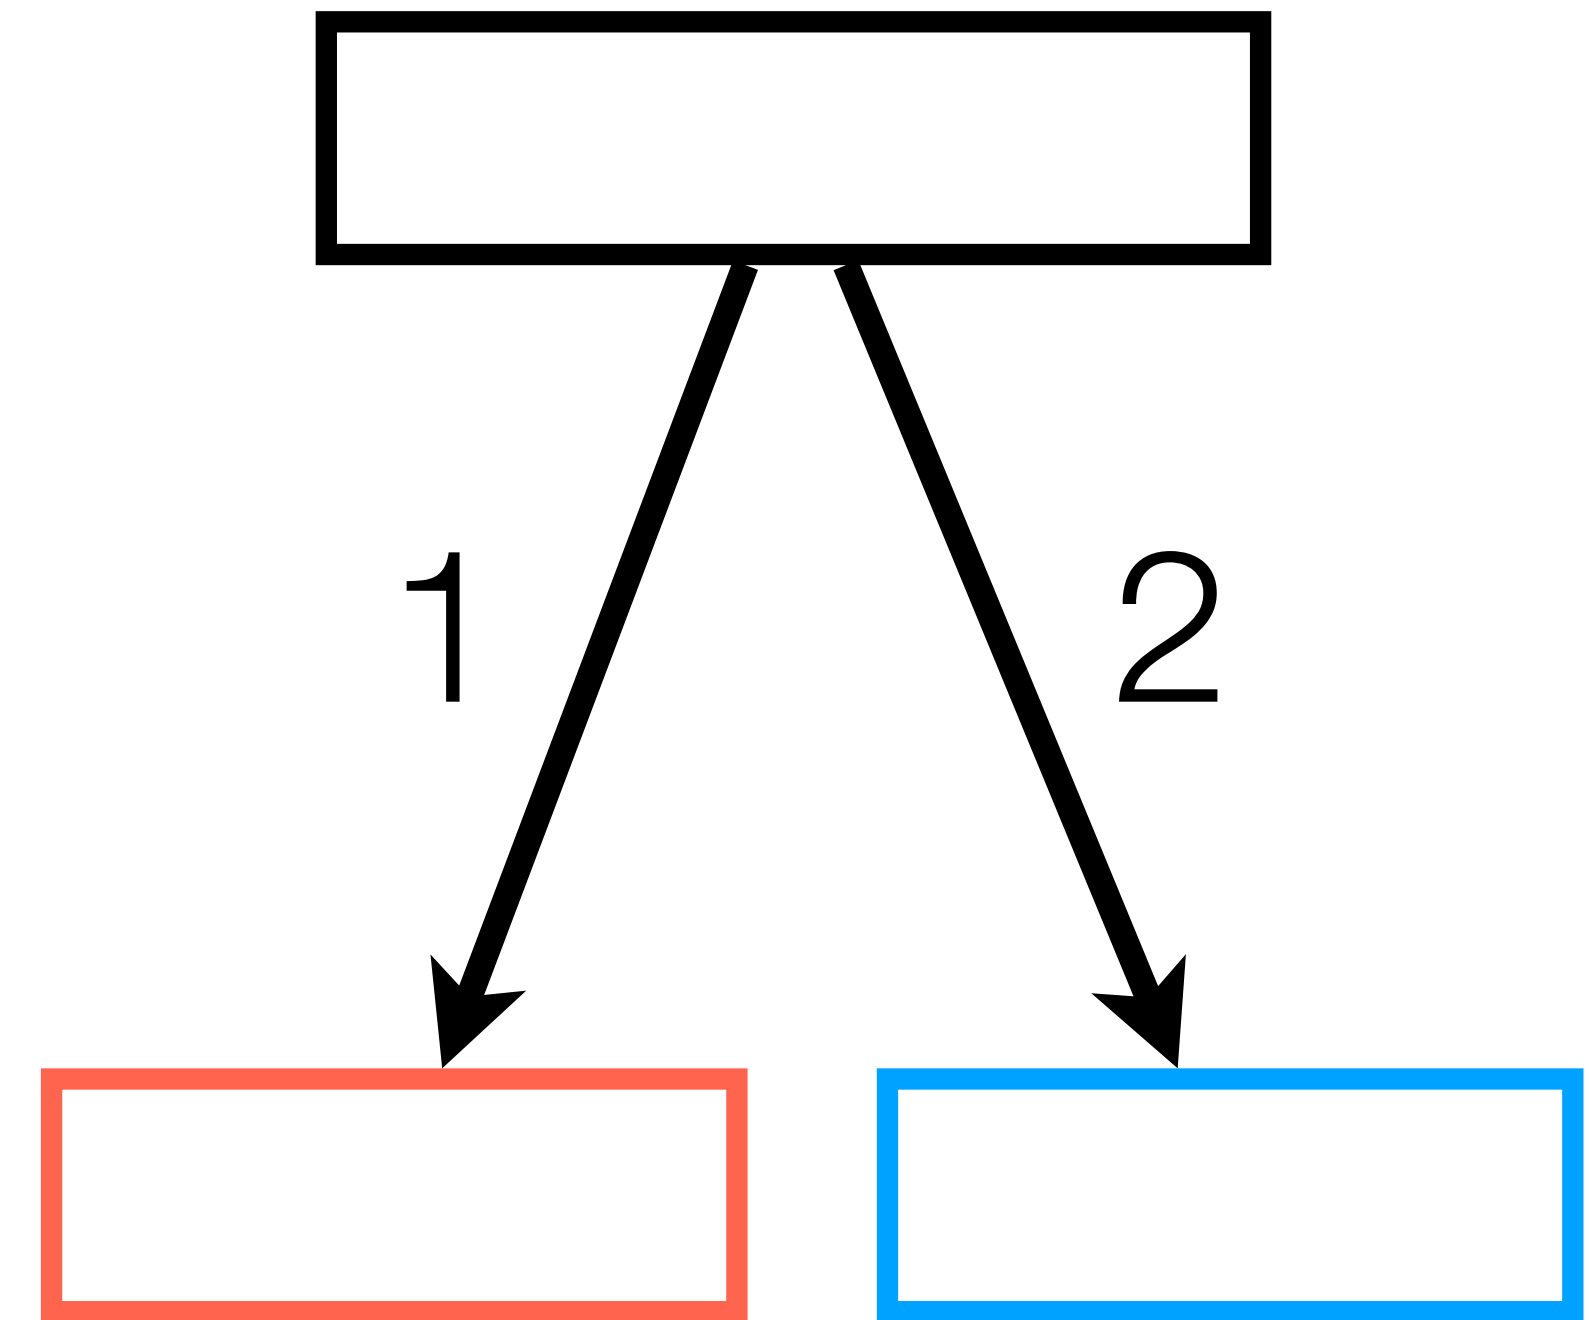
A PIFO tree manifests a *programming language*.



Key Insight

A PIFO tree manifests a *programming language*.

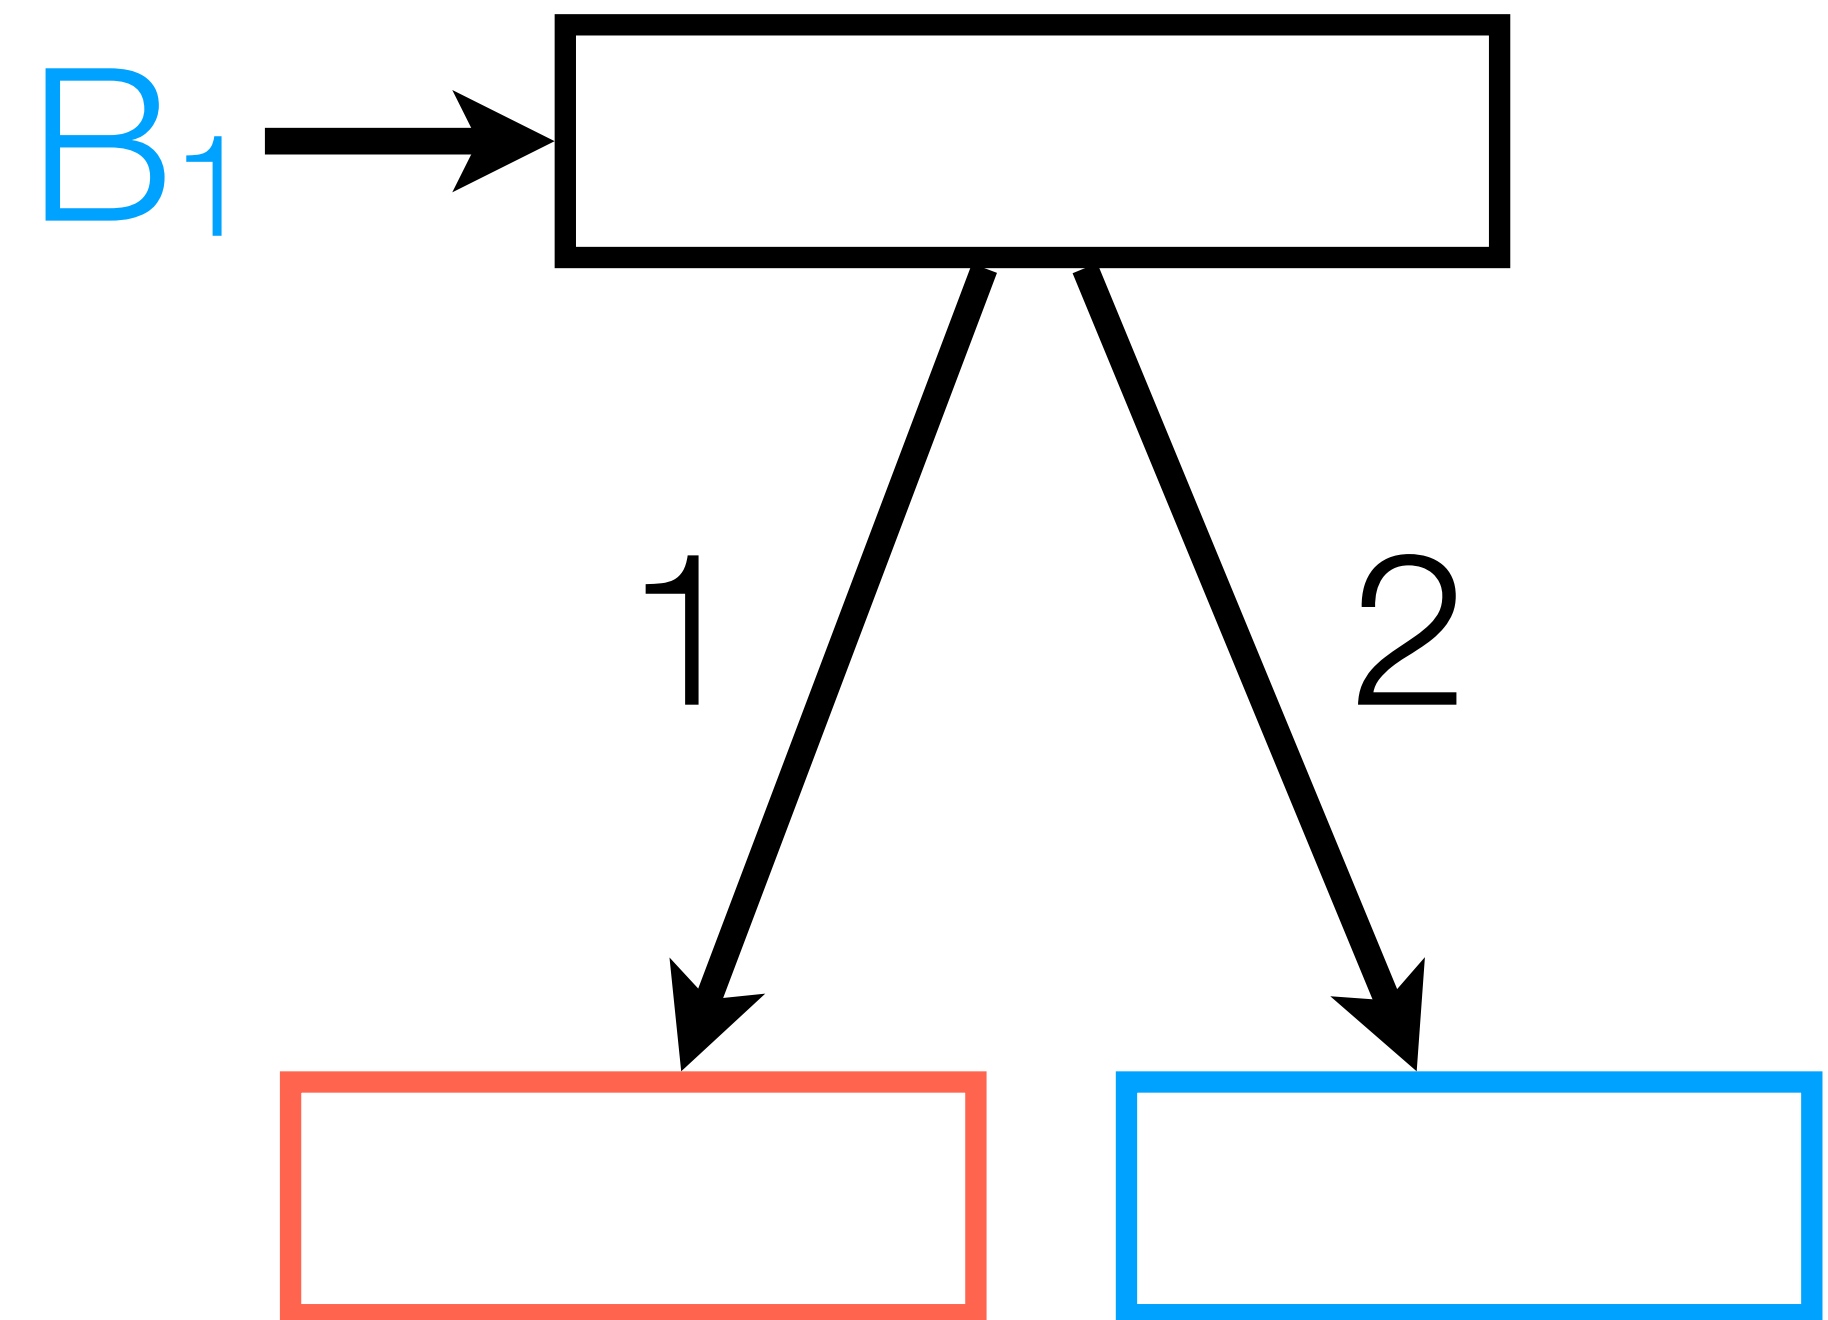
A program is precisely a *scheduling algorithm*.



Key Insight

A PIFO tree manifests a *programming language*.

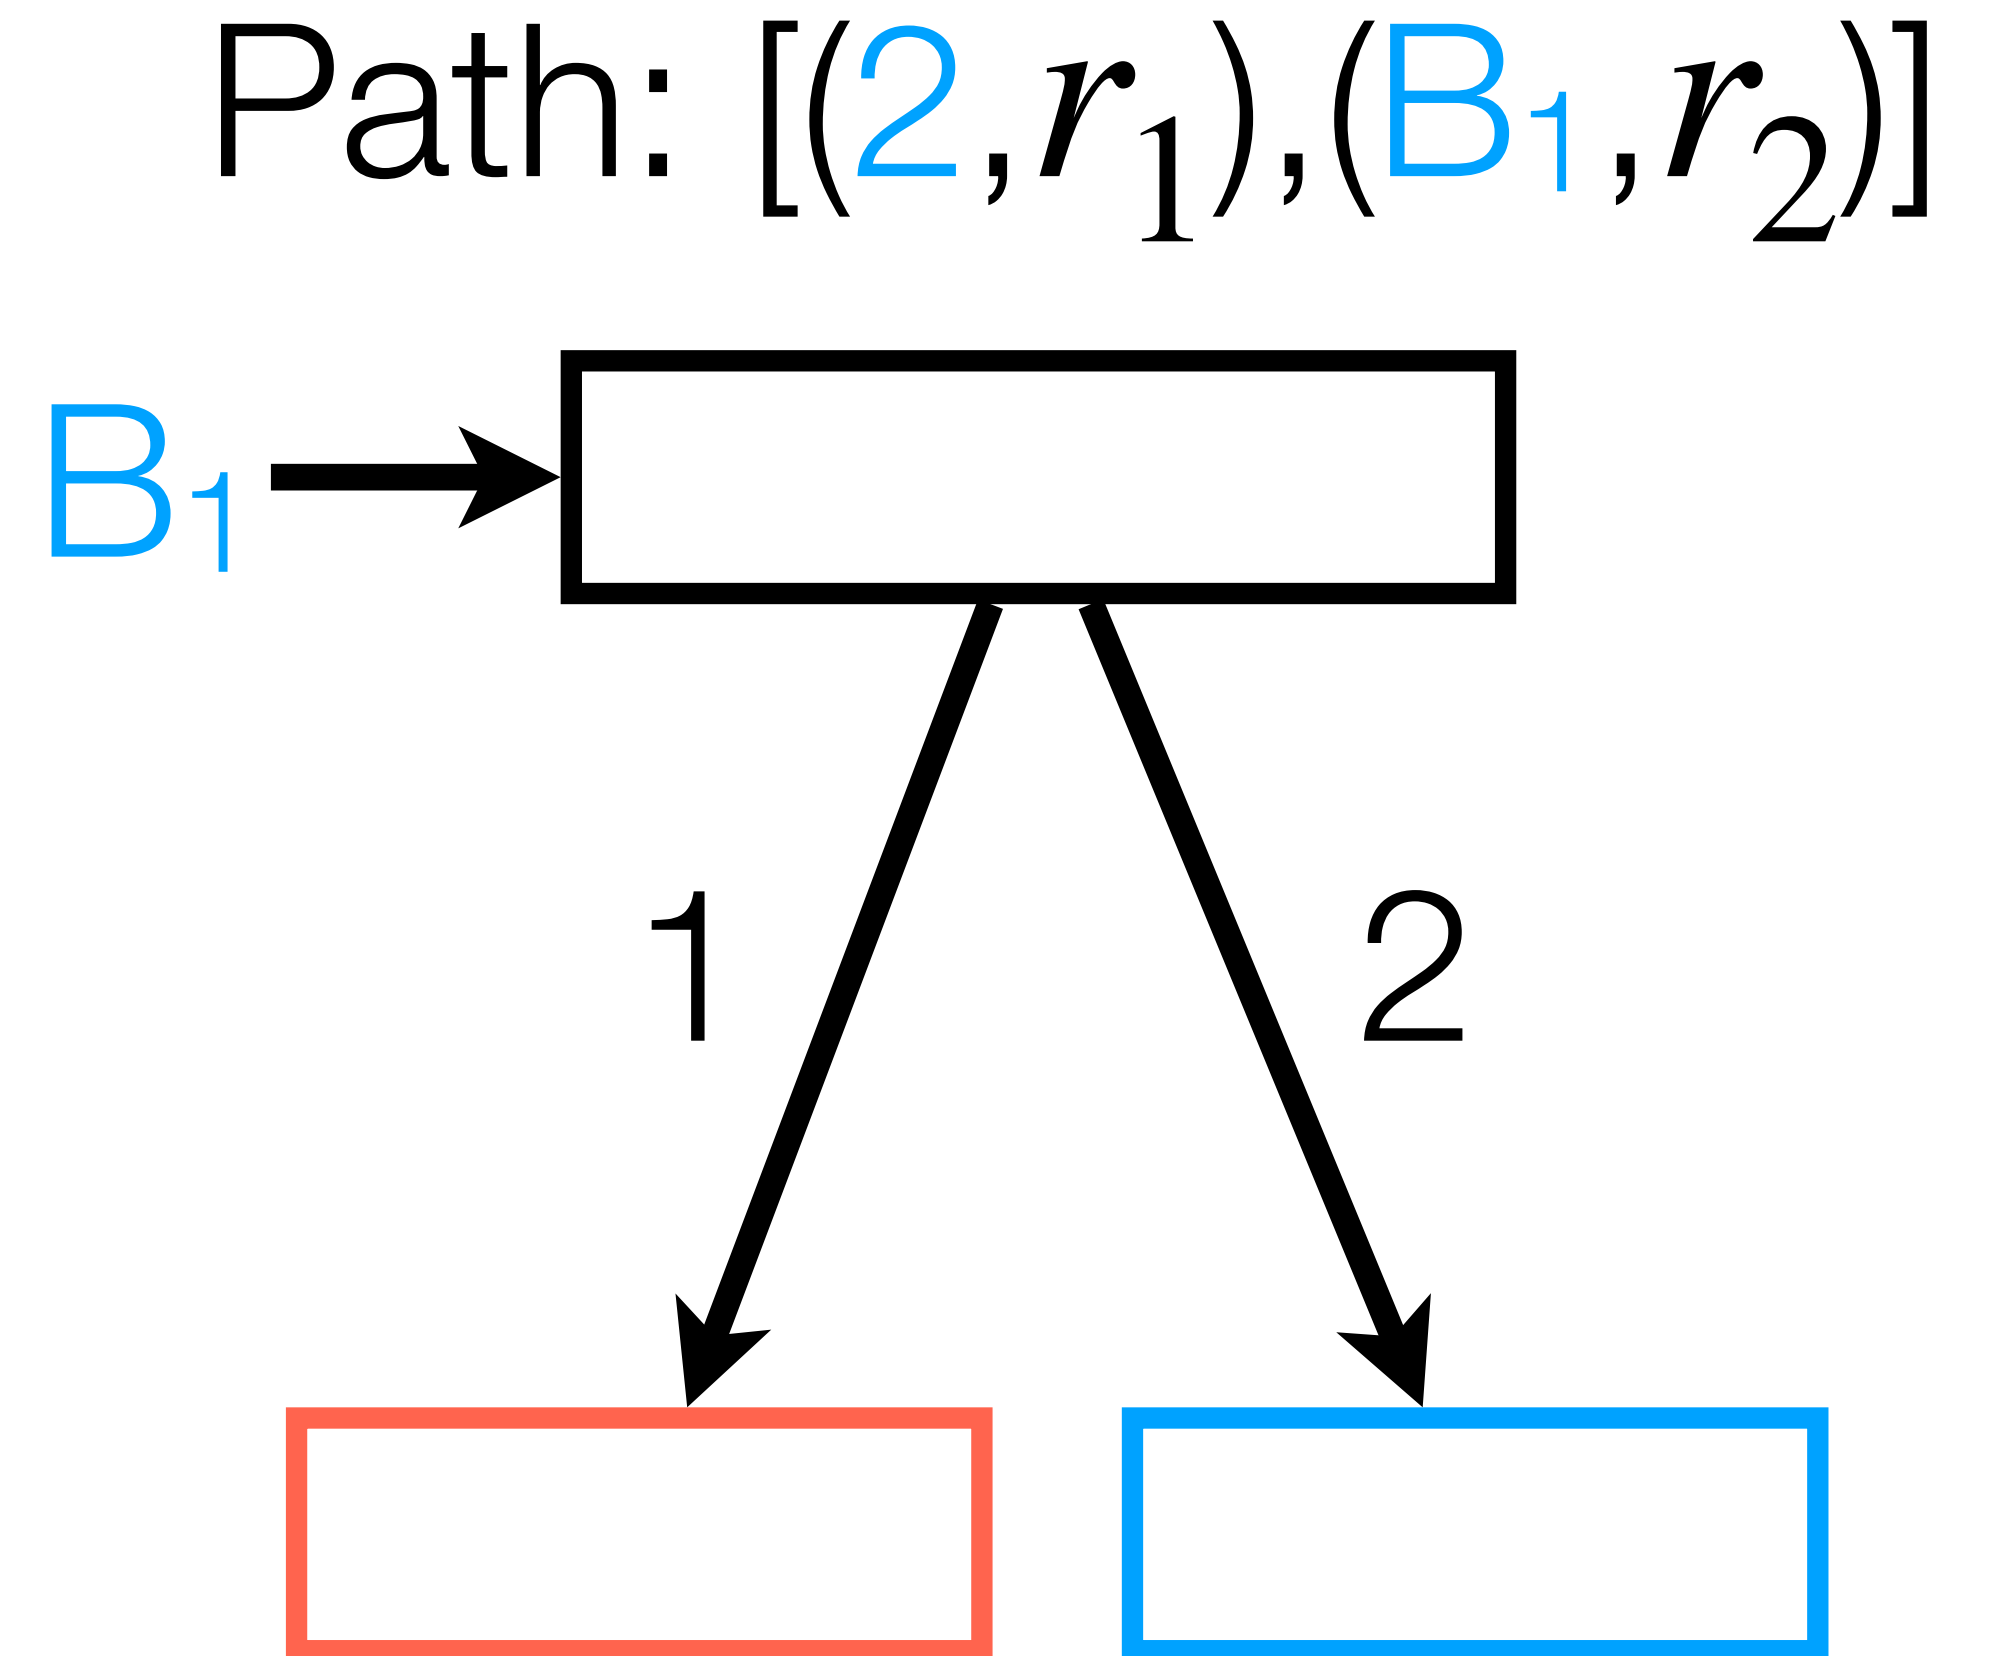
A program is precisely a *scheduling algorithm*.



Key Insight

A PIFO tree manifests a *programming language*.

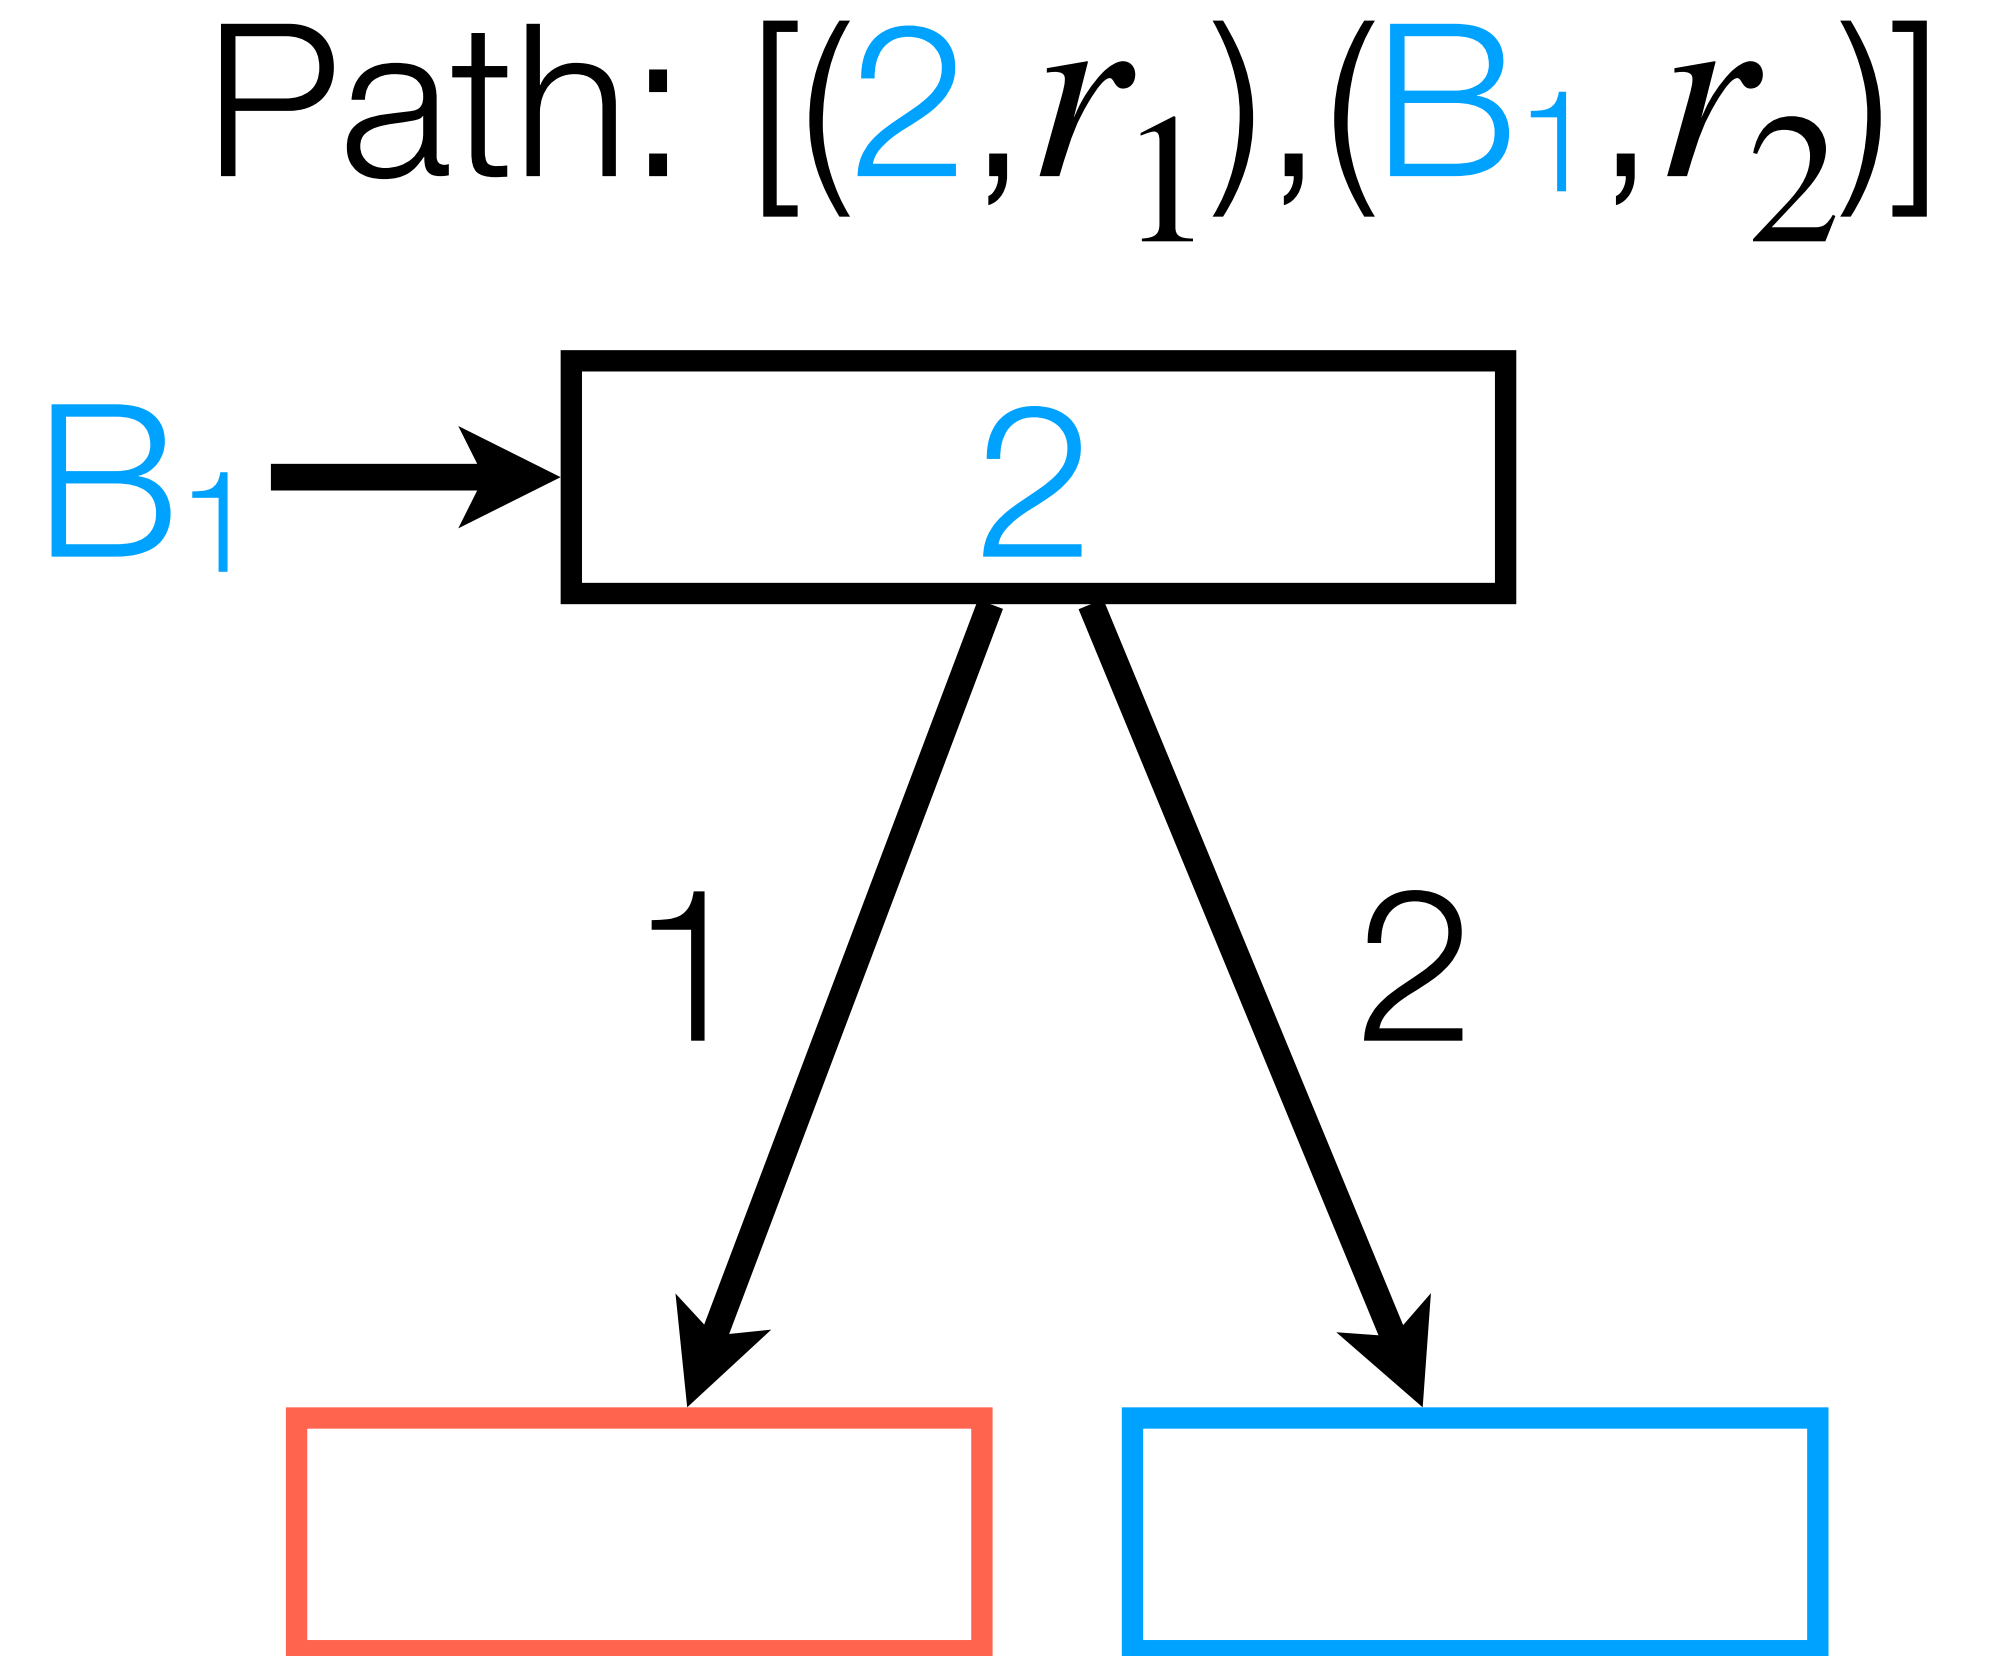
A program is precisely a *scheduling algorithm*.



Key Insight

A PIFO tree manifests a *programming language*.

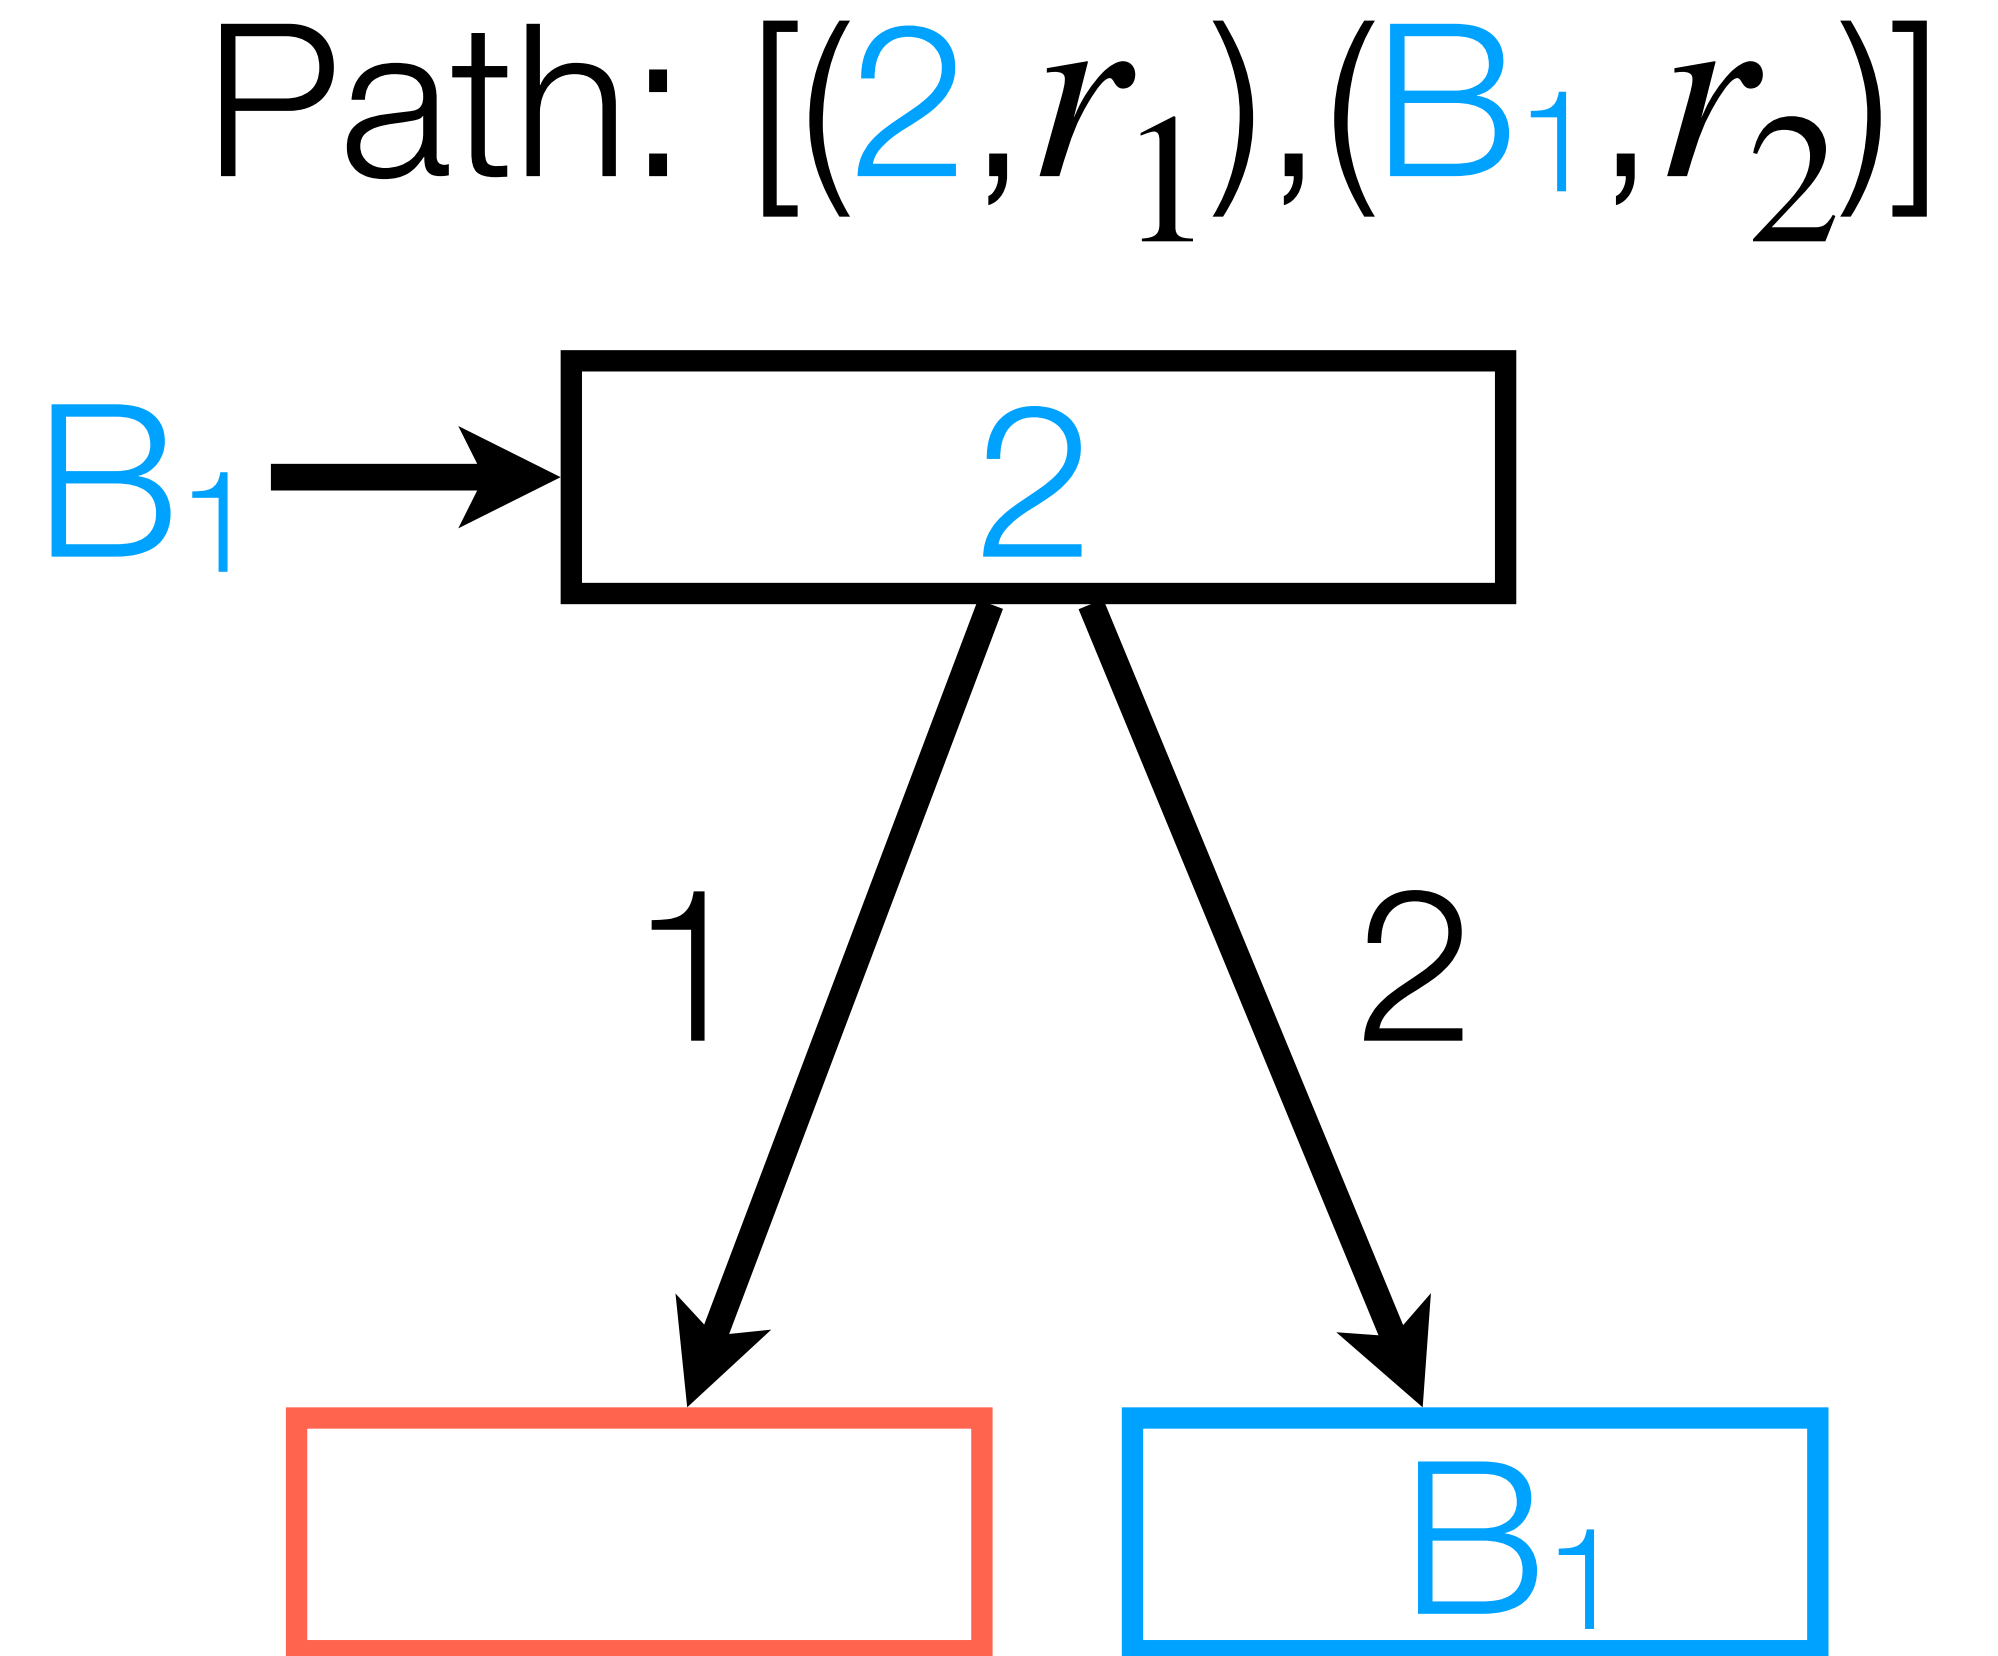
A program is precisely a *scheduling algorithm*.



Key Insight

A PIFO tree manifests a *programming language*.

A program is precisely a *scheduling algorithm*.



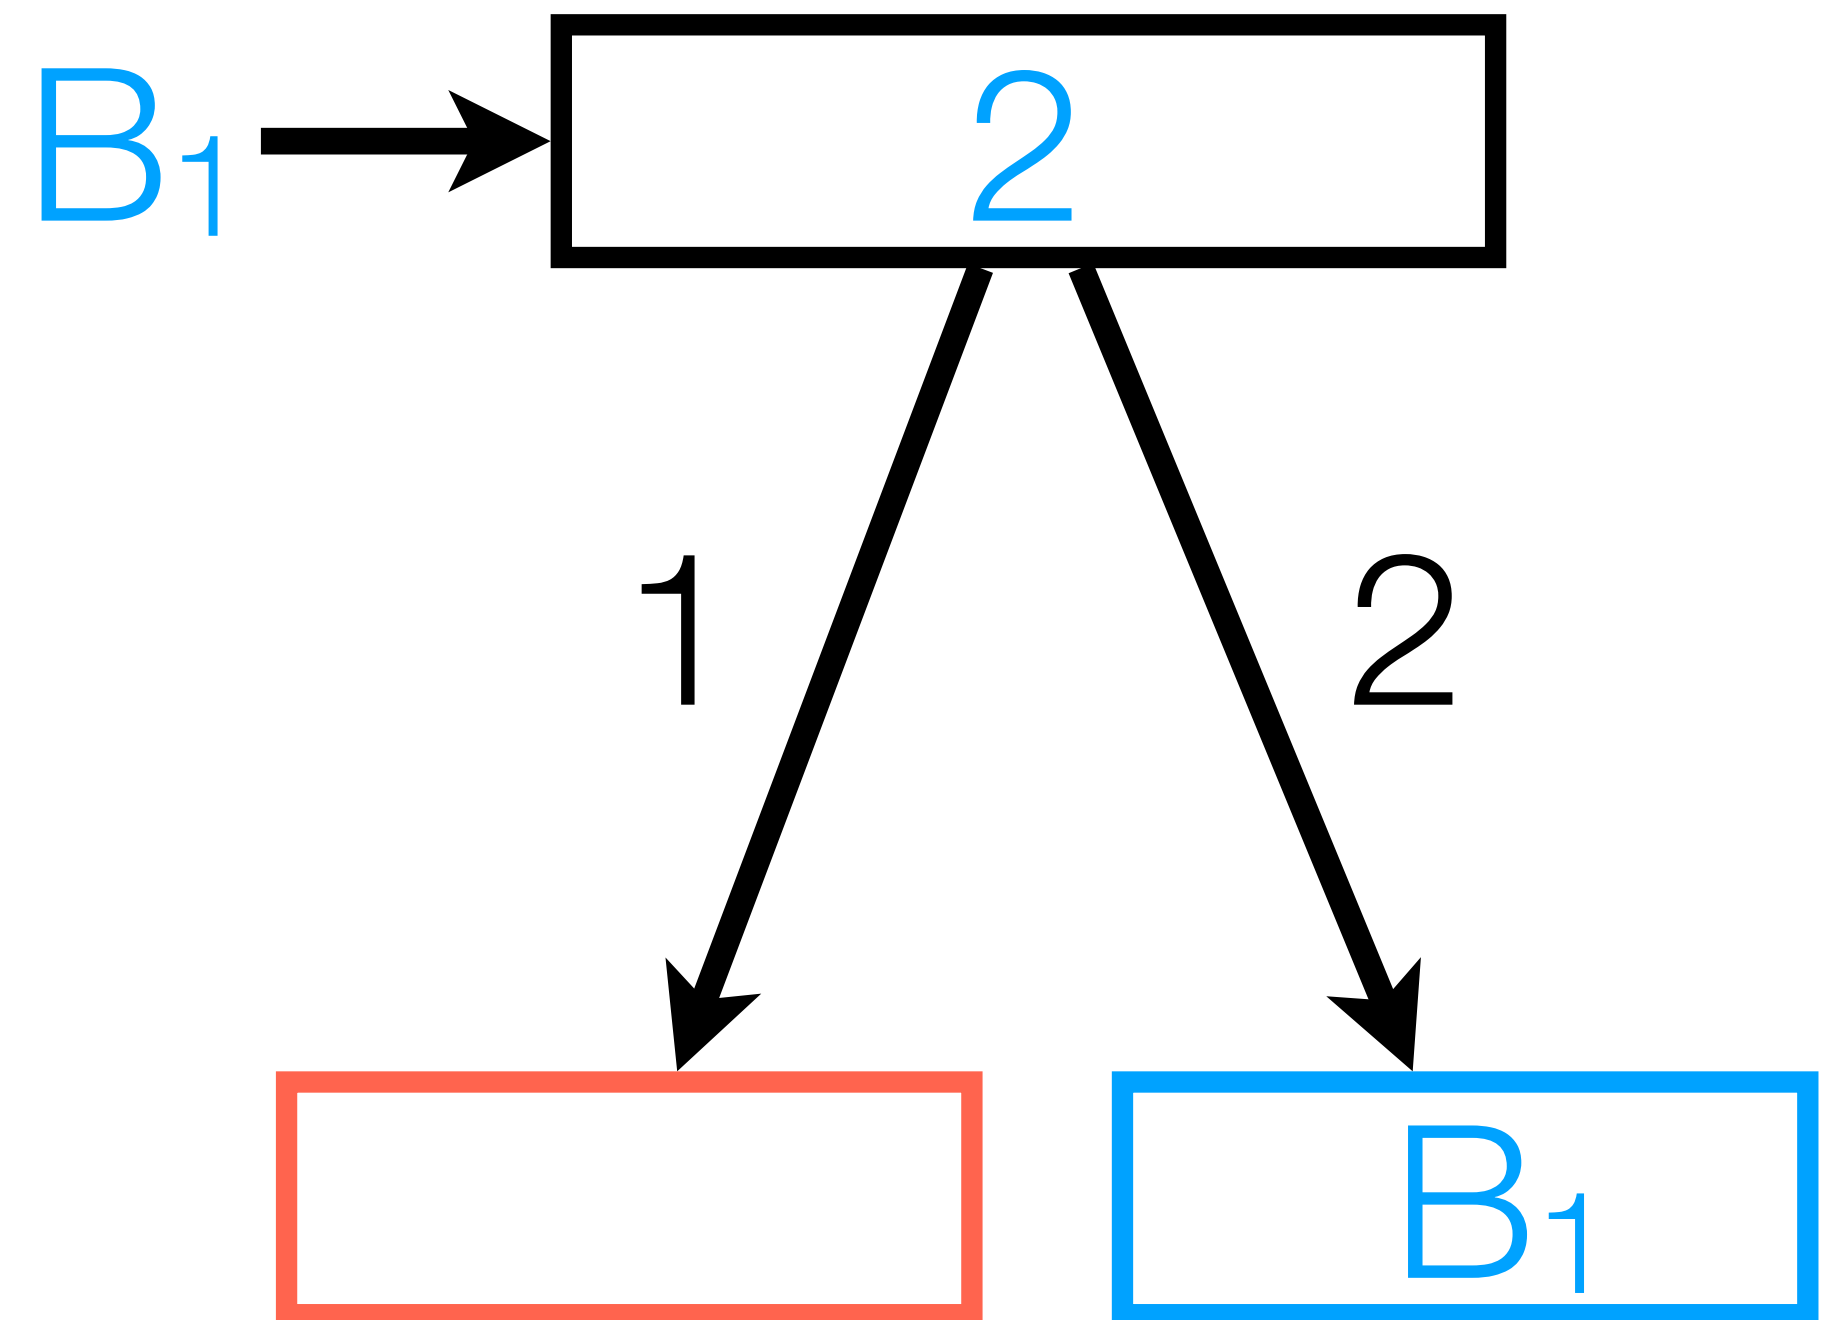
Key Insight

A PIFO tree manifests a *programming language*.

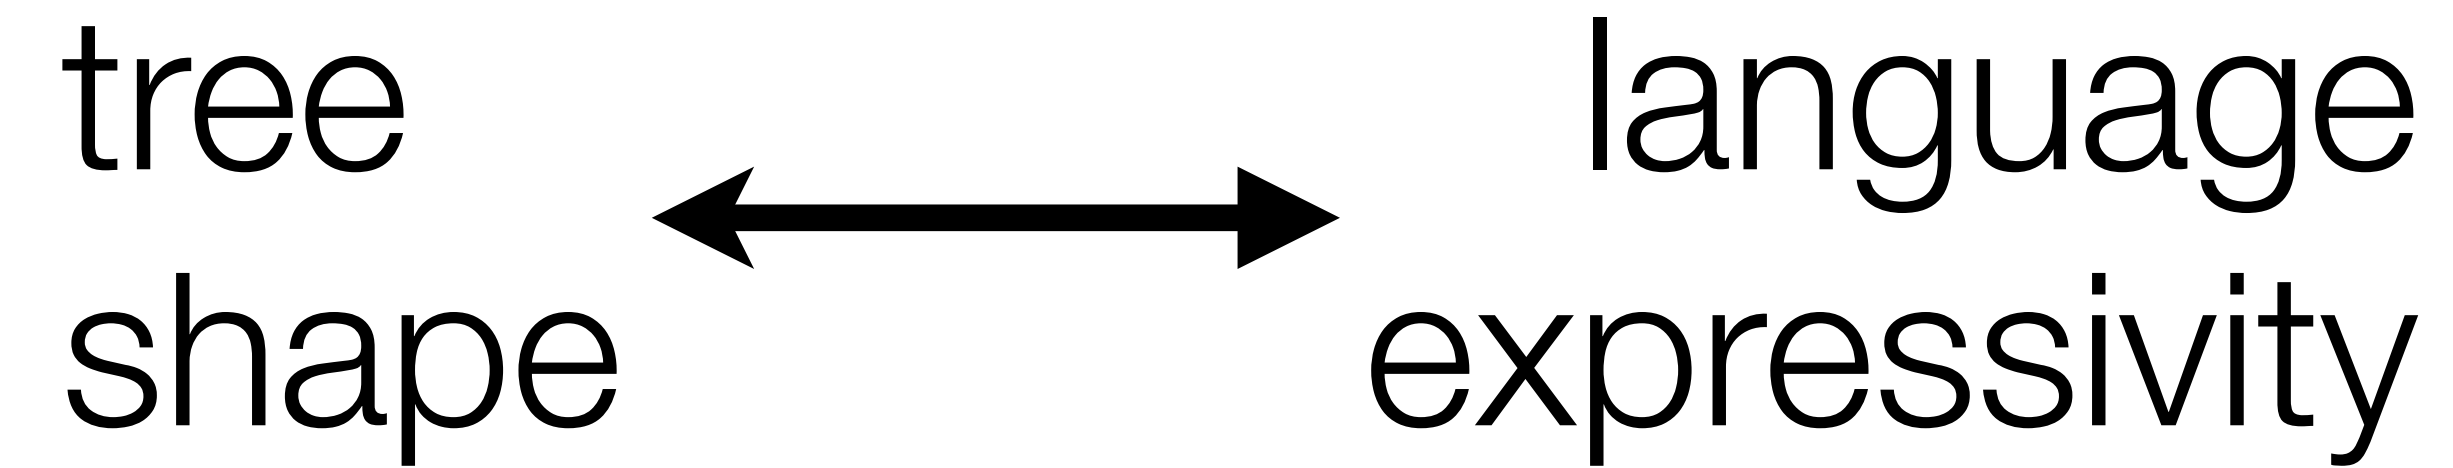
A program is precisely a *scheduling algorithm*.

tree shape \longleftrightarrow language expressivity

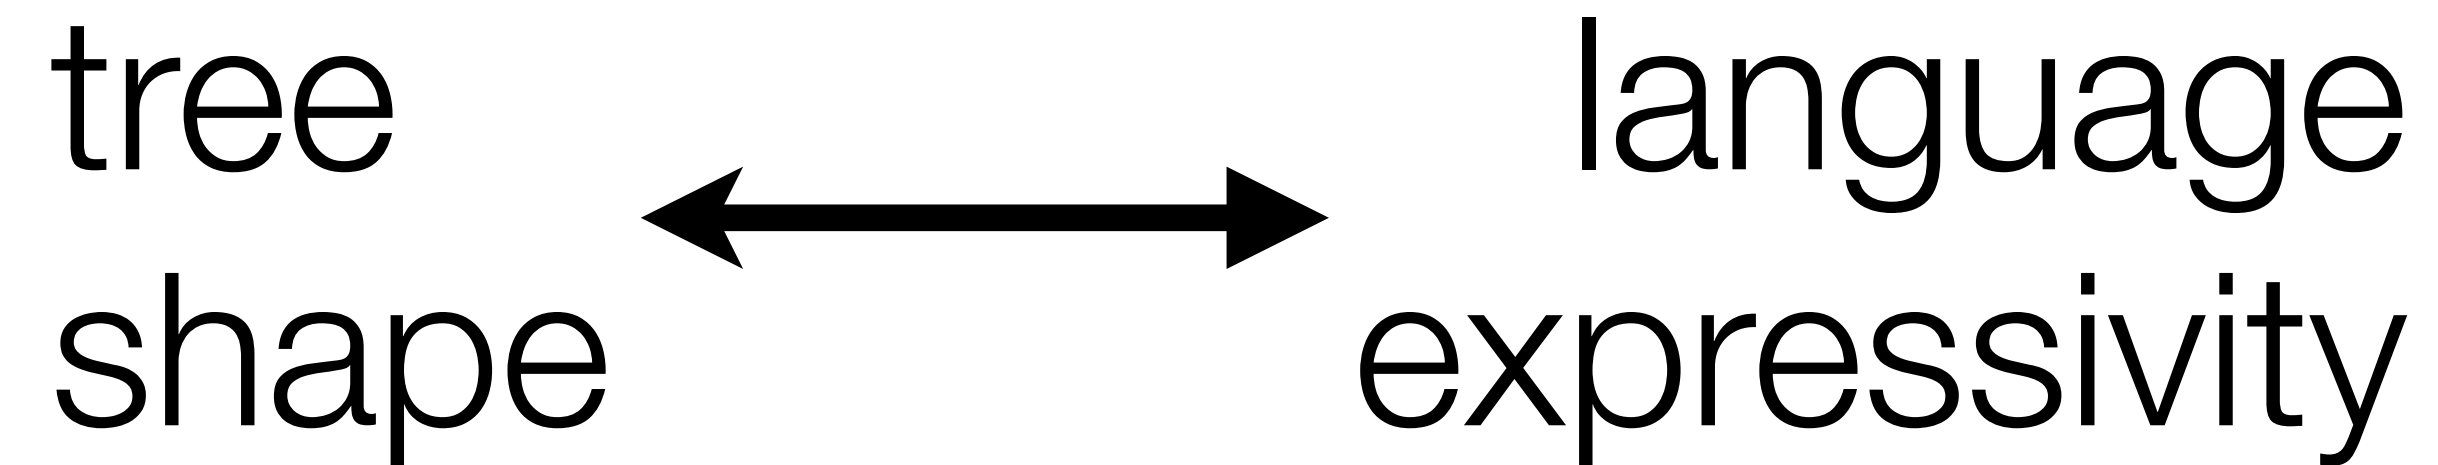
Path: $[(2, r_1), (B_1, r_2)]$



Which leads to some very PL-ey questions:

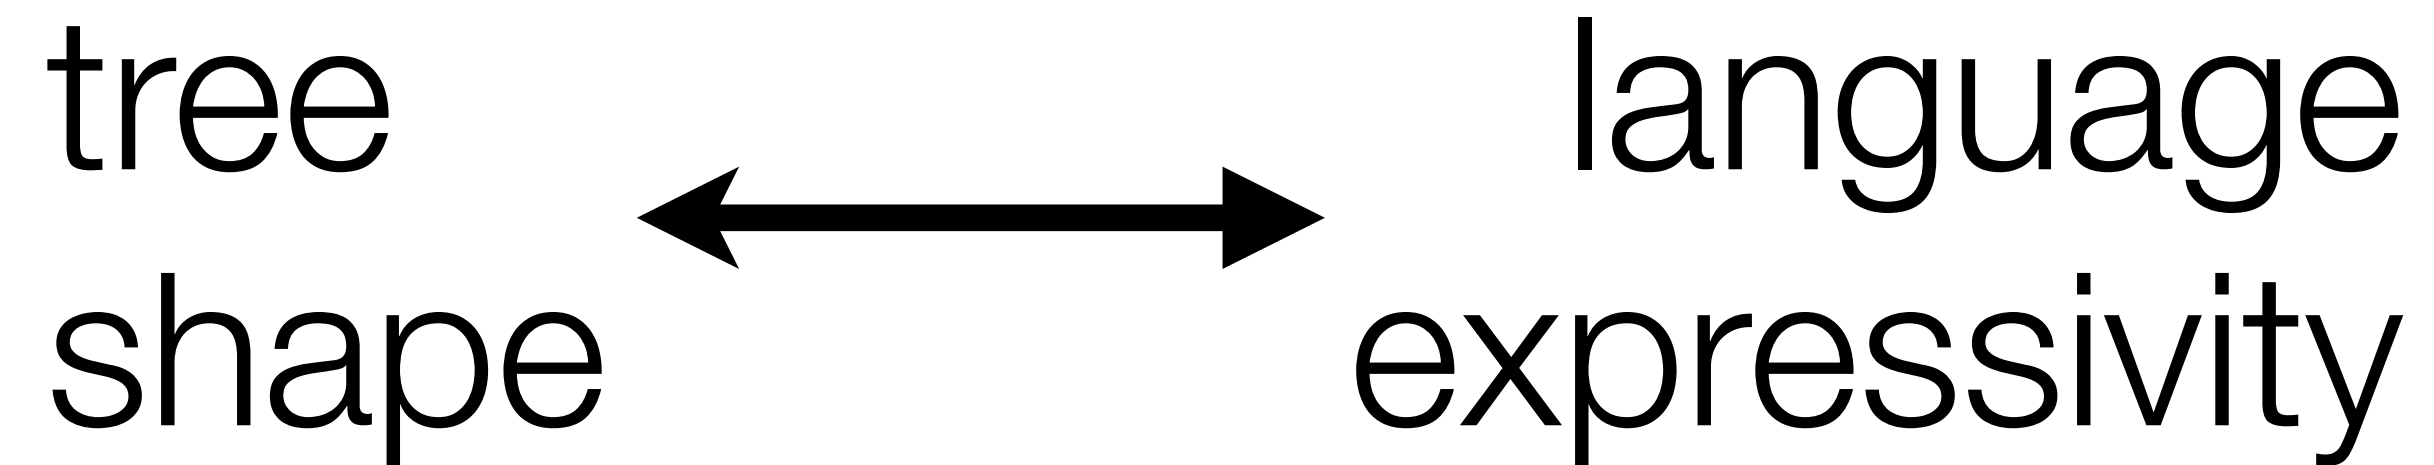


Which leads to some very PL-ey questions:



Compare expressivity of languages?

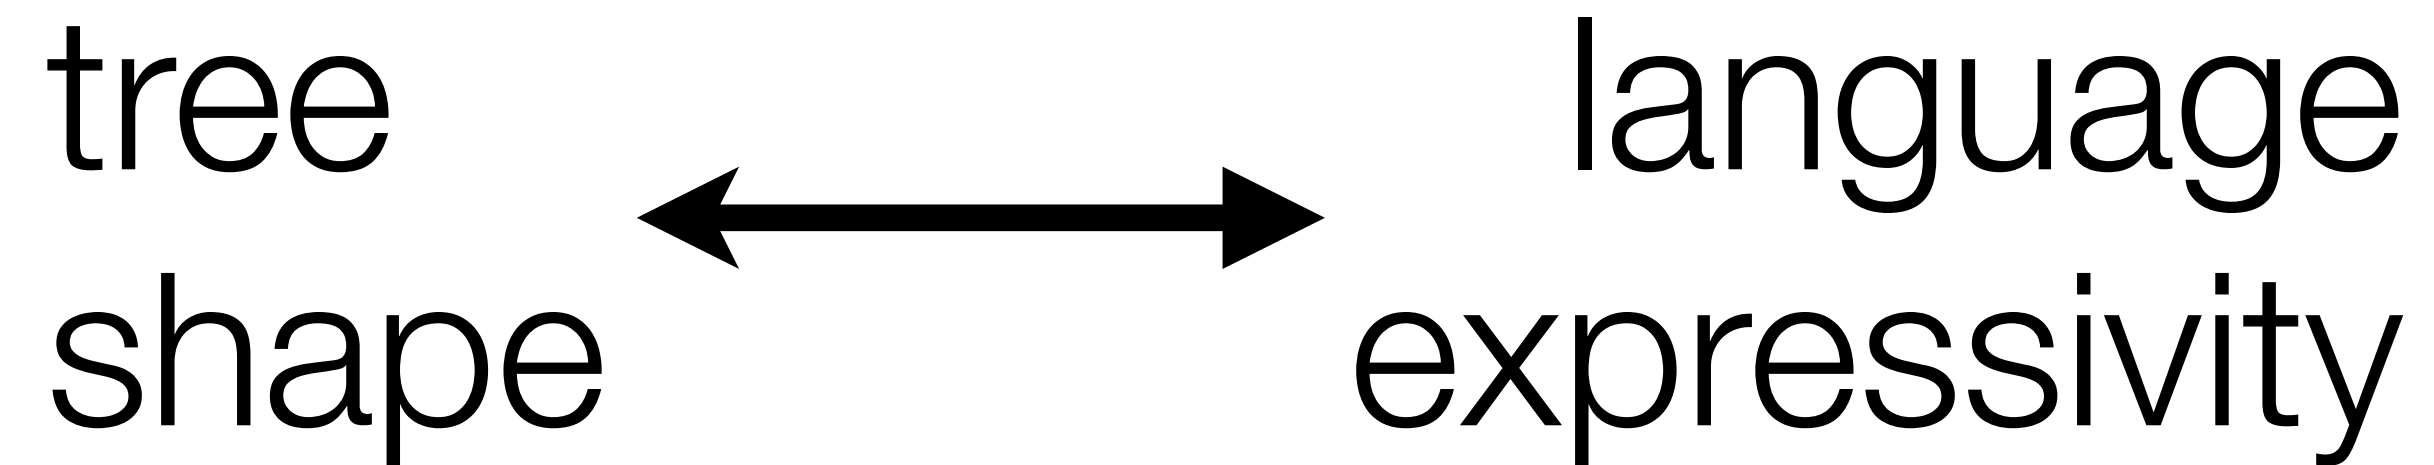
Which leads to some very PL-ey questions:



Compare expressivity of languages?

Compare expressivity of *trees*?

Which leads to some very PL-ey questions:



Compare expressivity of languages?

Compare expressivity of *trees*?

Compile a program so it runs against a new tree?

~~No general way to deploy our gadget.~~



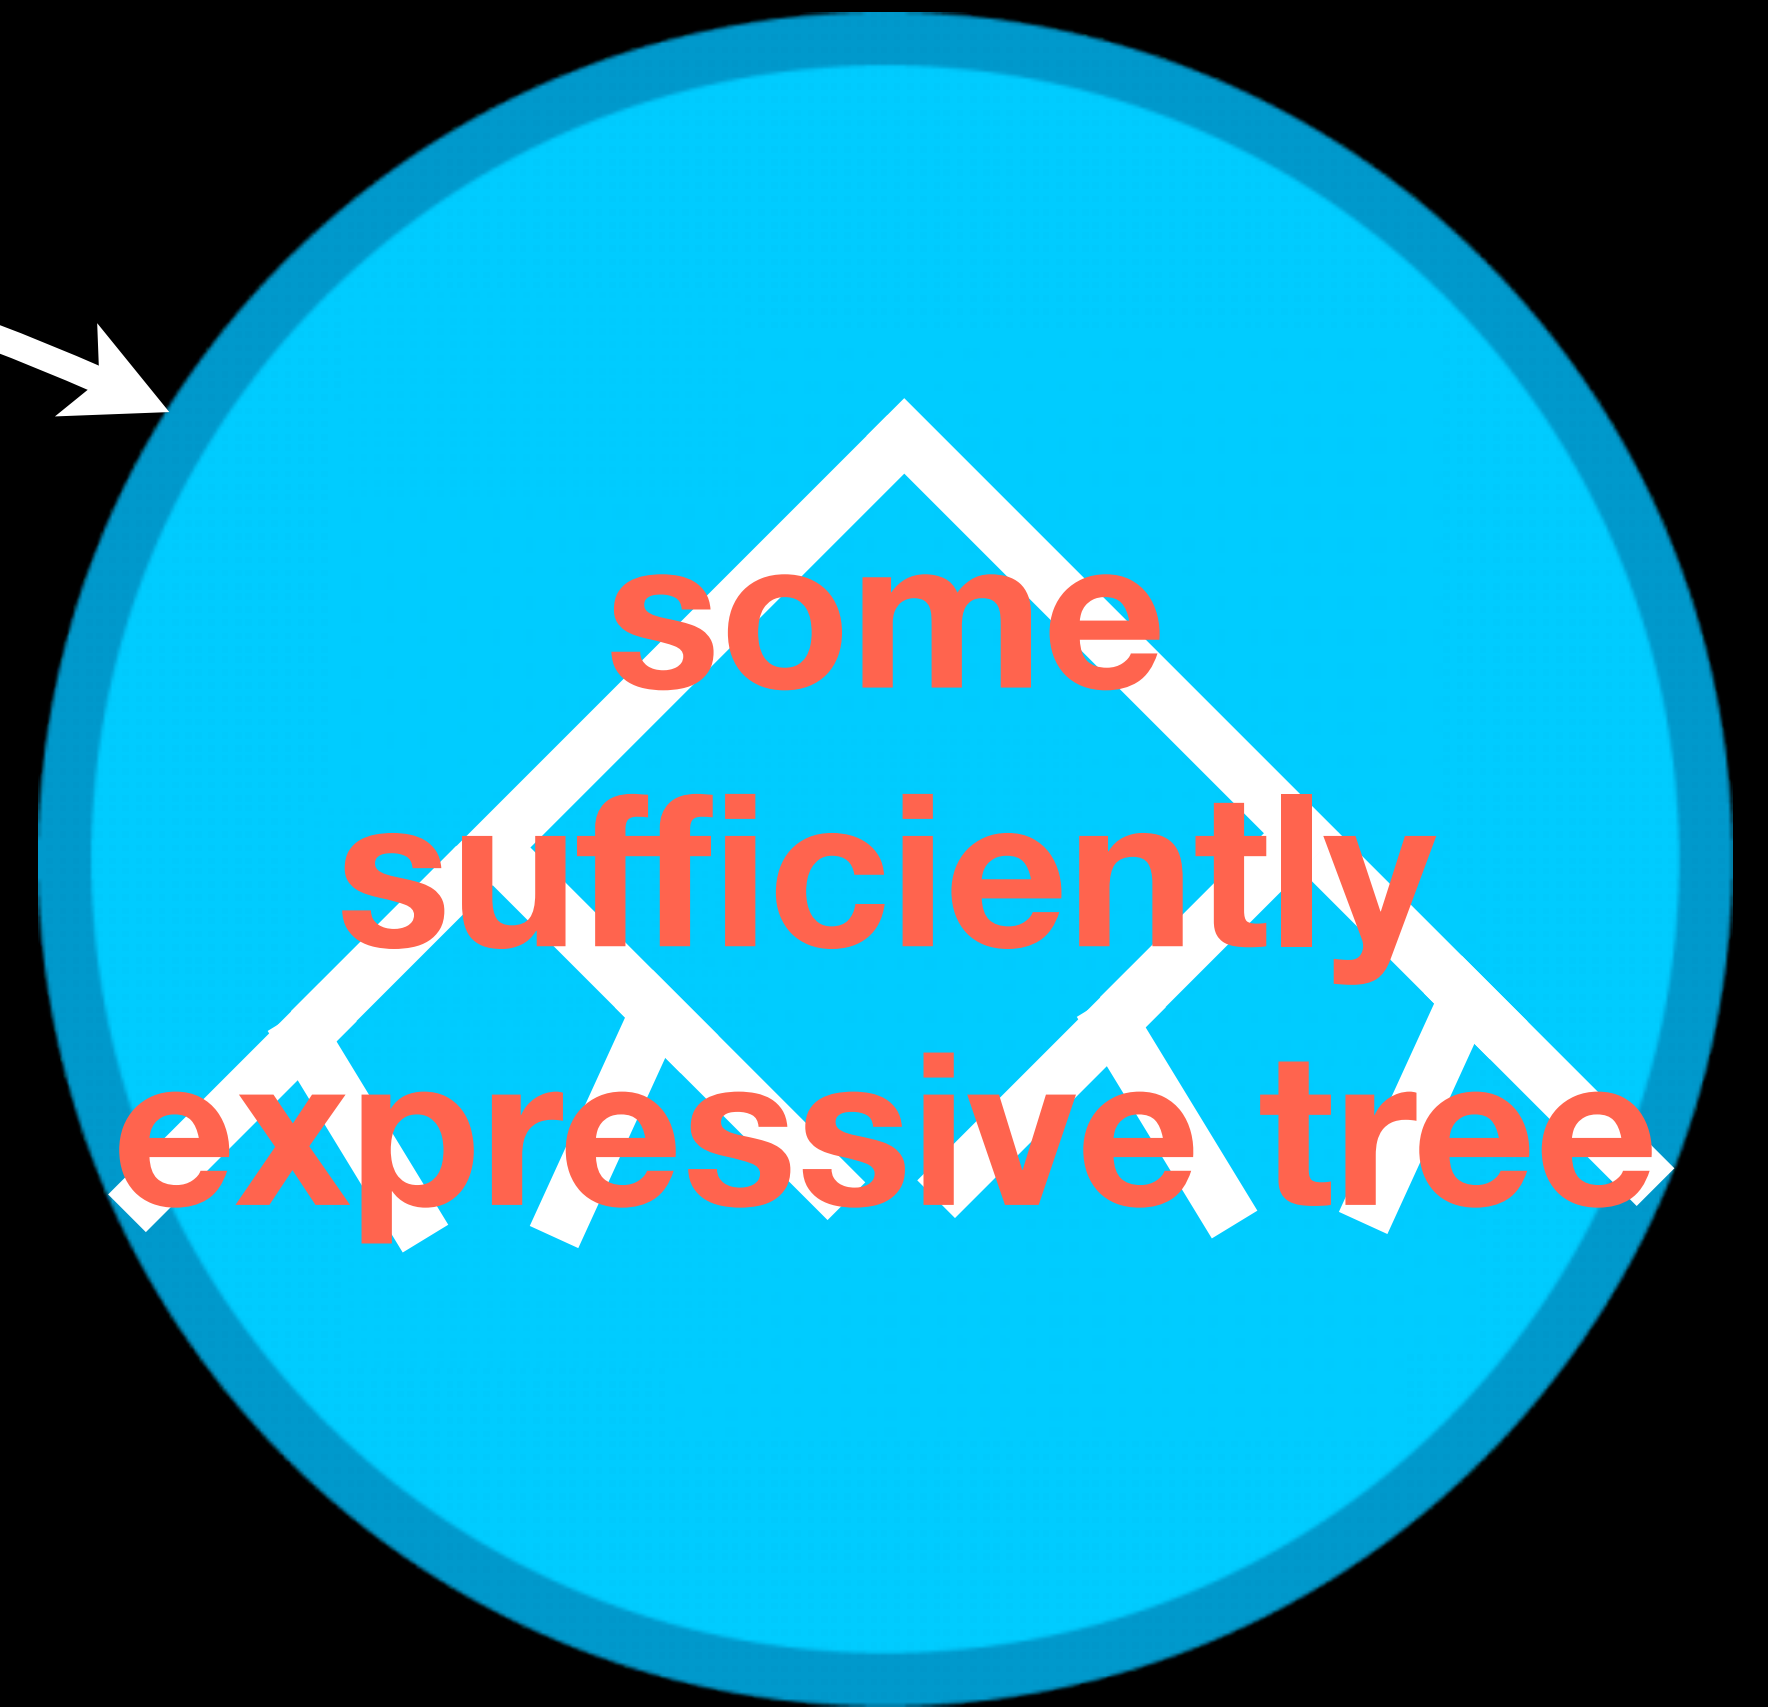
A human needs a
range of trees.

The hardware wants
to support *one* tree.

~~No general way to deploy our gadget.~~

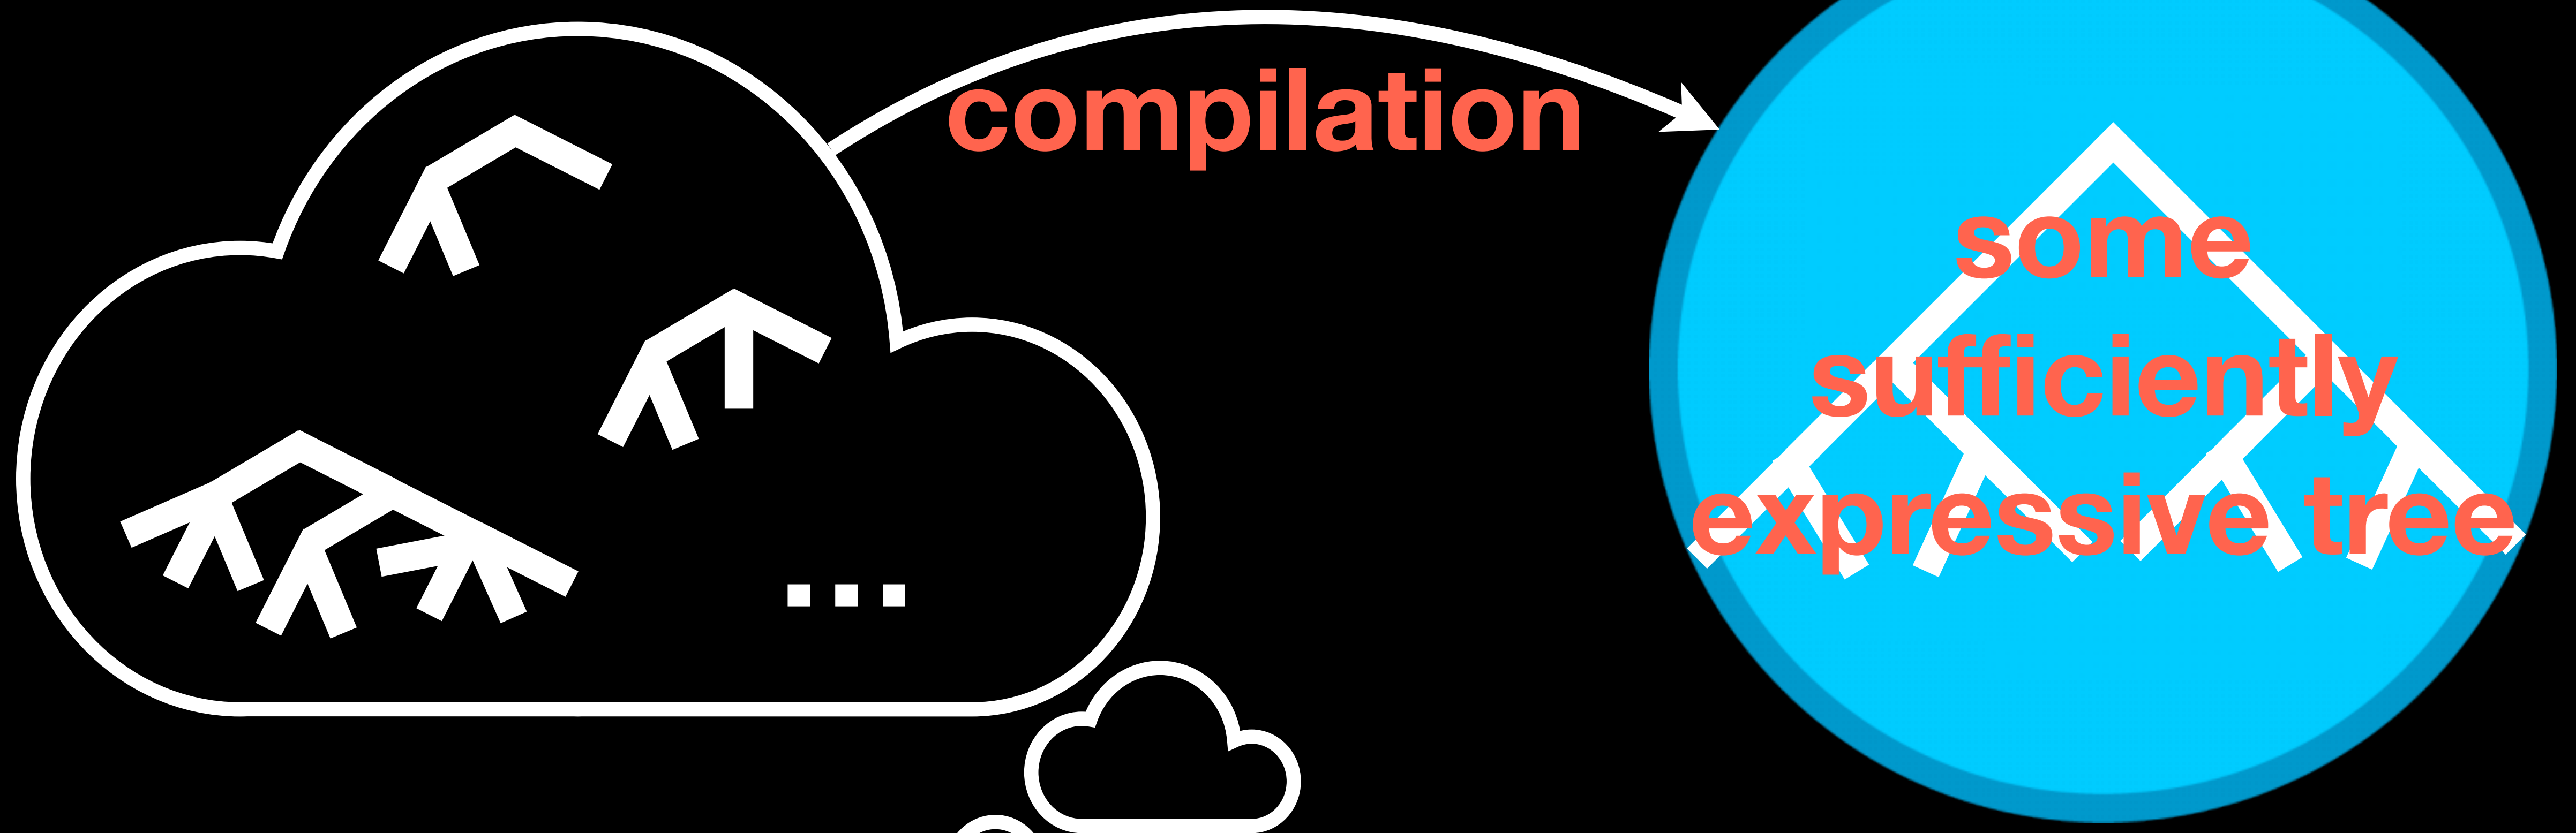


A human needs a
range of trees.



The hardware wants
to support *one* tree.

~~No general way to deploy our gadget.~~



A human needs a *range* of trees.

The hardware wants to support *one* tree.

Contributions

Contributions

Formal model of PIFO trees

Contributions

Formal model of PIFO trees

General theorems of expressiveness
w.r.t. tree shape

Contributions

Formal model of PIFO trees

General theorems of expressiveness
w.r.t. tree shape

Compiler

Contributions

Formal model of PIFO trees

General theorems of expressiveness
w.r.t. tree shape

Compiler

Simulator

Expressivity of trees

Trees with more leaves are more expressive.

Taller trees are more expressive.

Expressivity of trees

Trees with more leaves are more expressive.

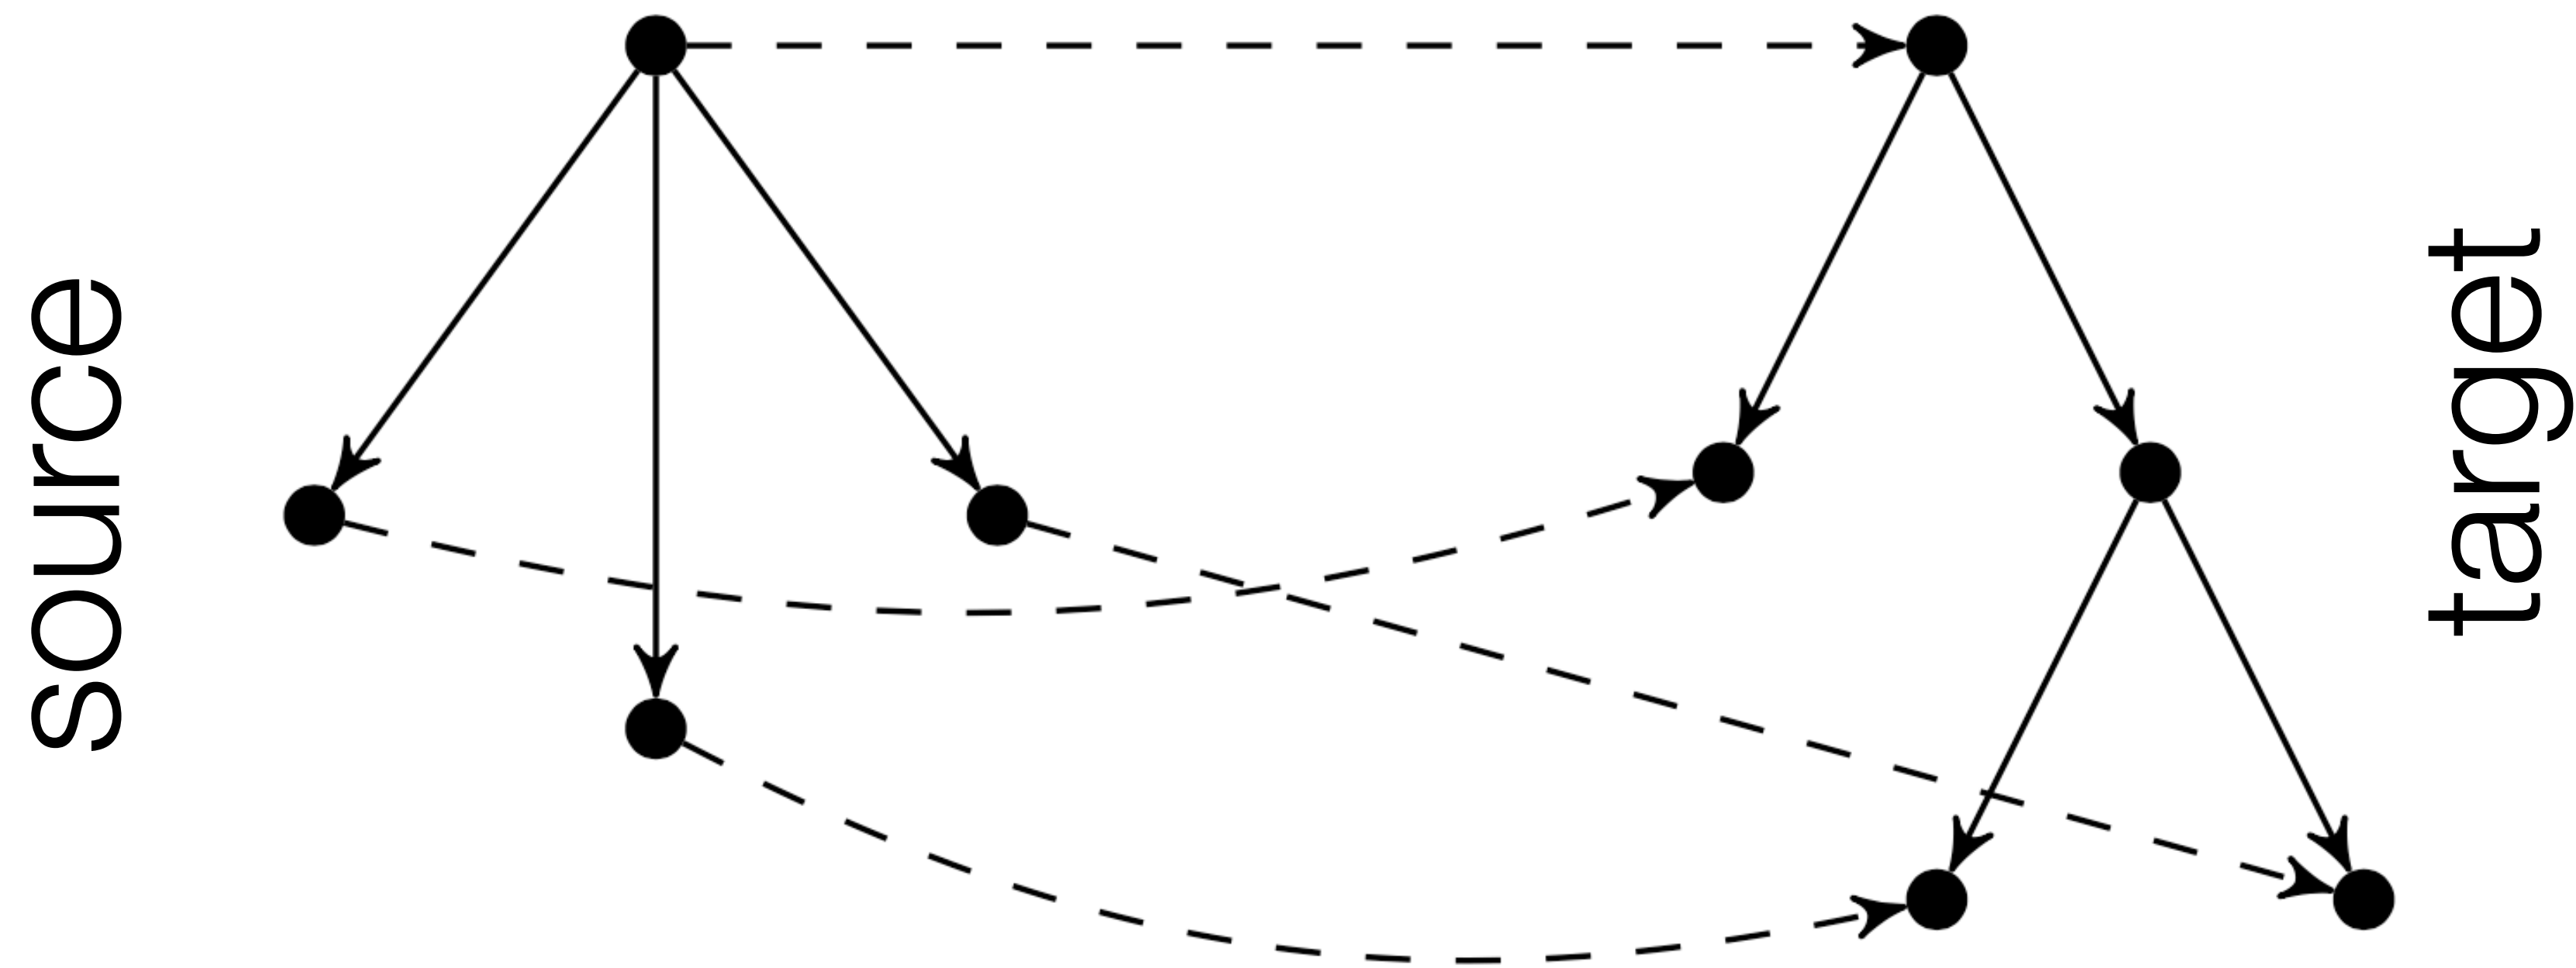
Taller trees are more expressive.

Captured elegantly by:

Homomorphic embedding.

Map root to root, leaves to leaves. Respect ancestry.

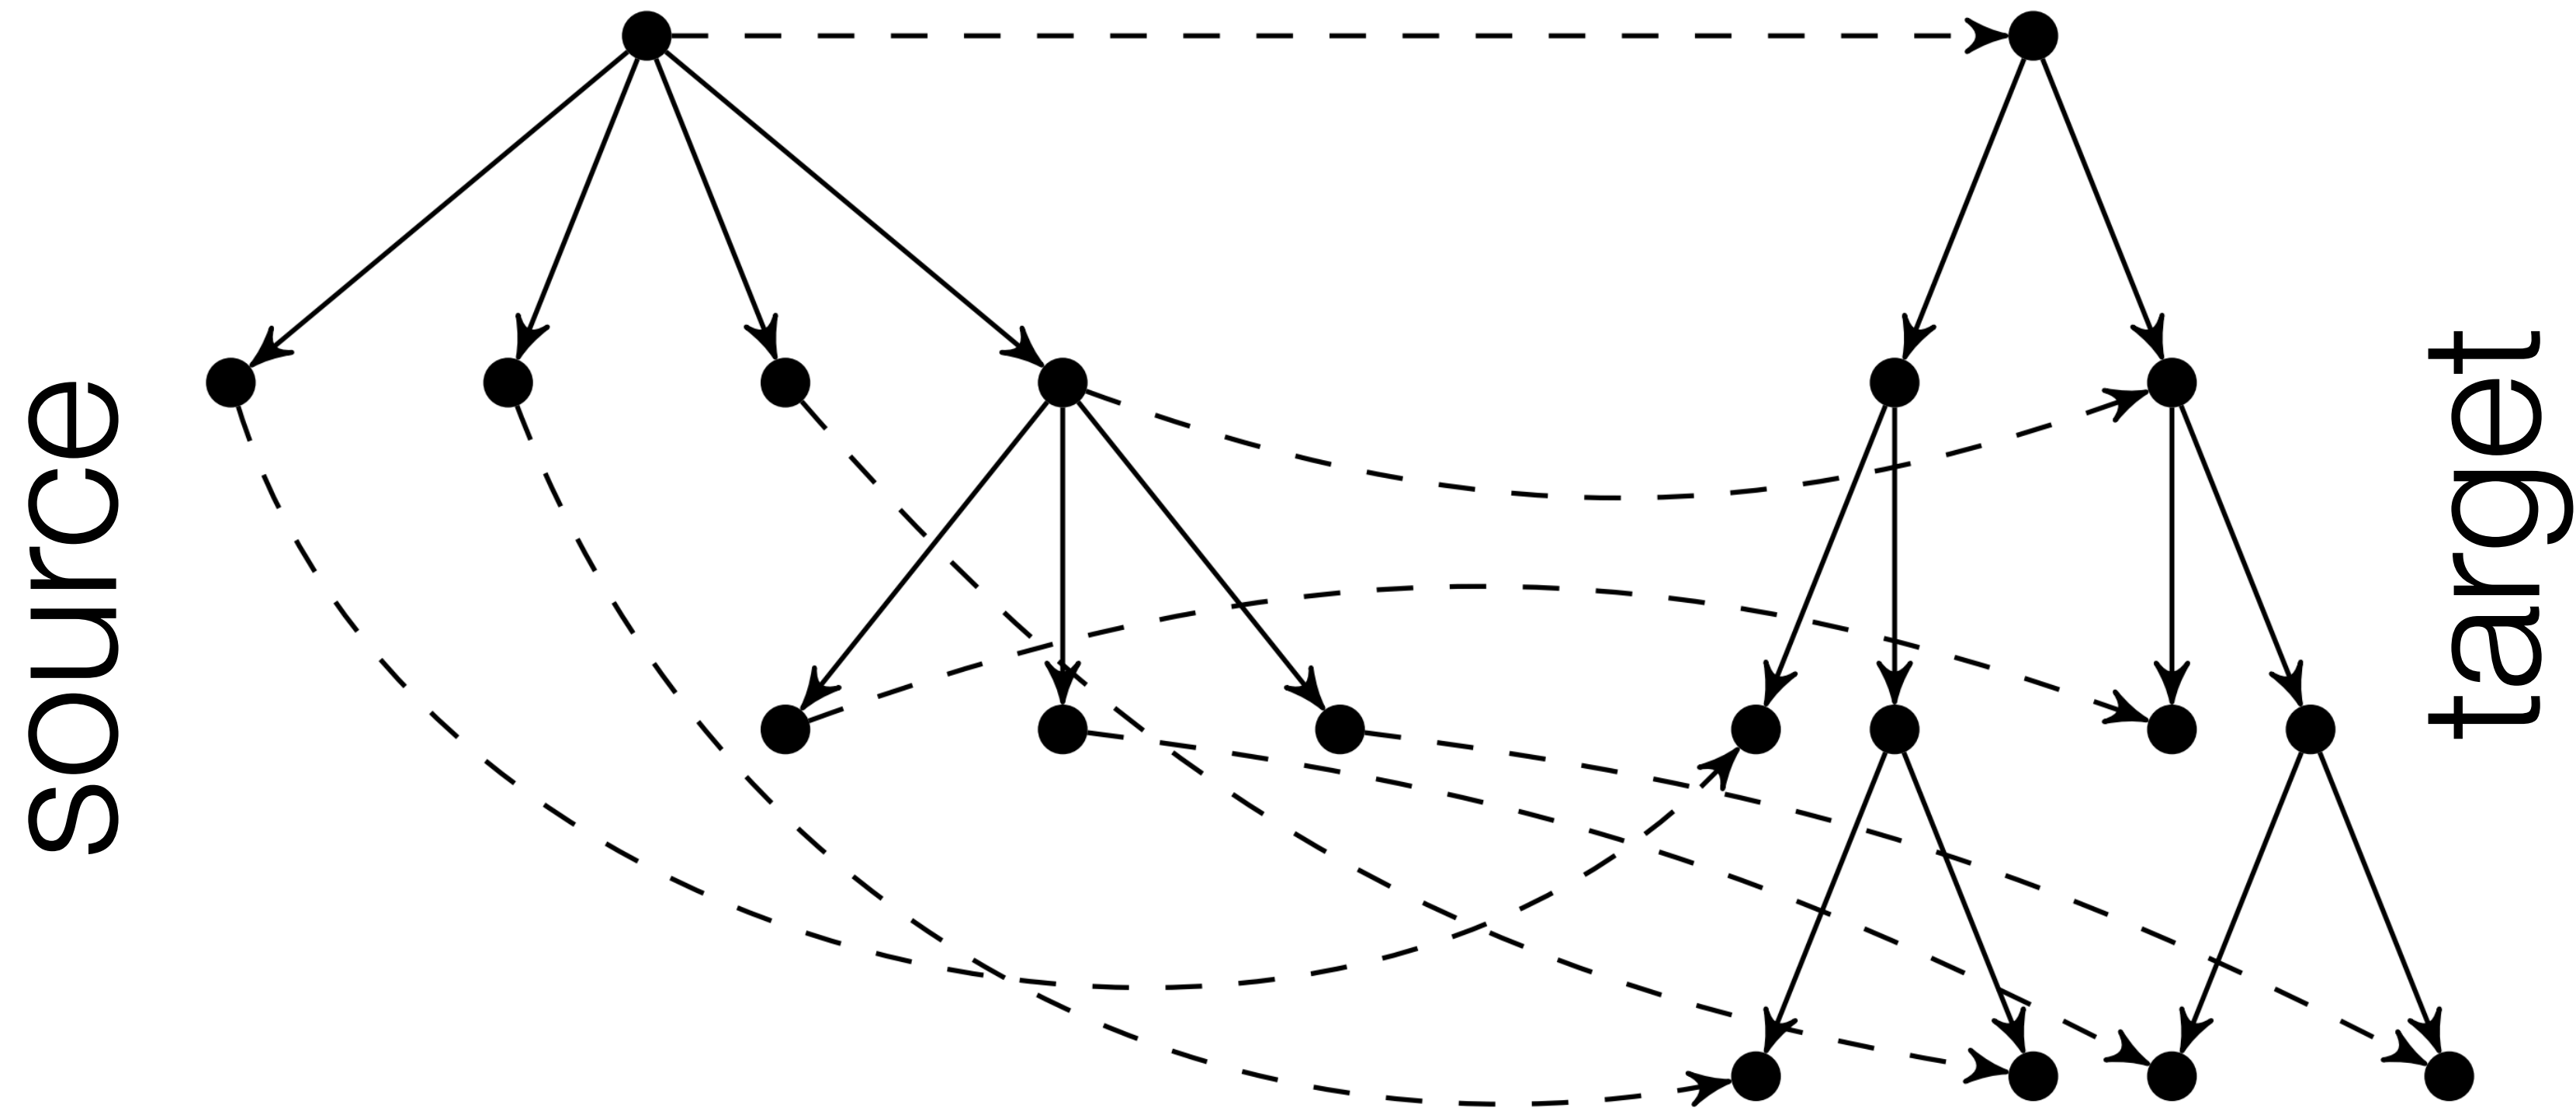
Expressivity of trees



Homomorphic embedding.

Map root to root, leaves to leaves. Respect ancestry.

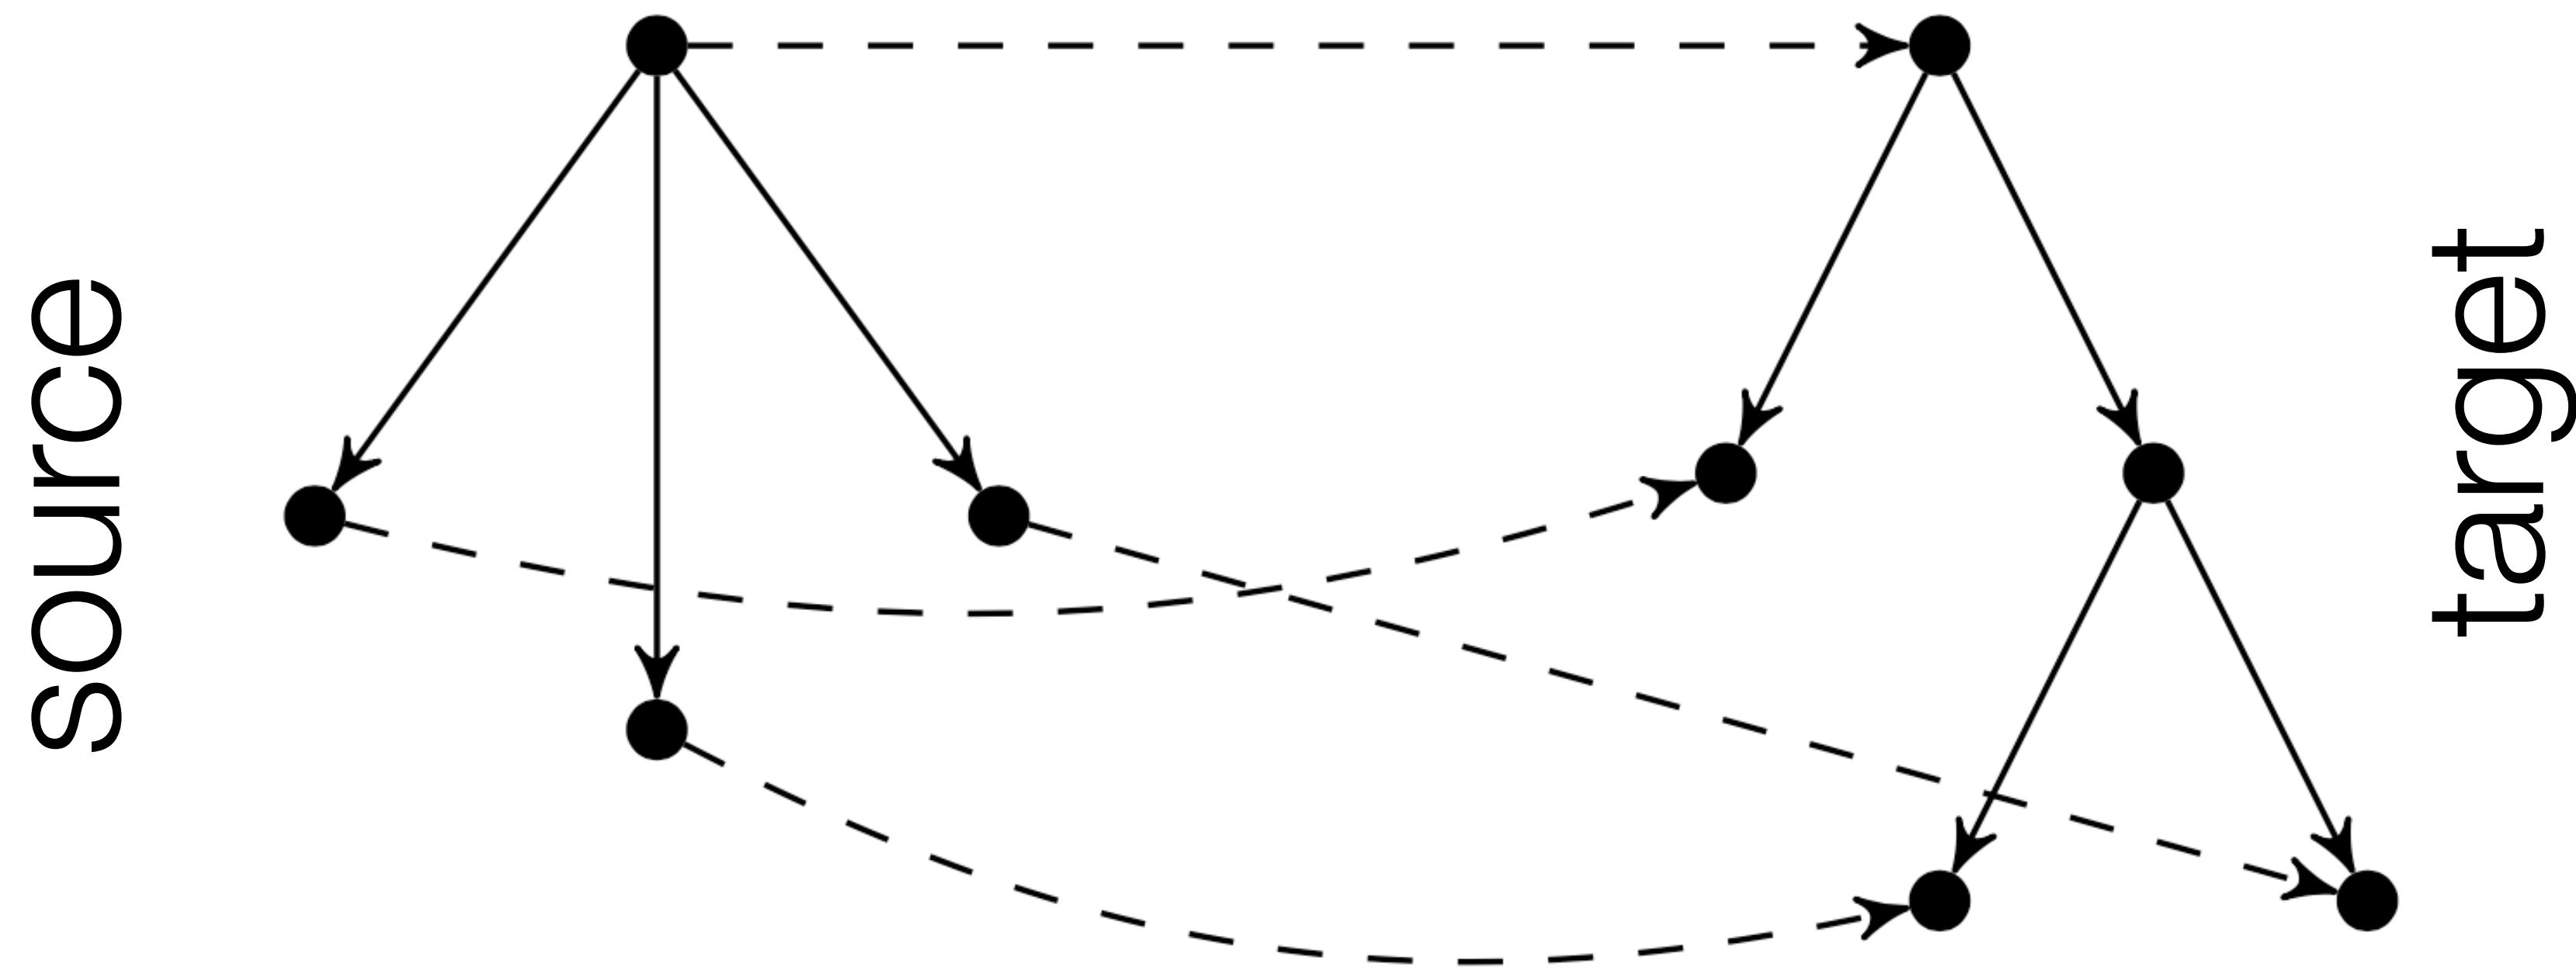
Expressivity of trees



Homomorphic embedding.

Map root to root, leaves to leaves. Respect ancestry.

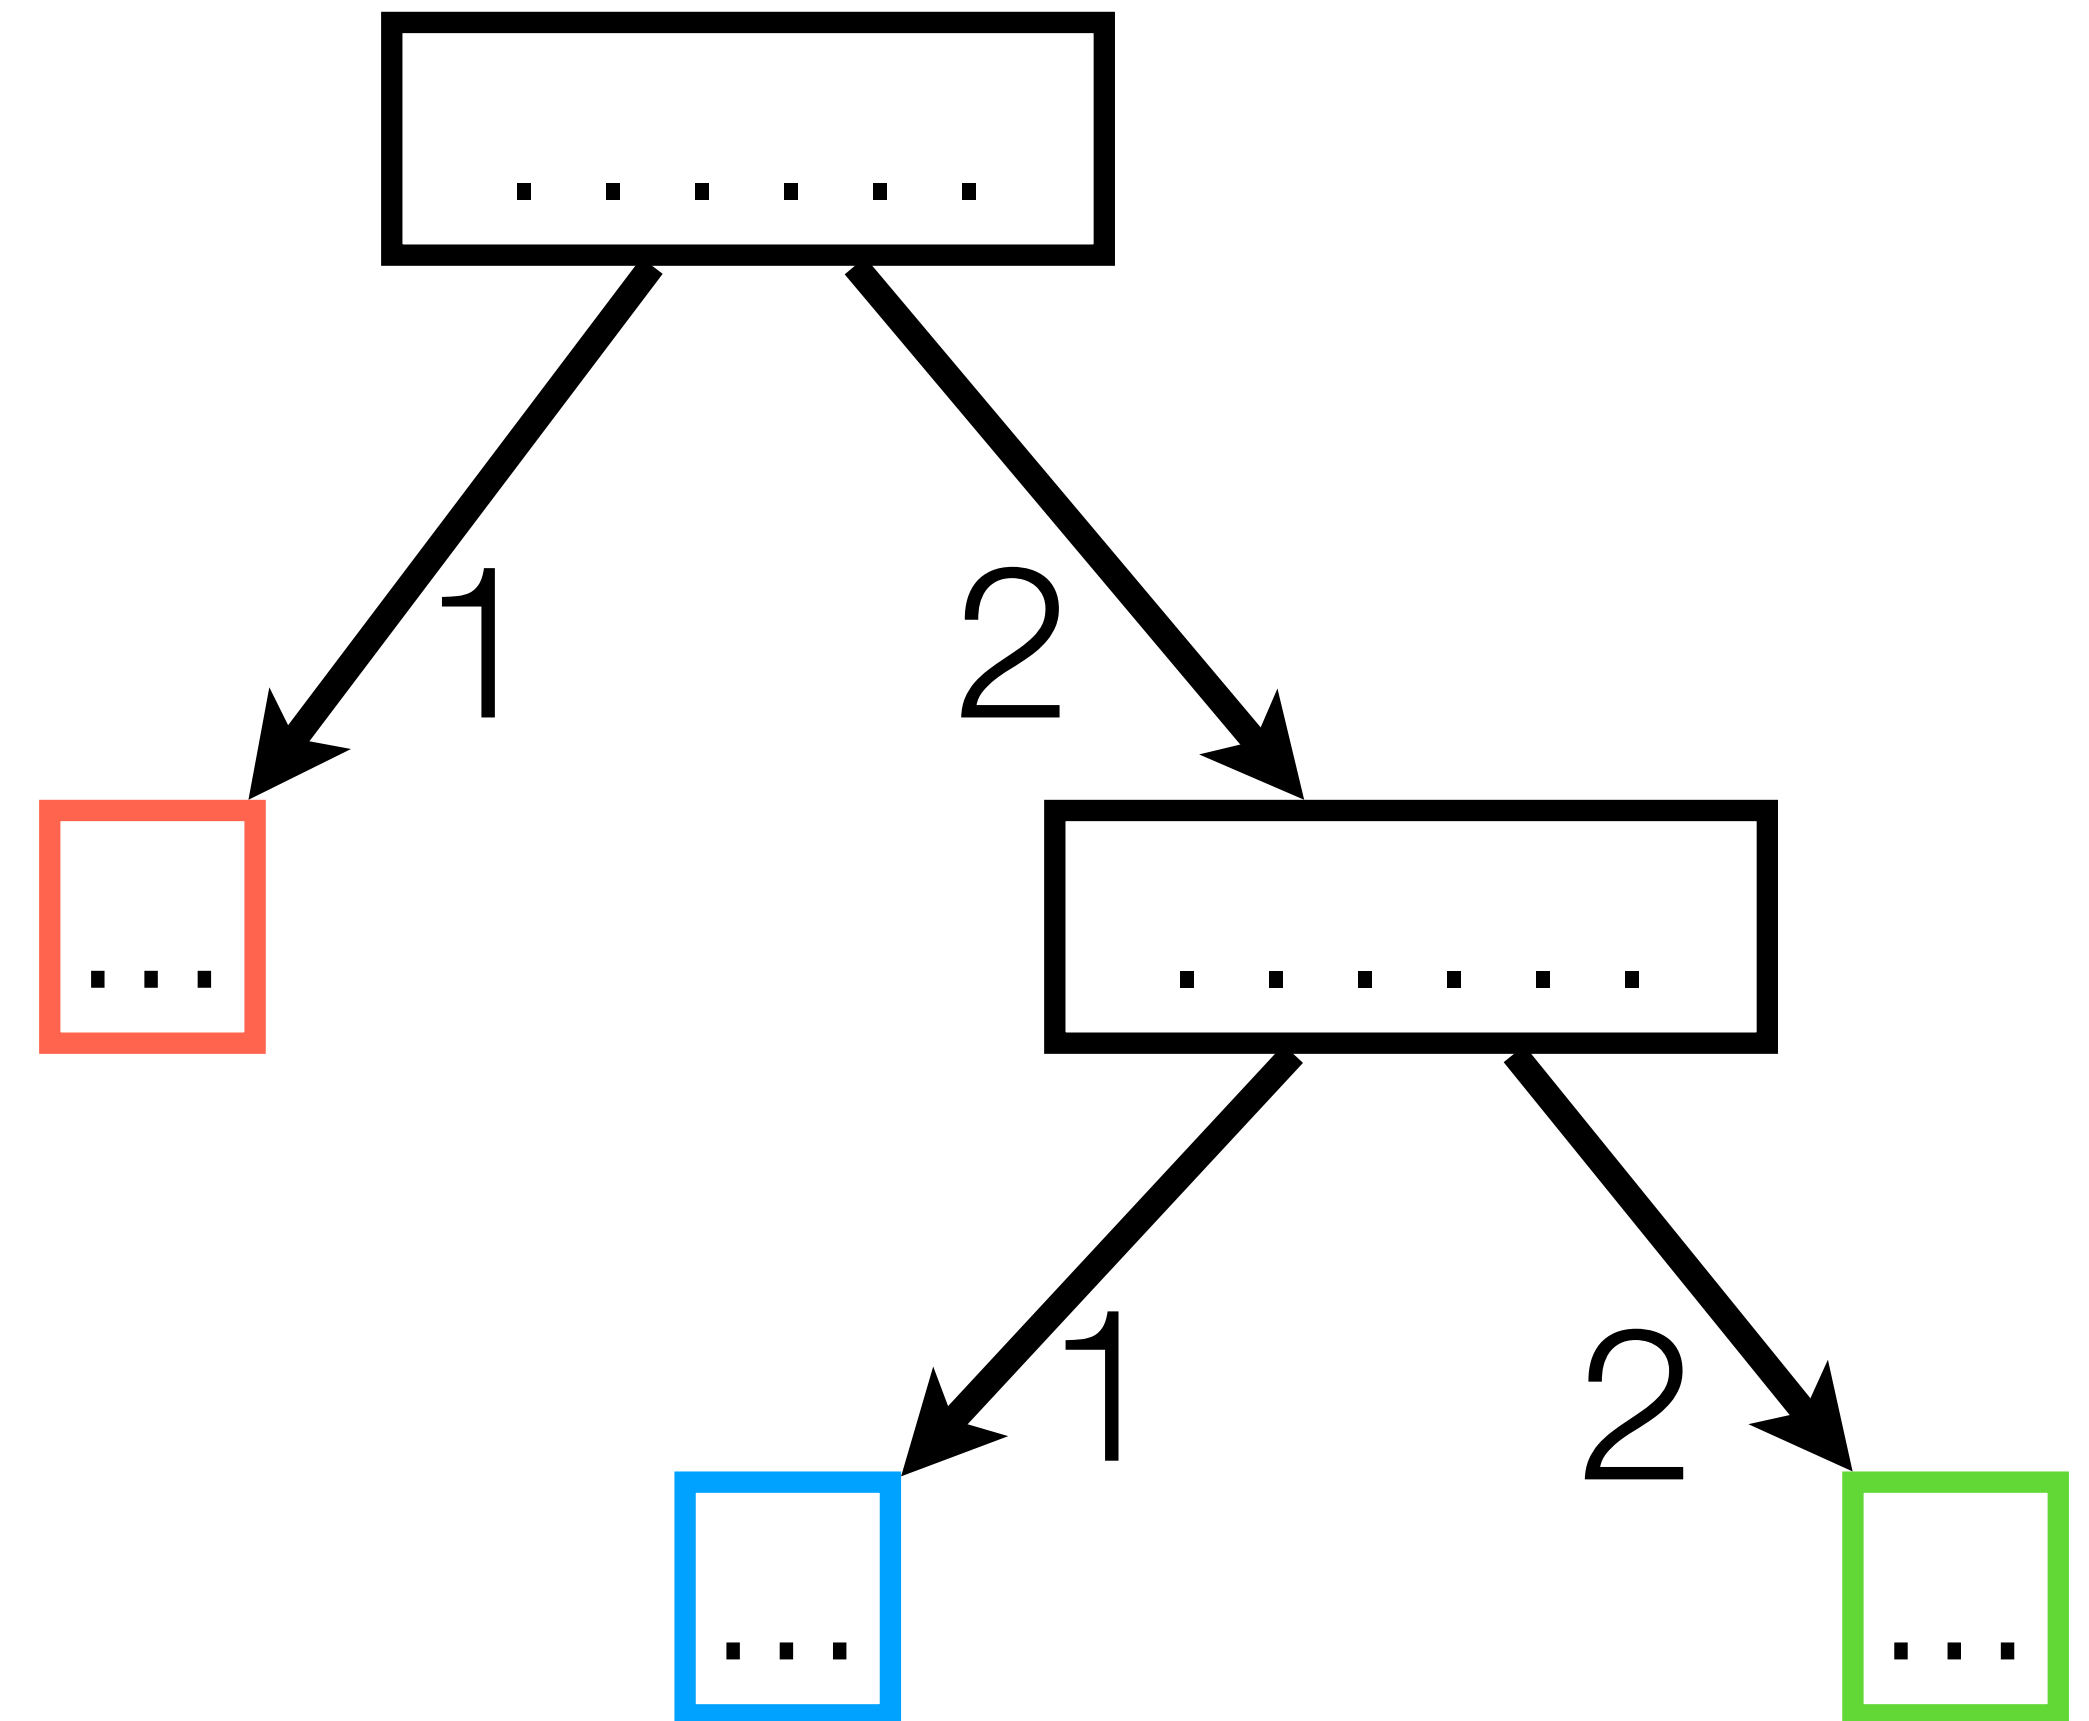
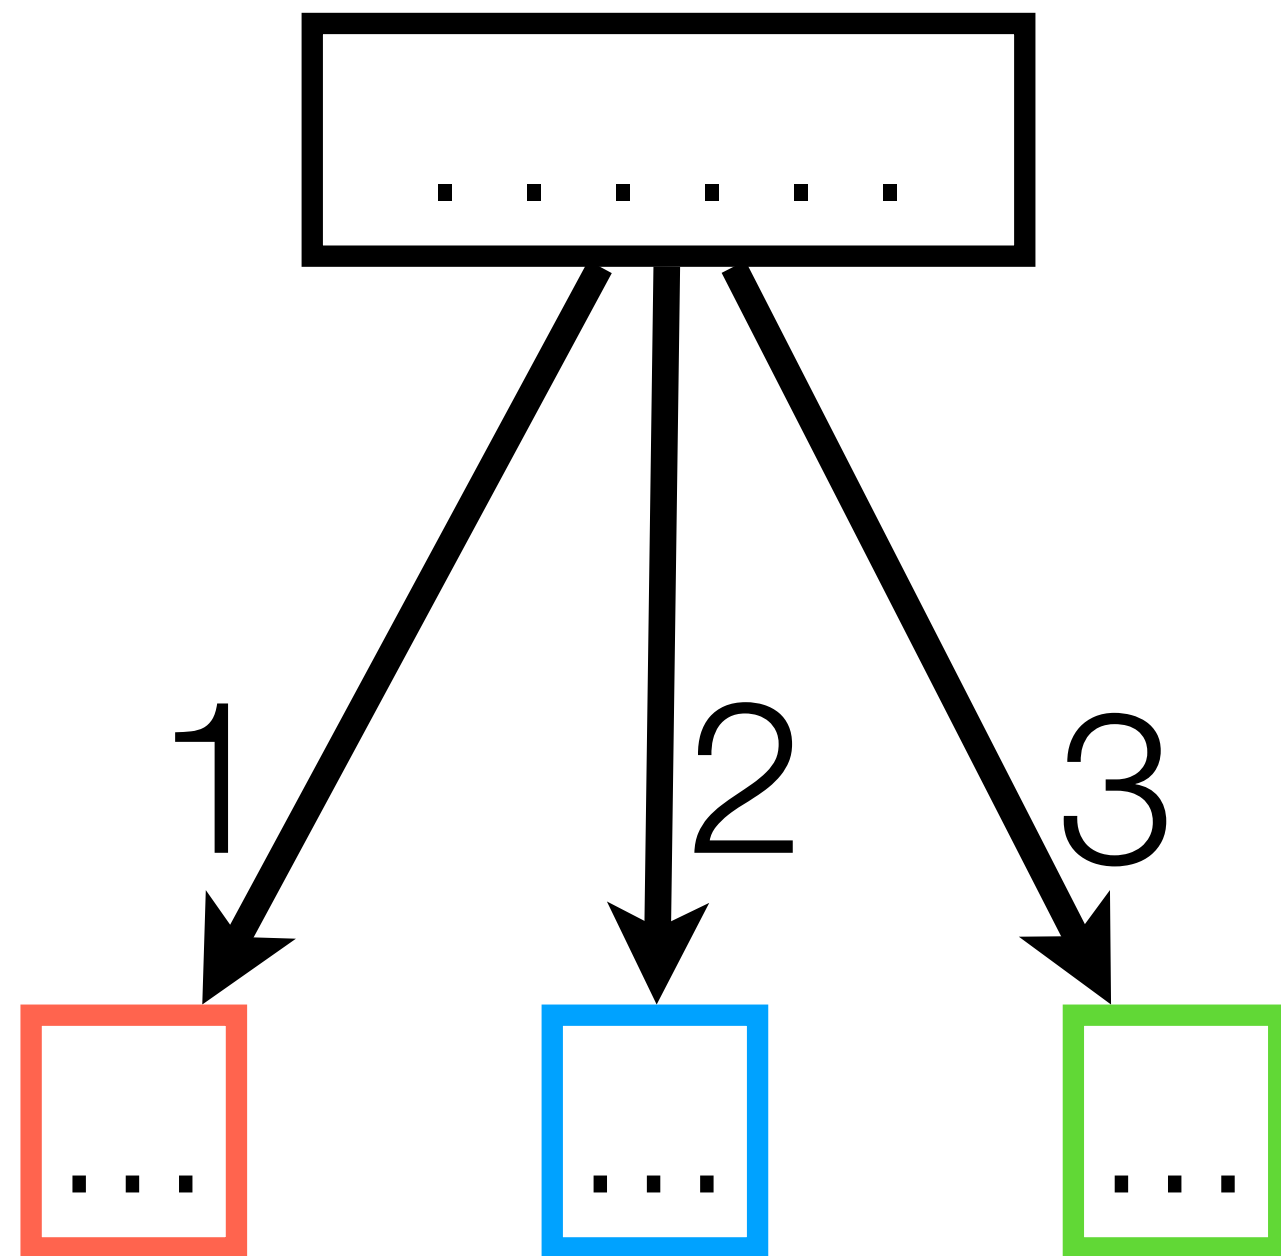
Expressivity of trees



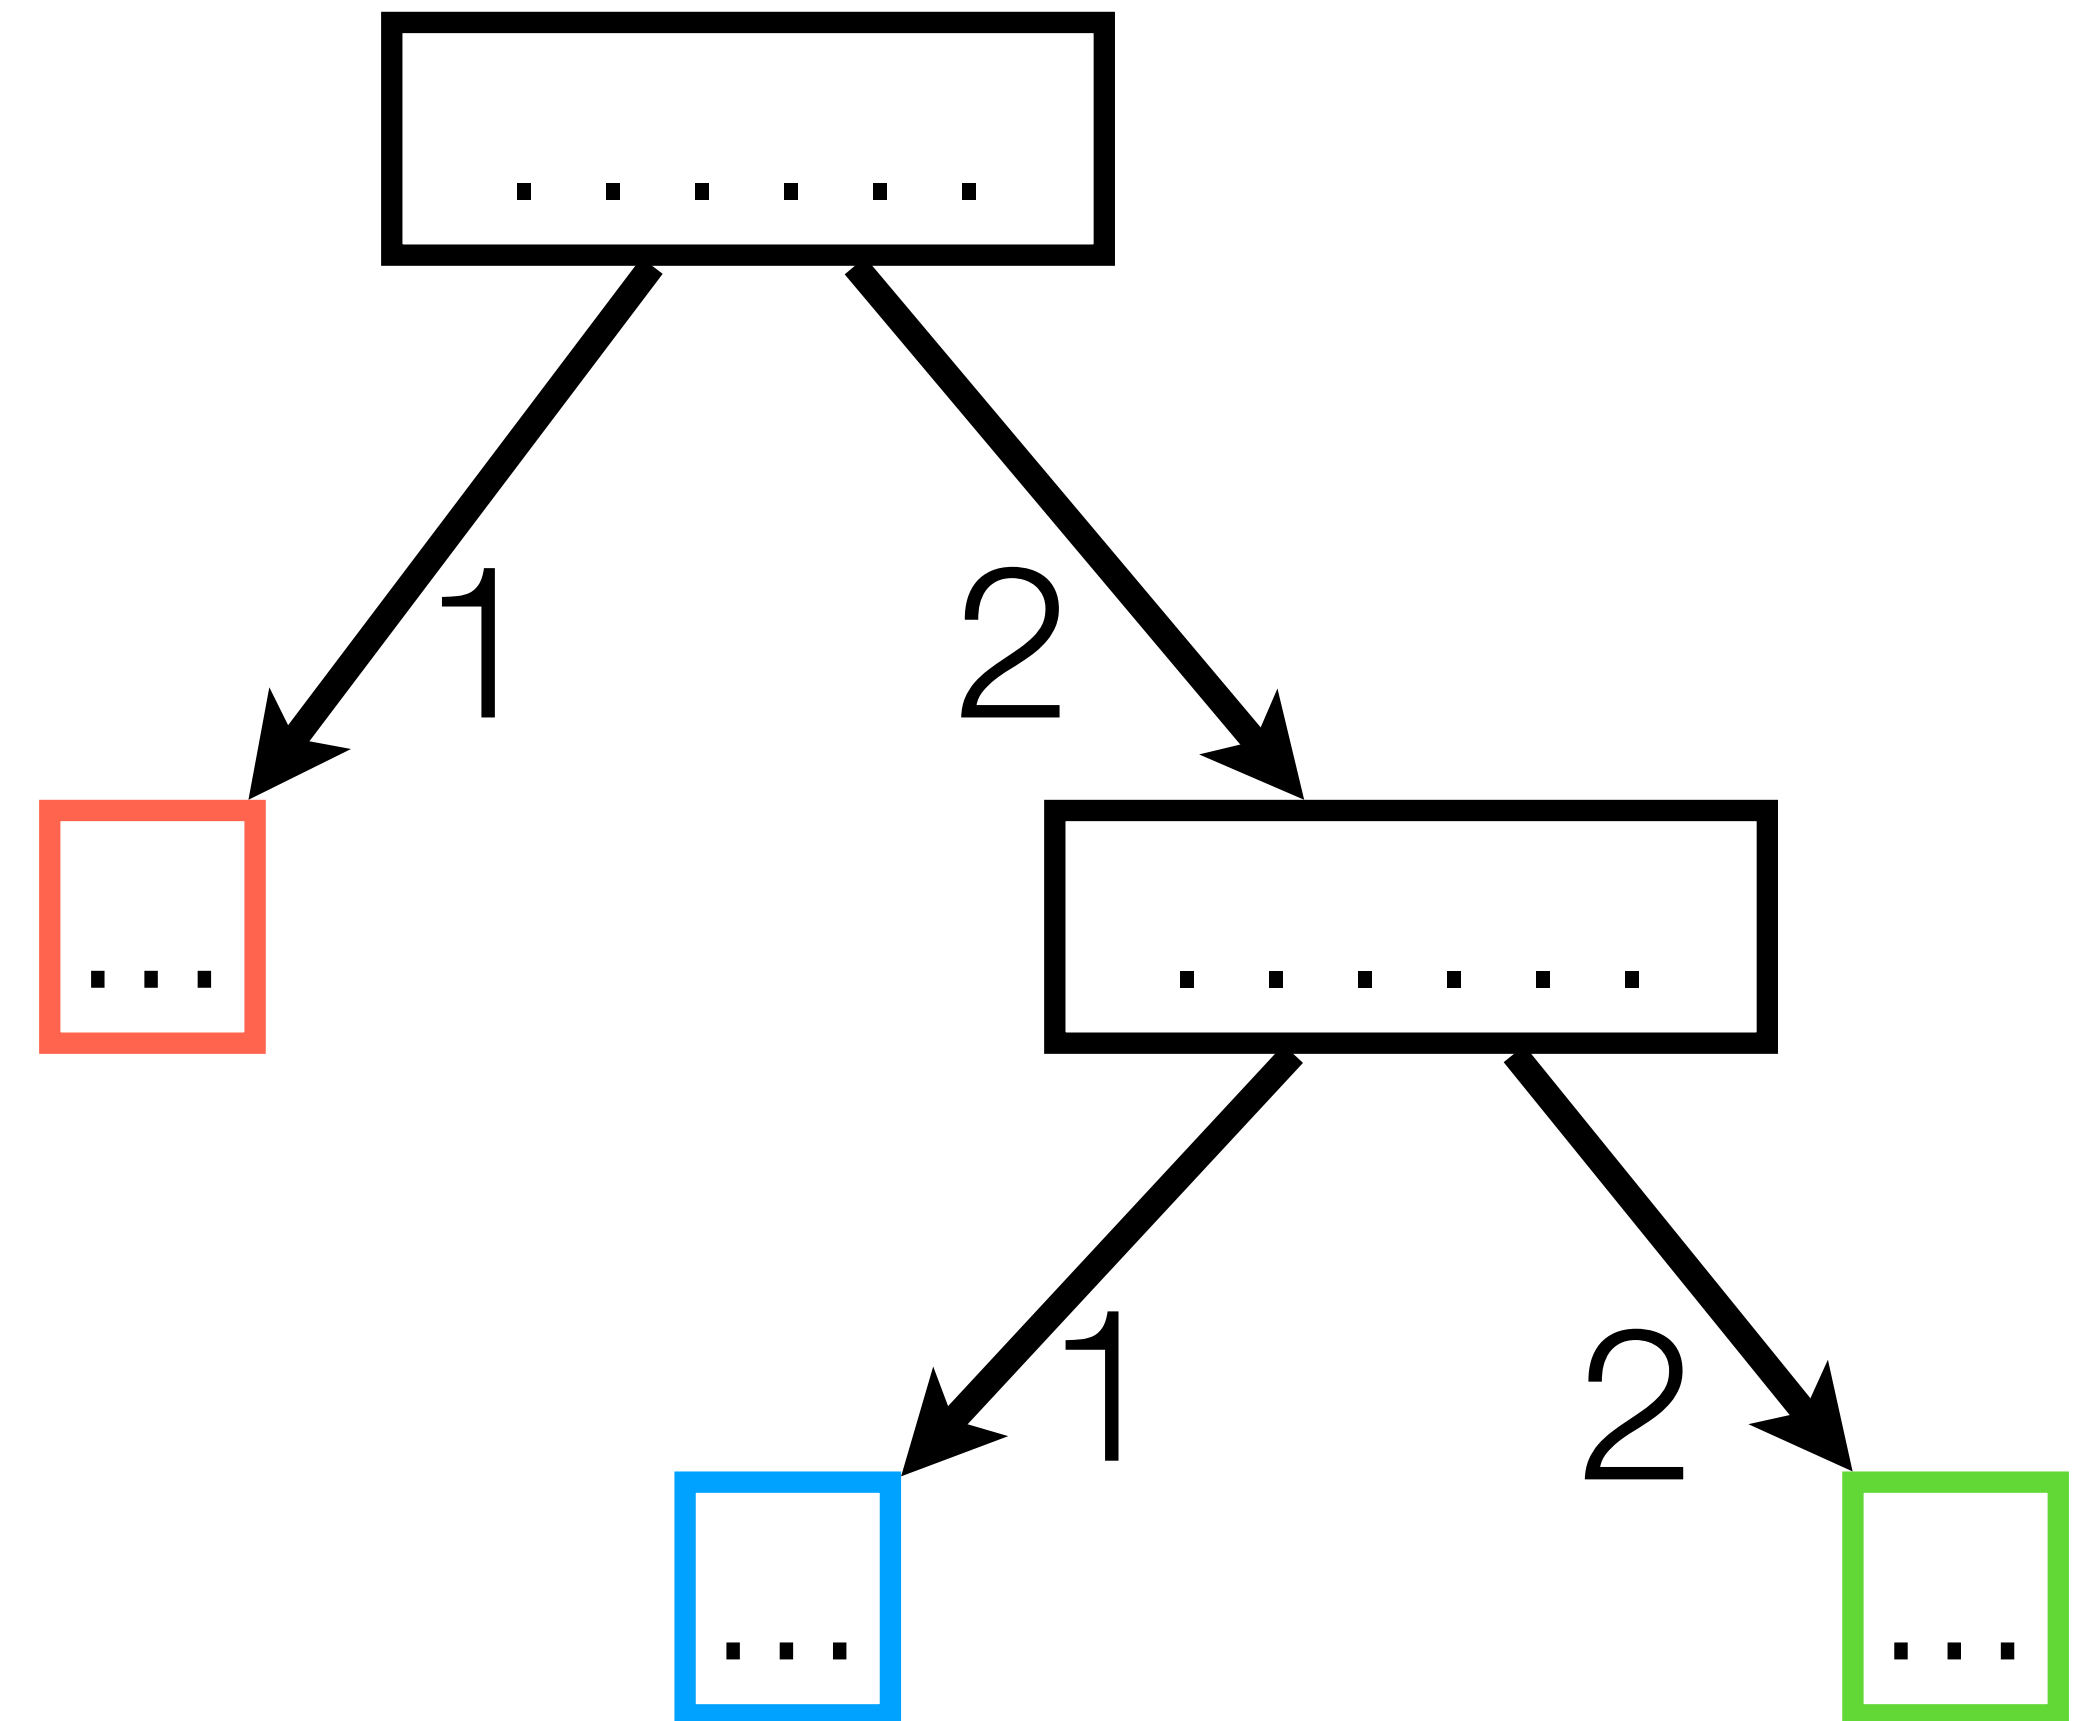
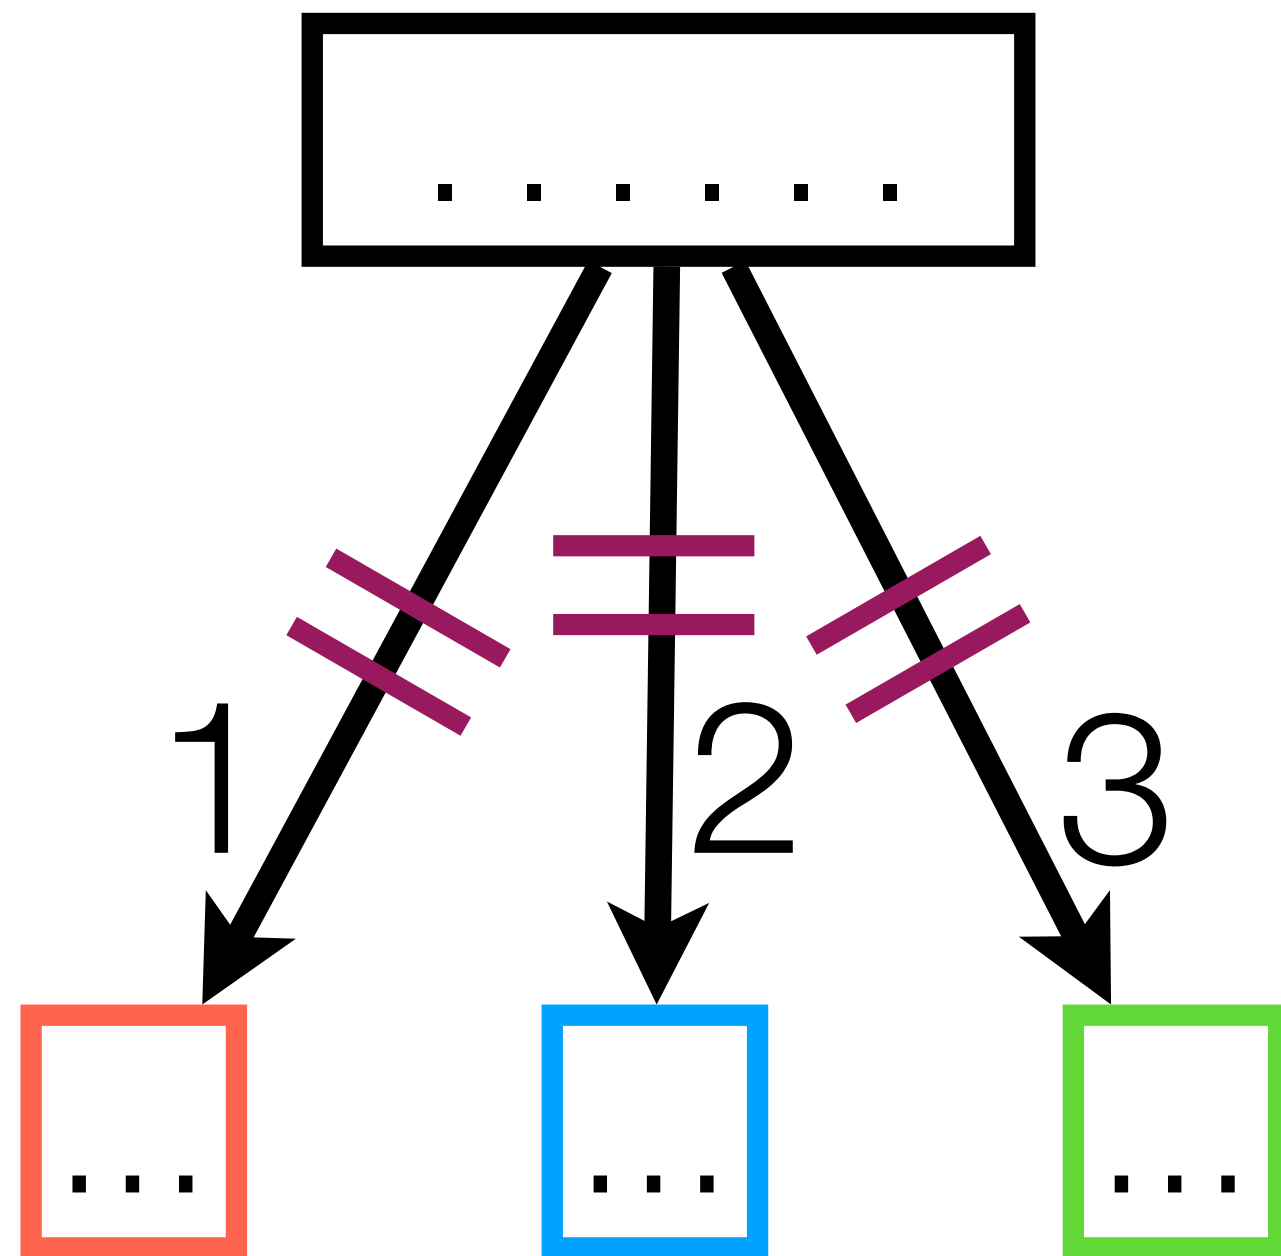
Homomorphic embedding.

Map root to root, leaves to leaves. Respect ancestry.

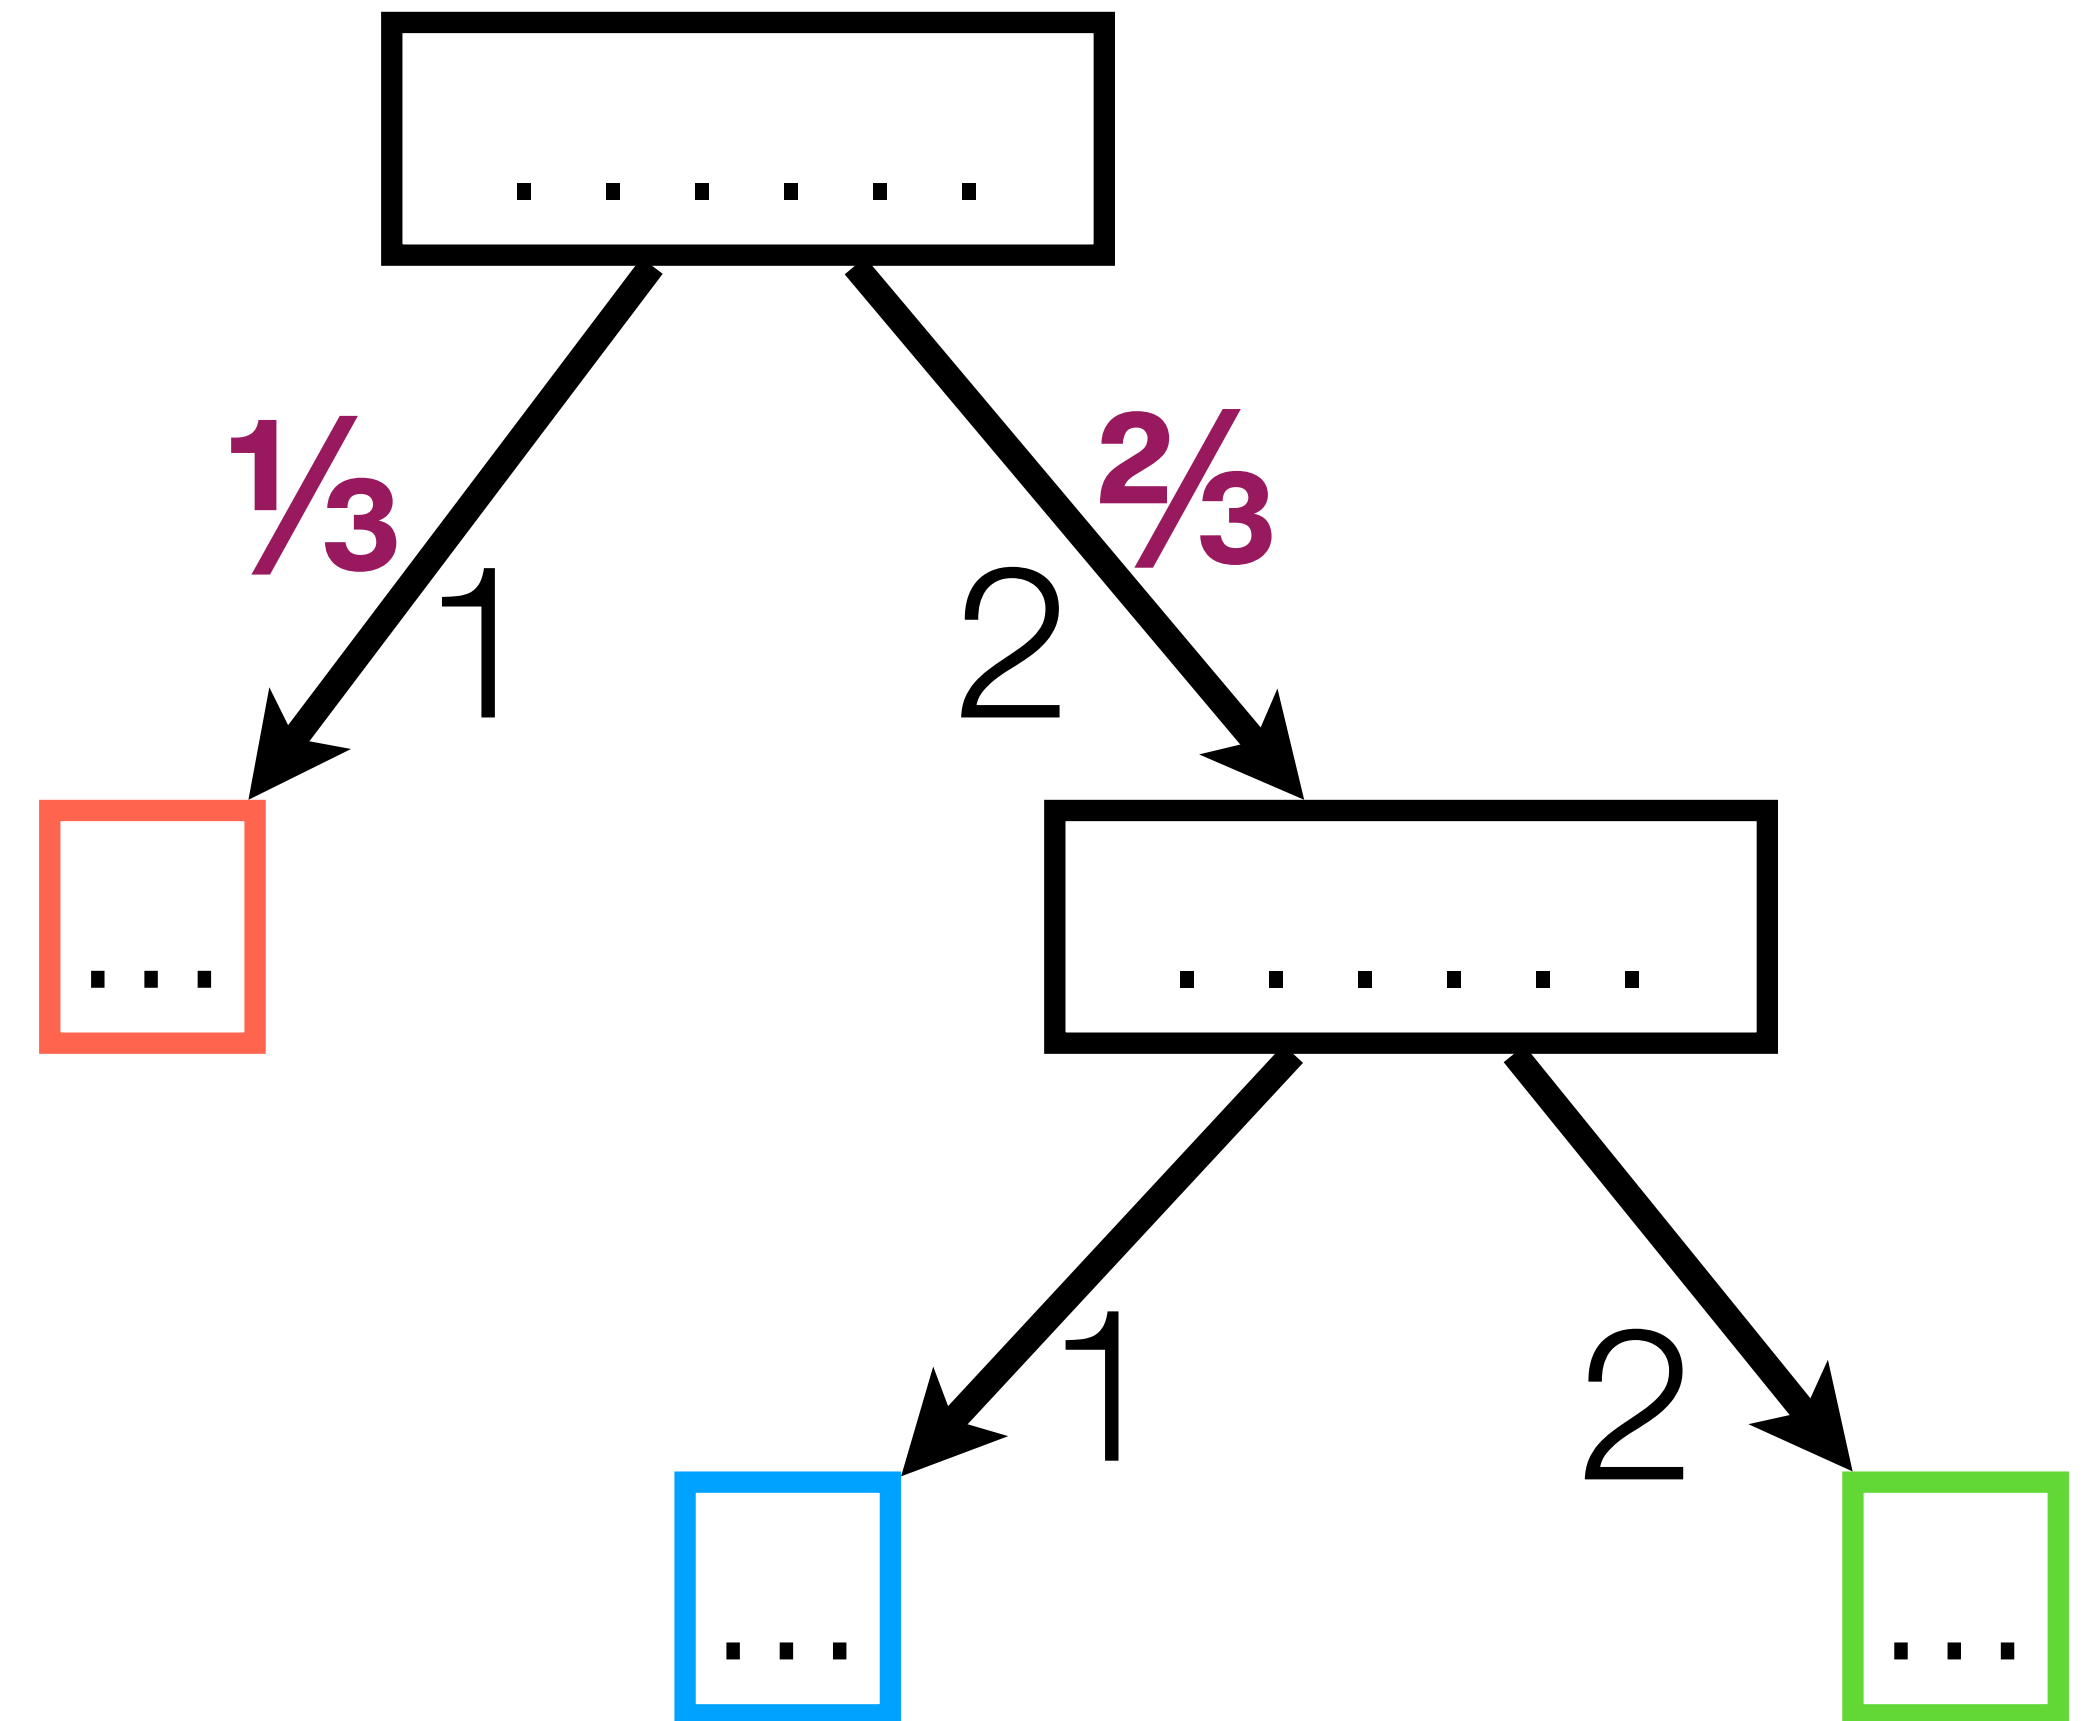
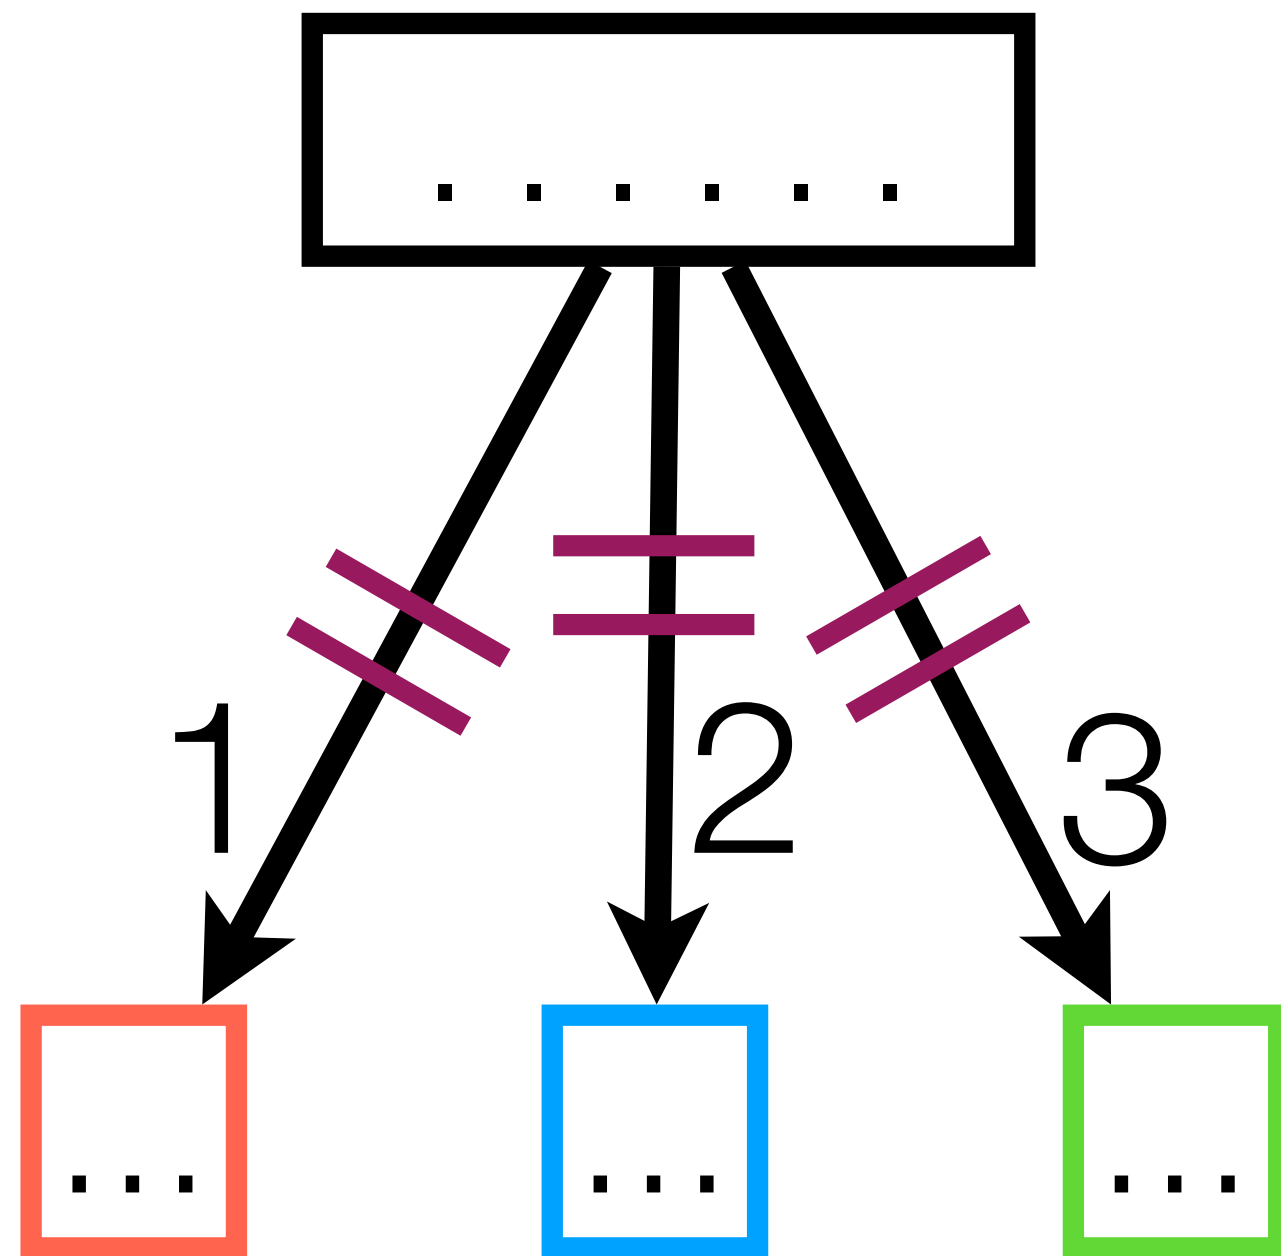
Compiling programs



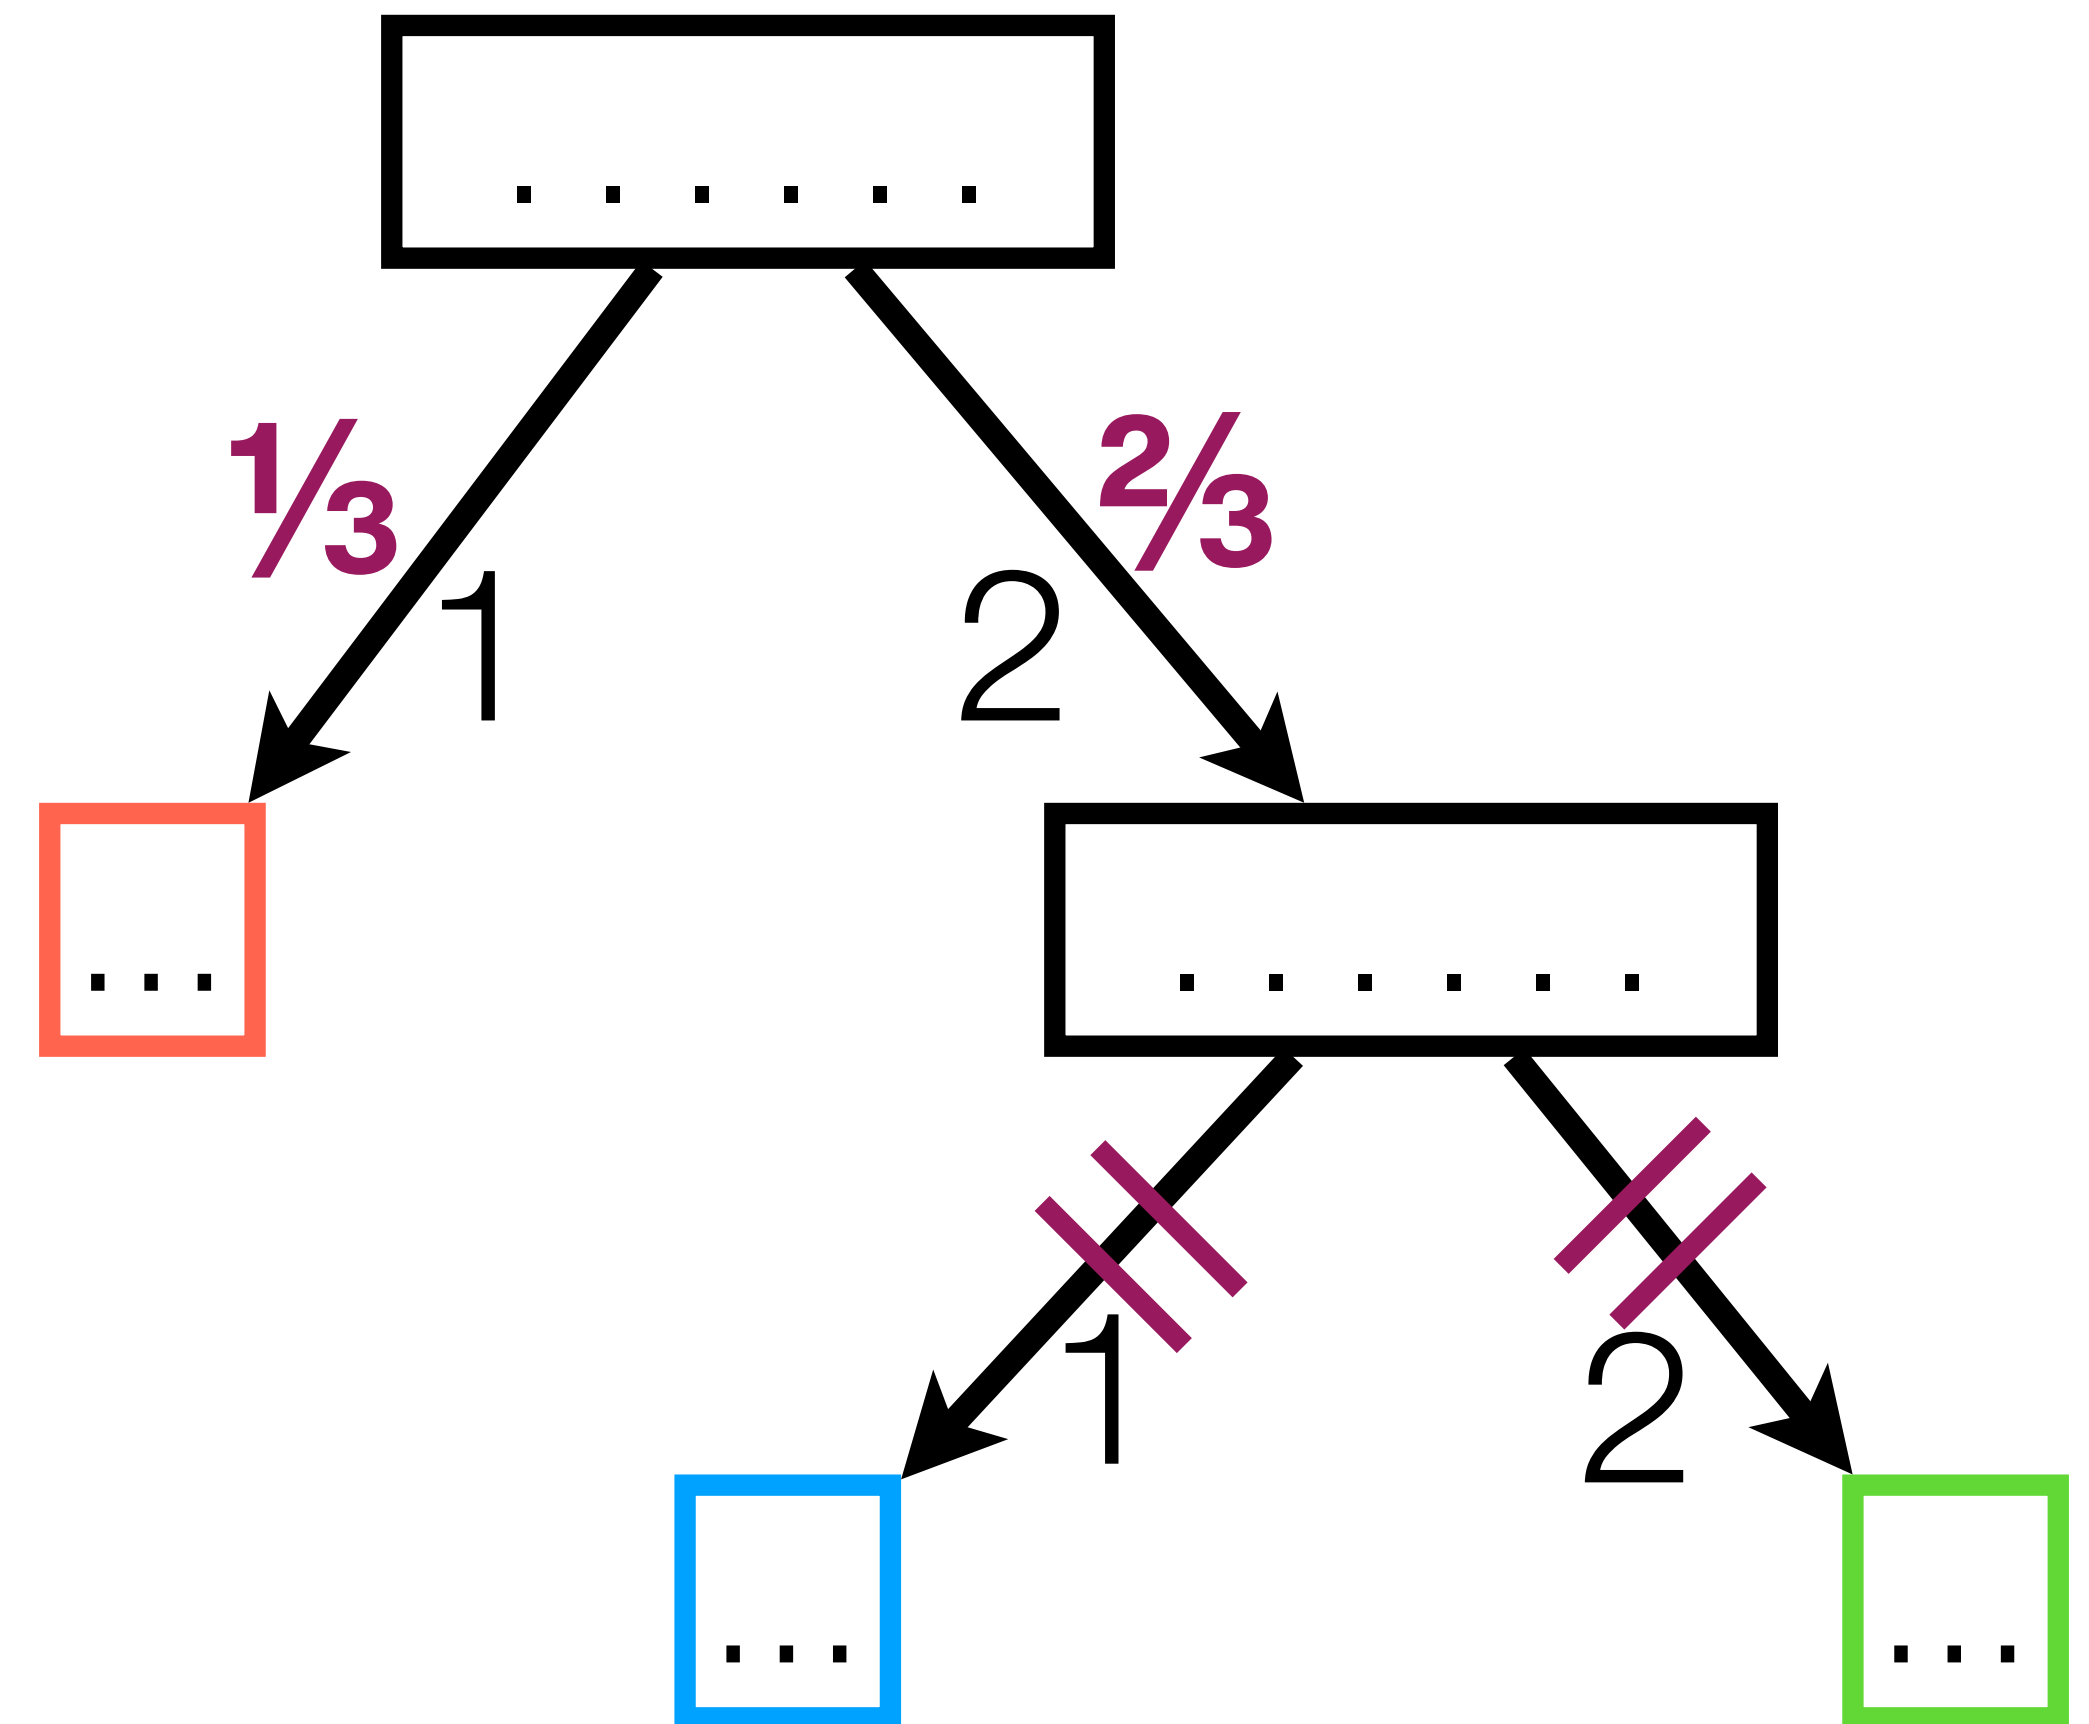
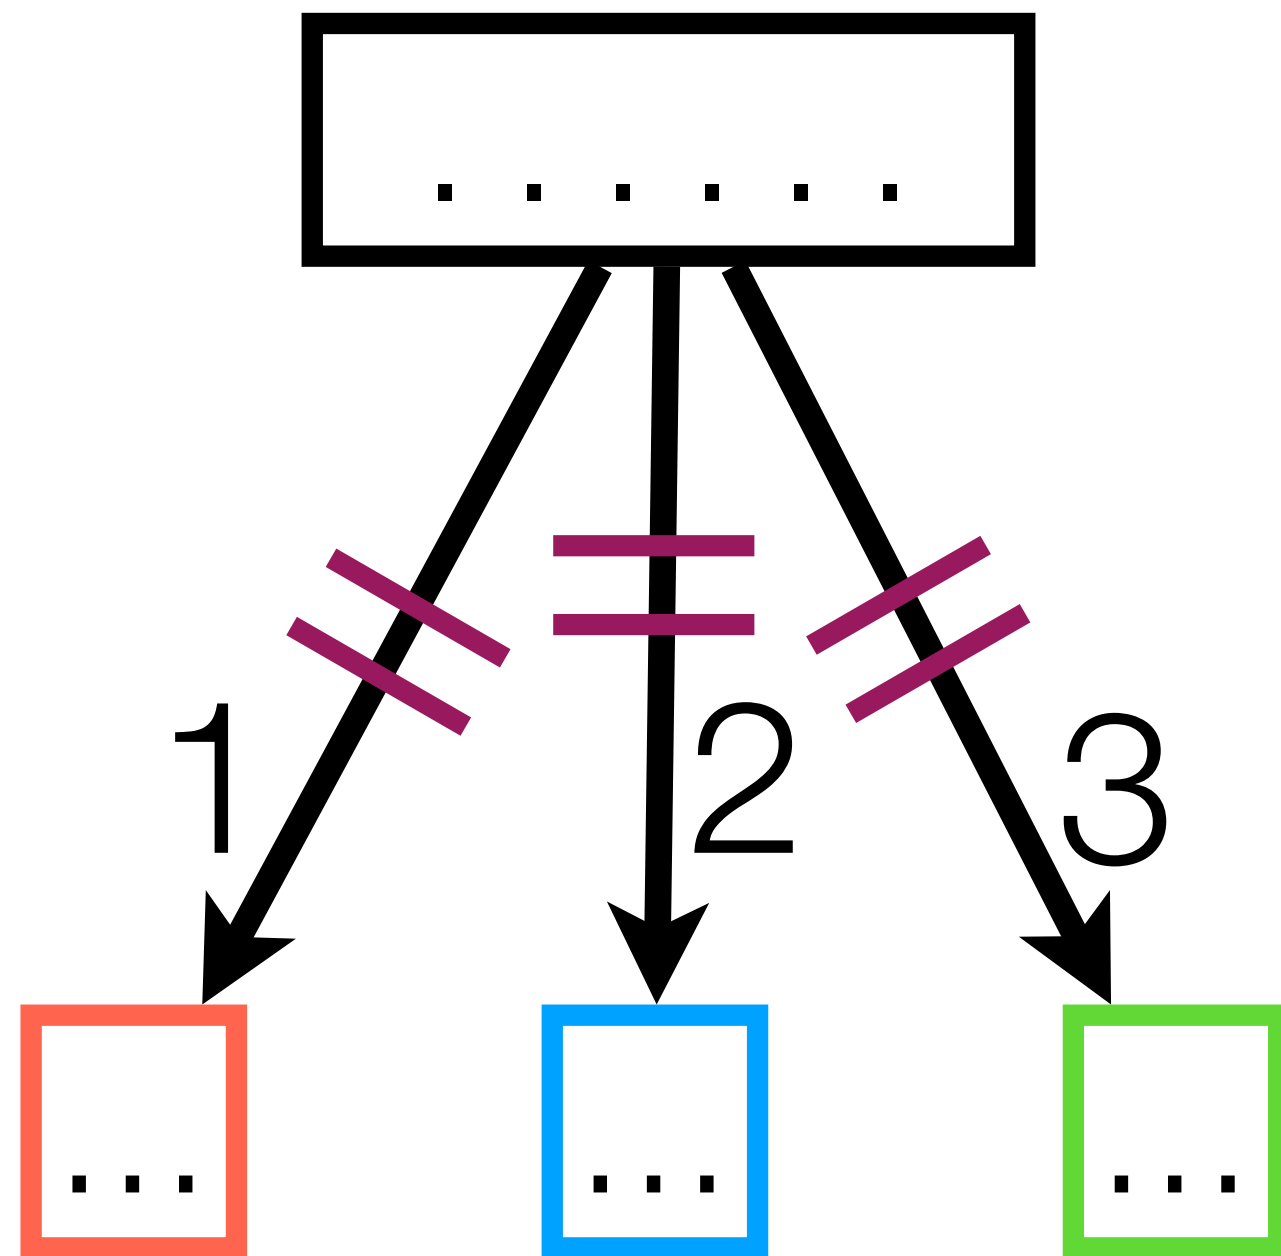
Compiling programs



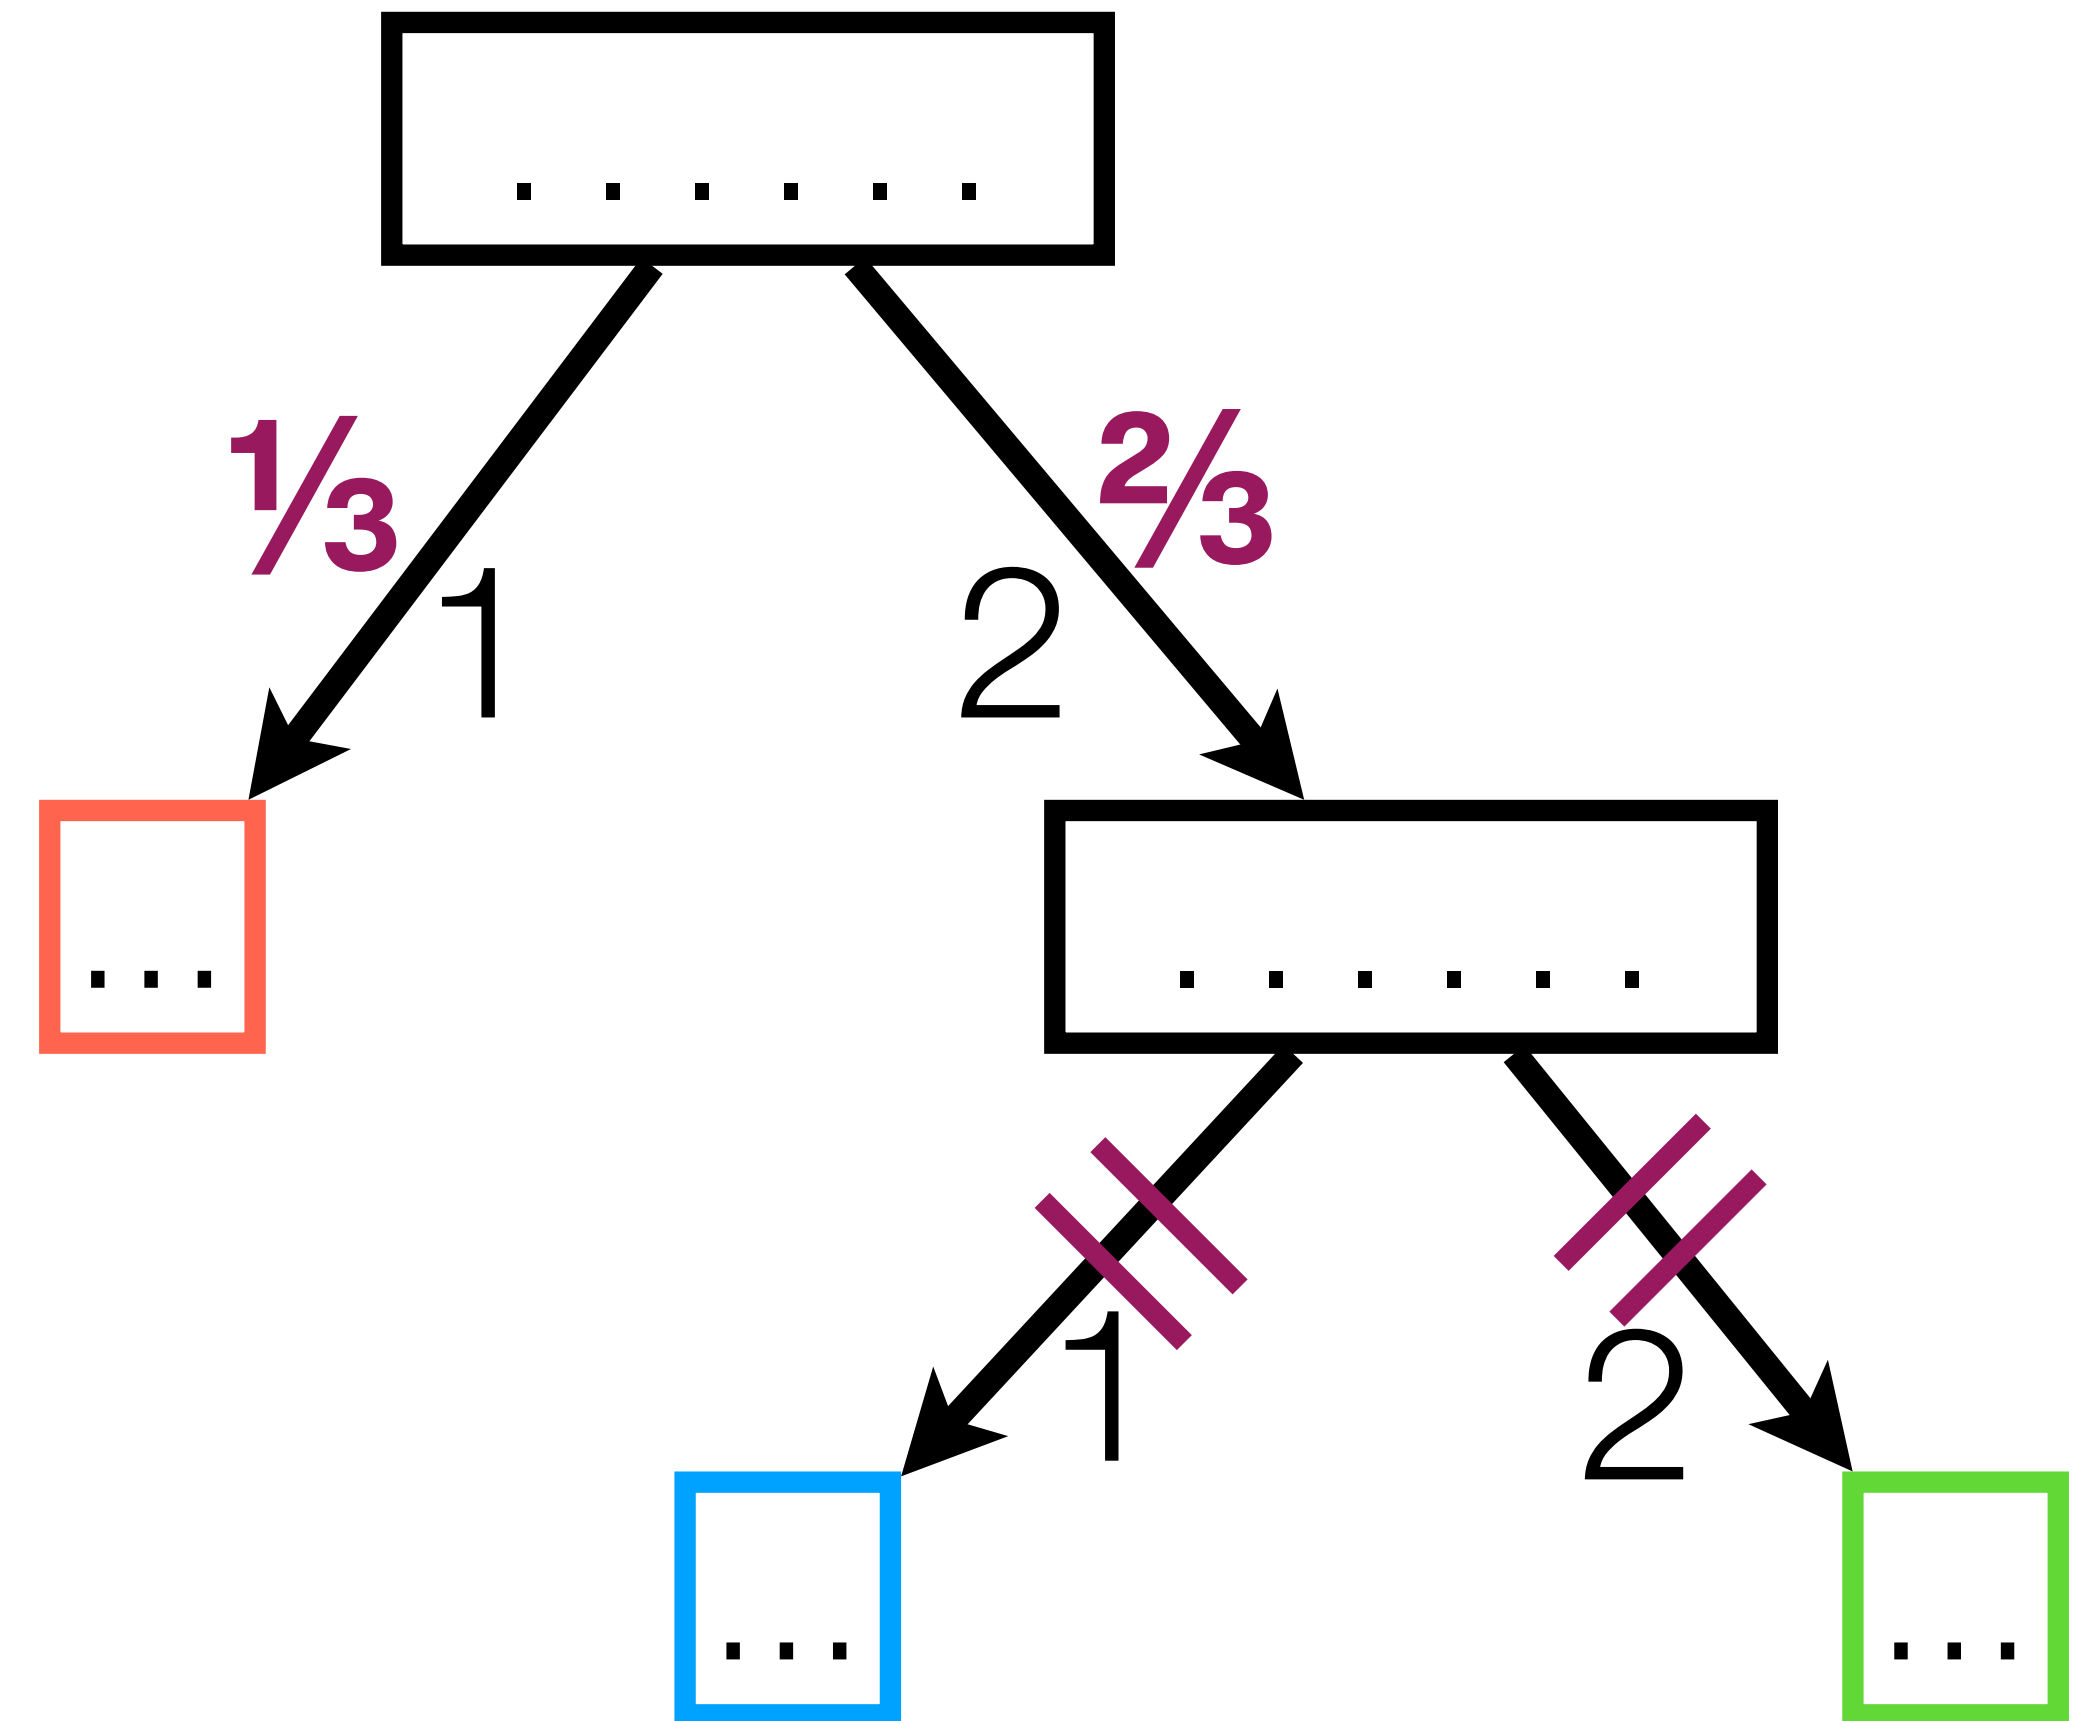
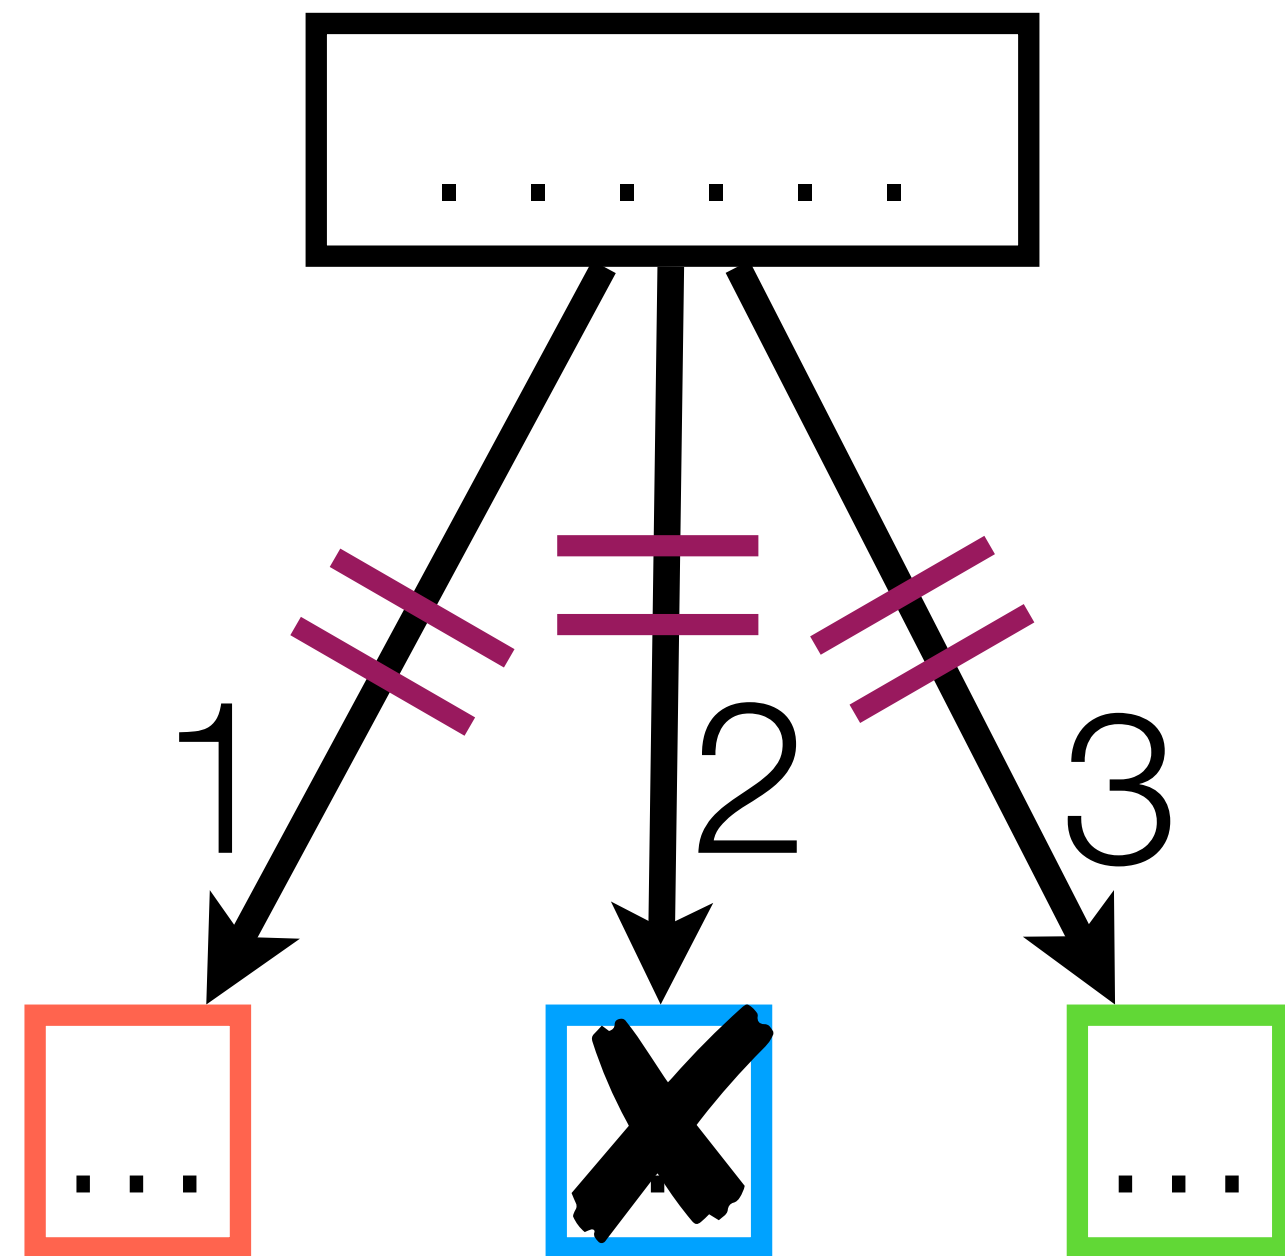
Compiling programs



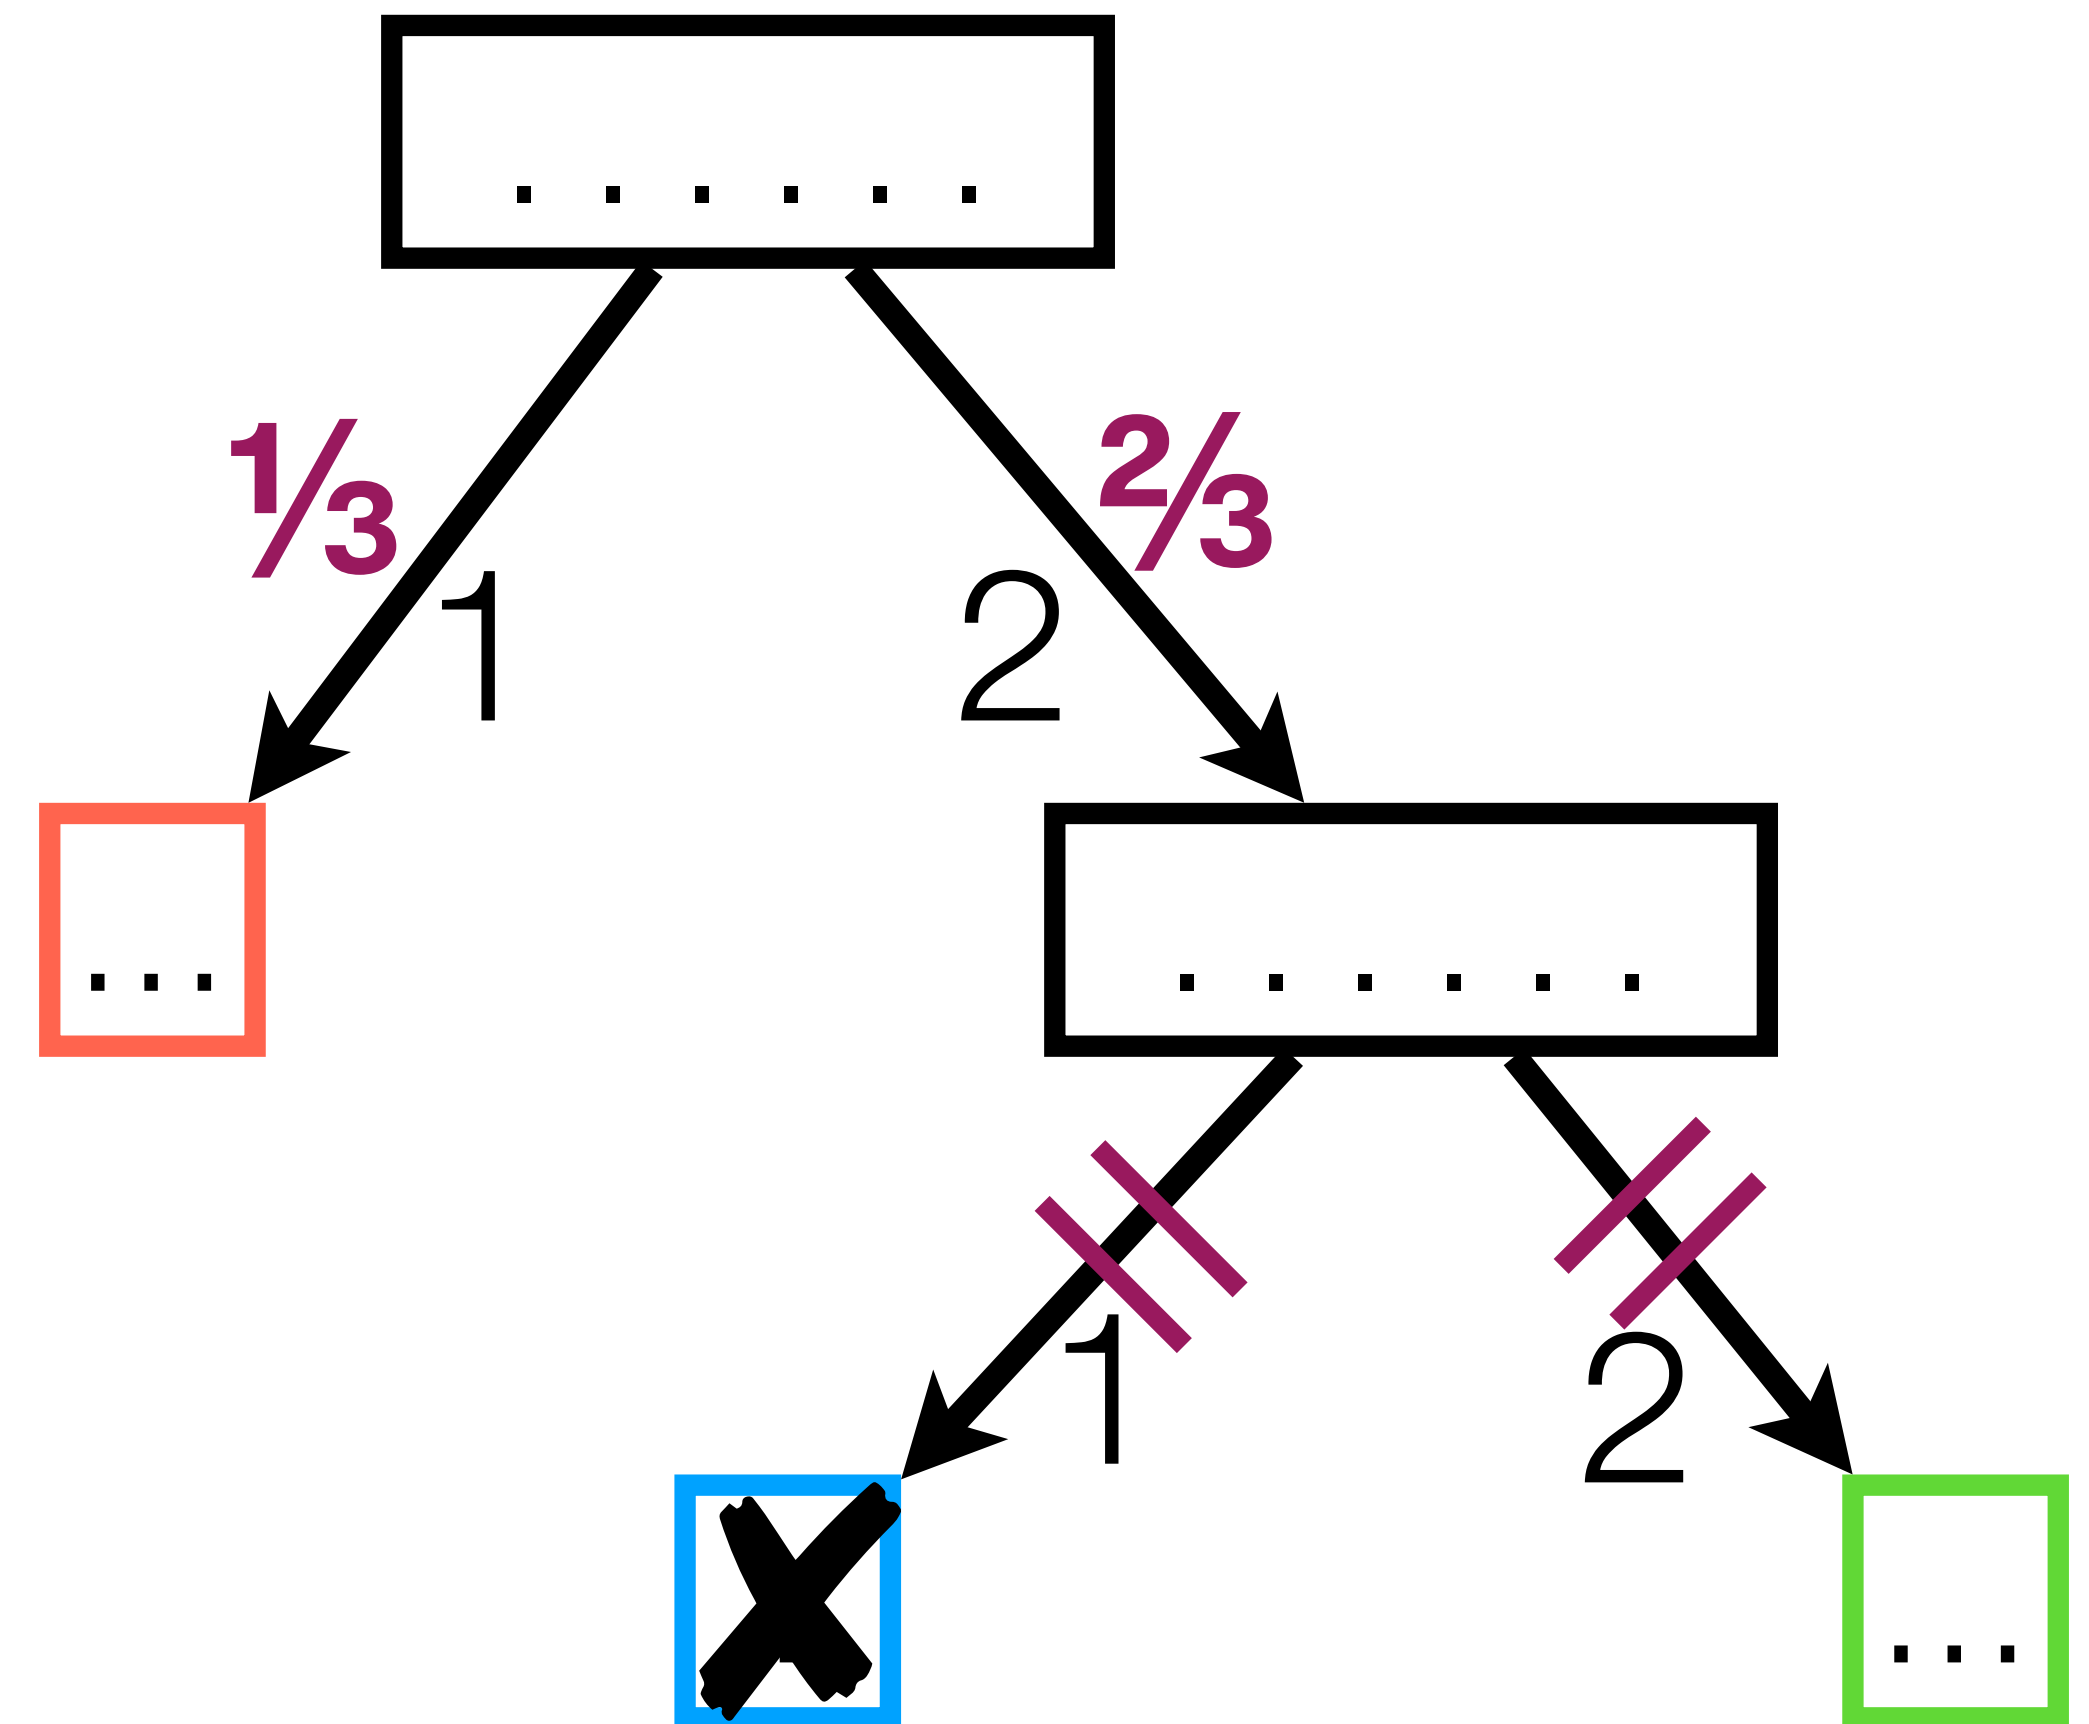
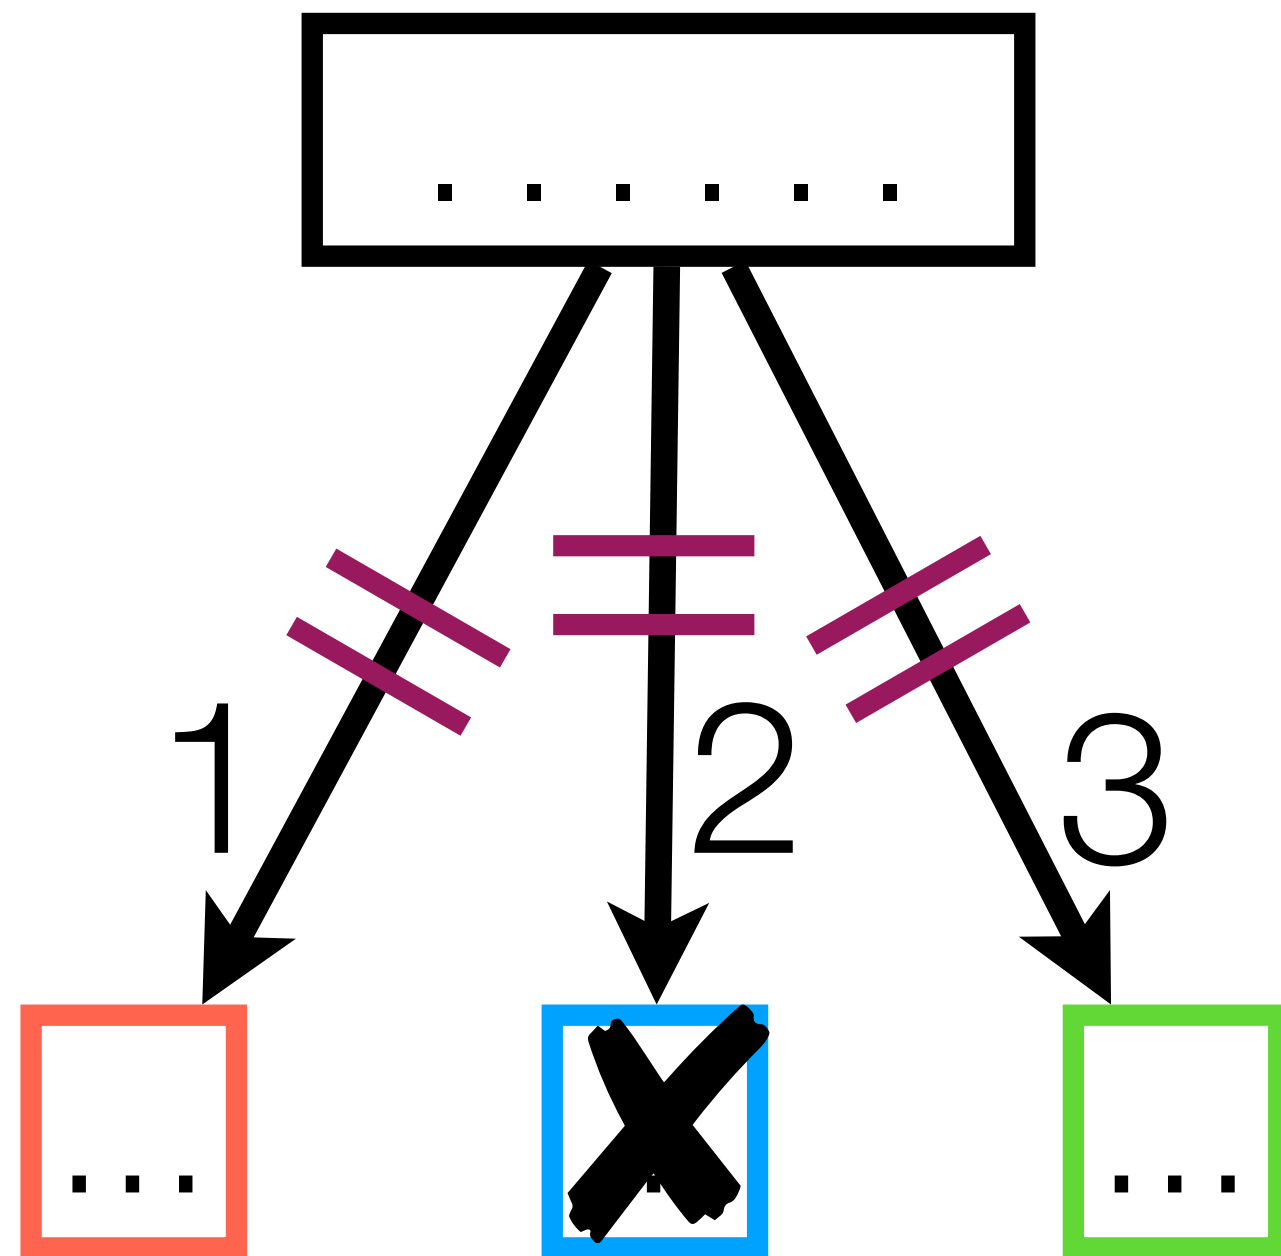
Compiling programs



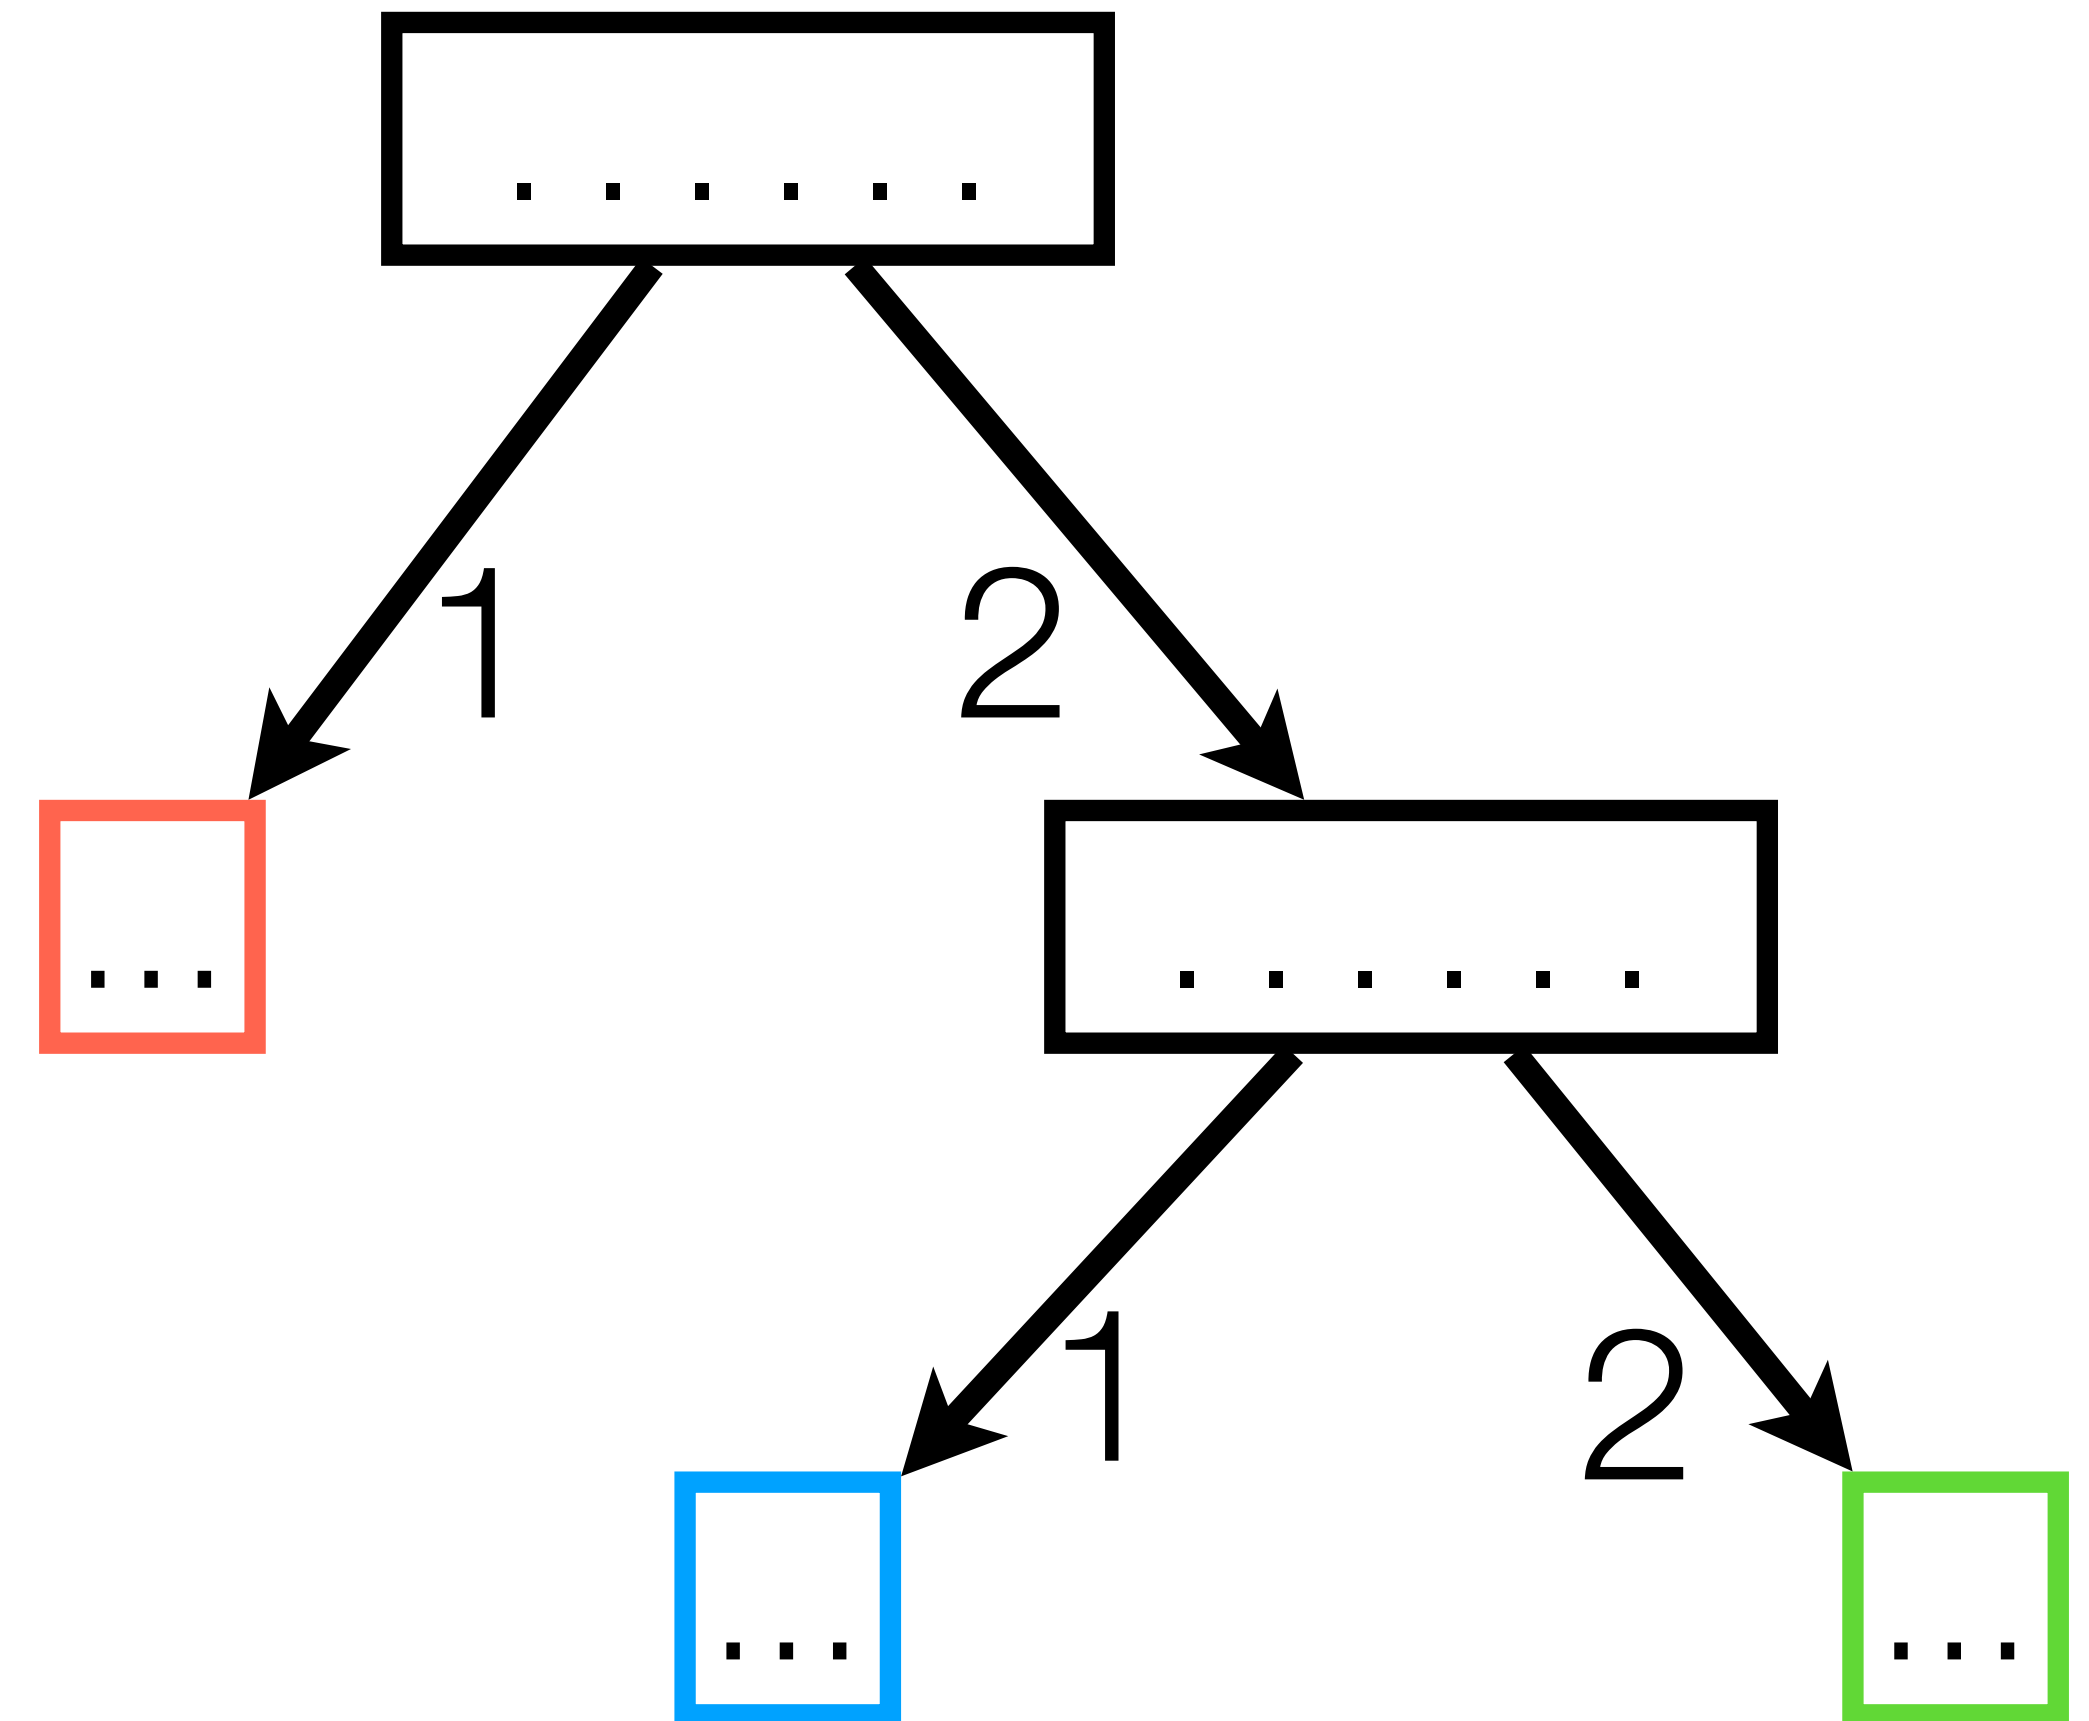
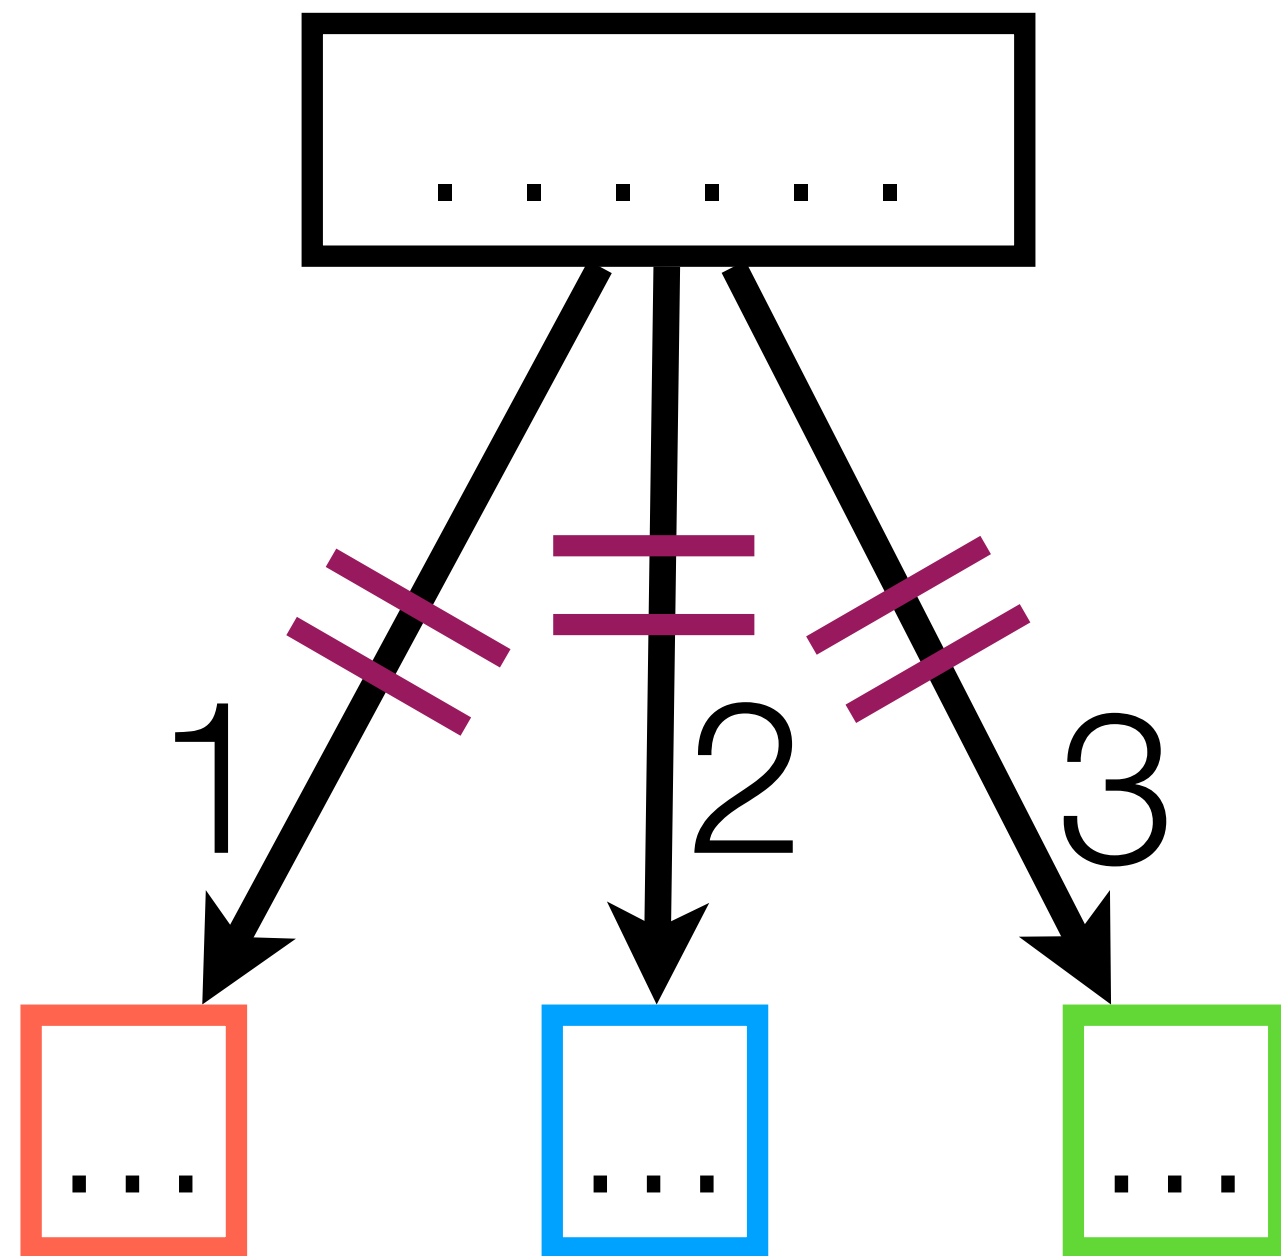
Compiling programs



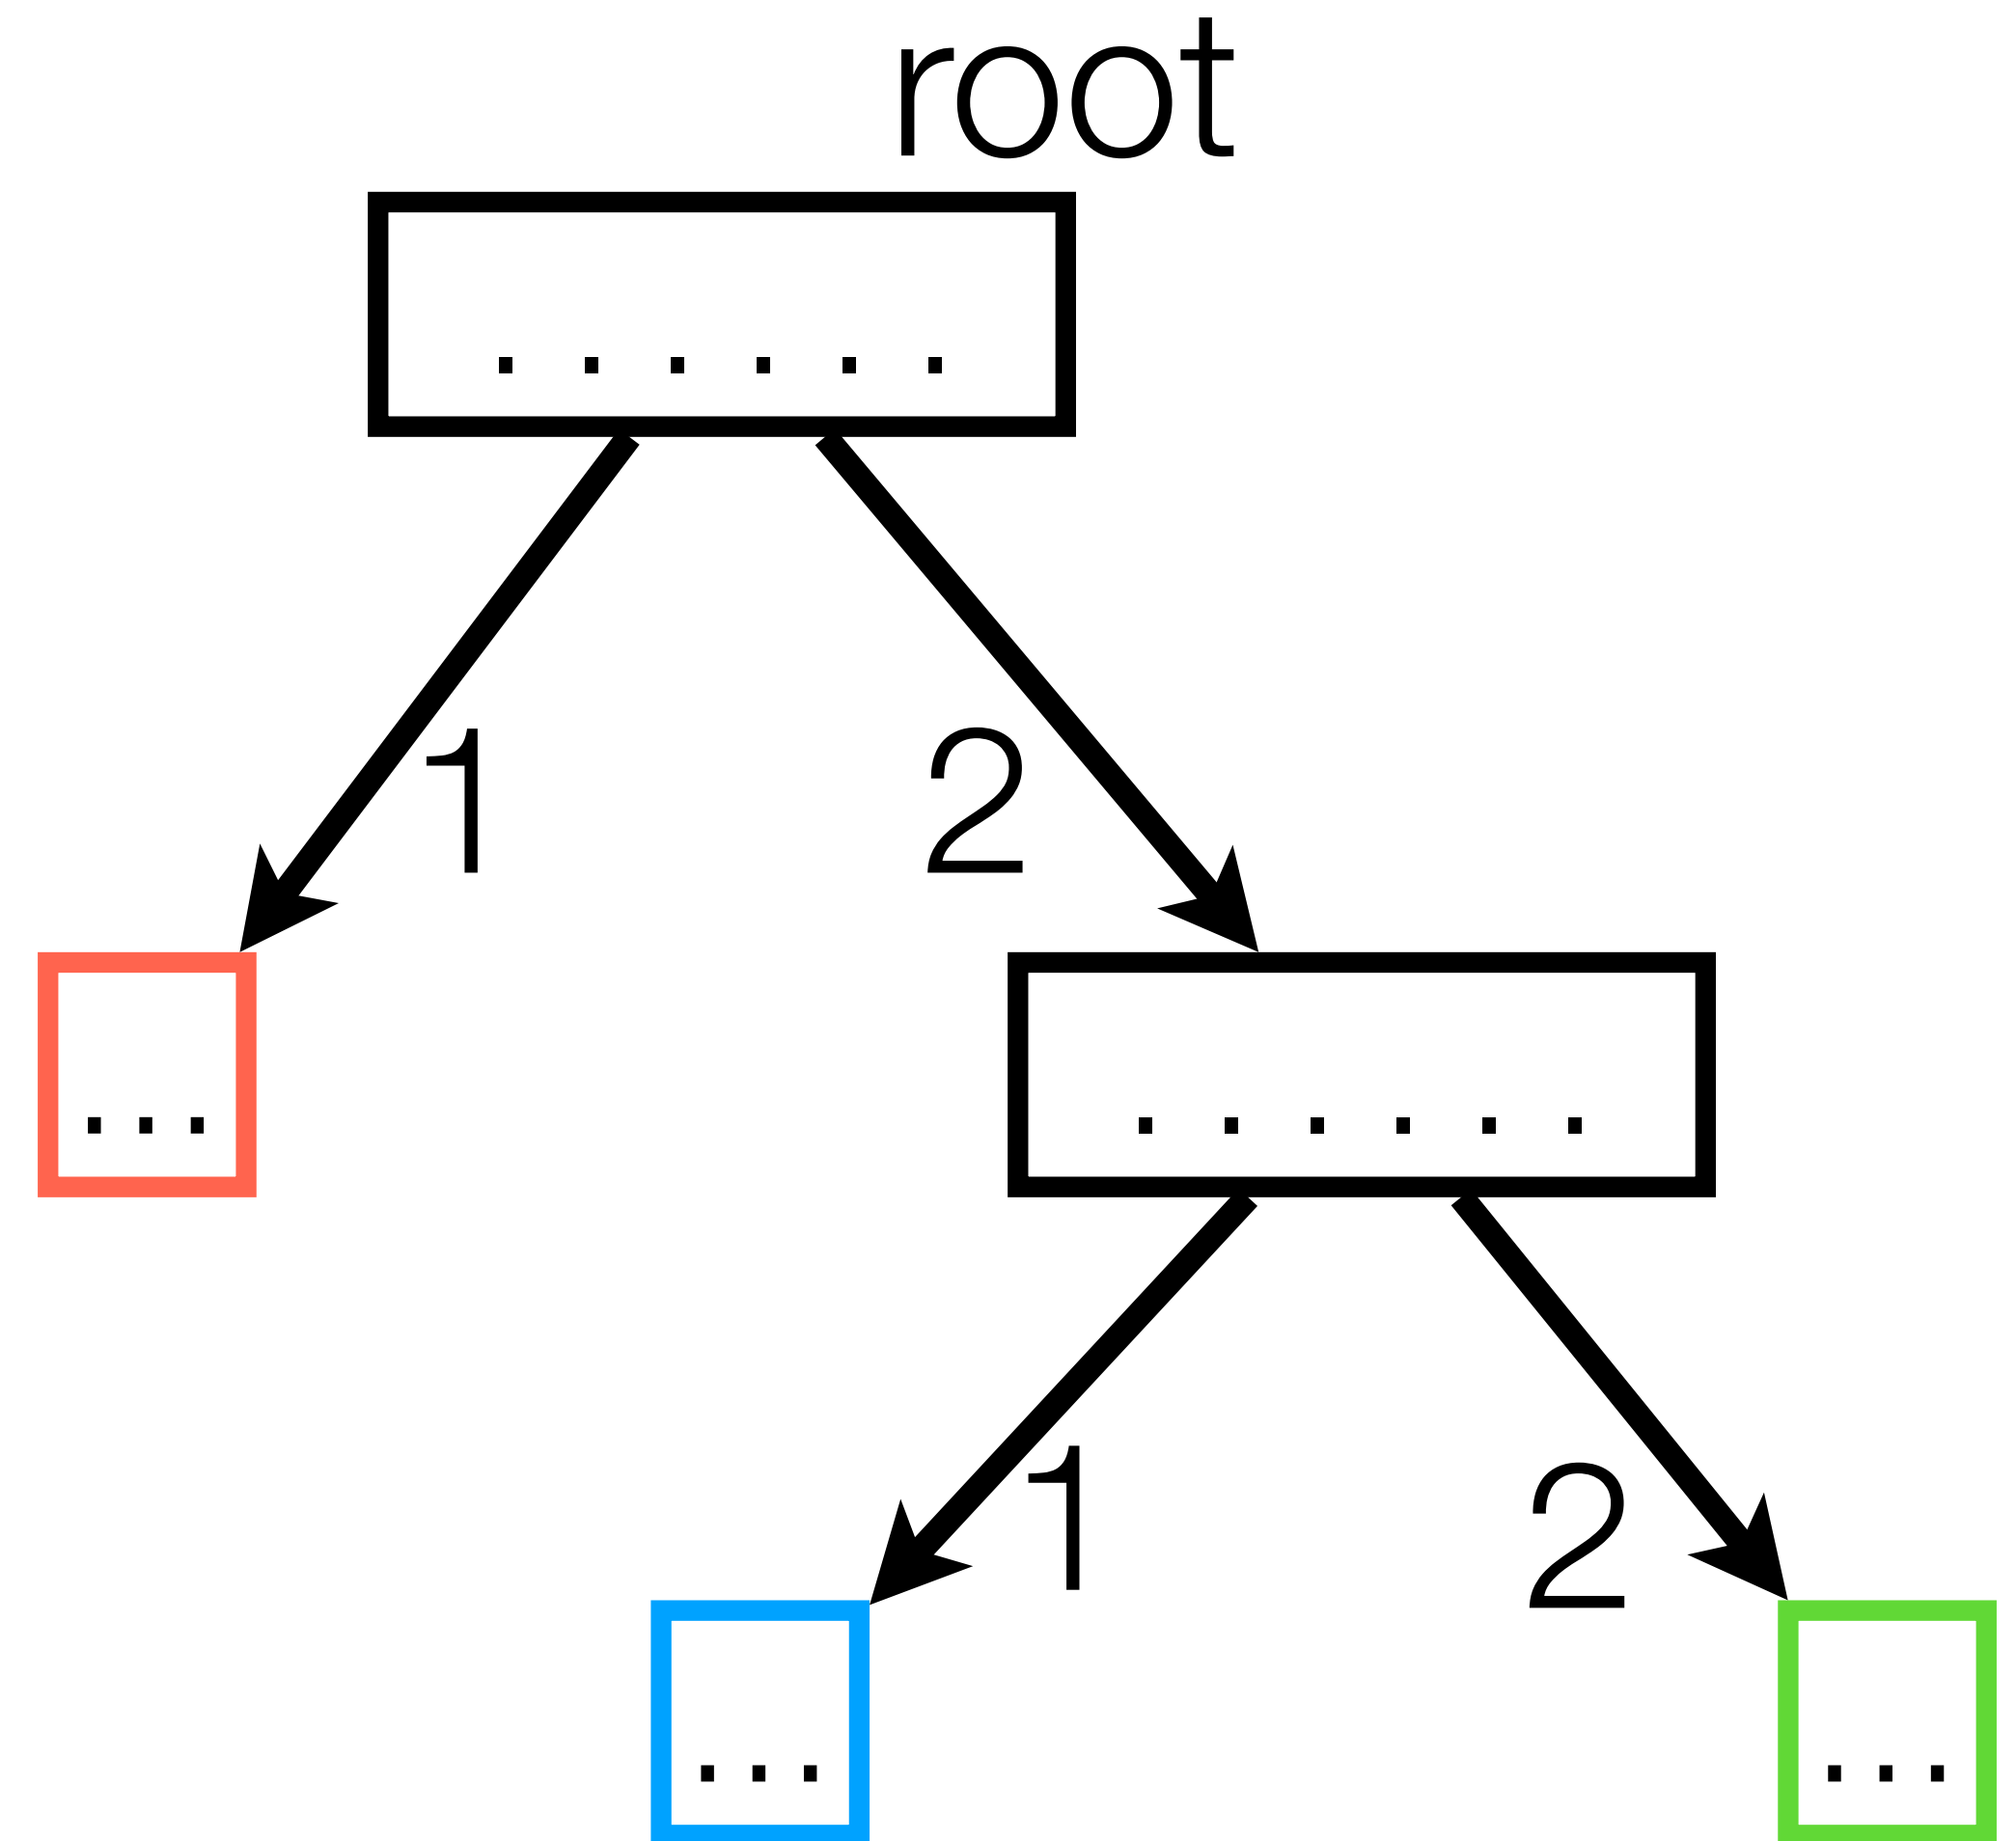
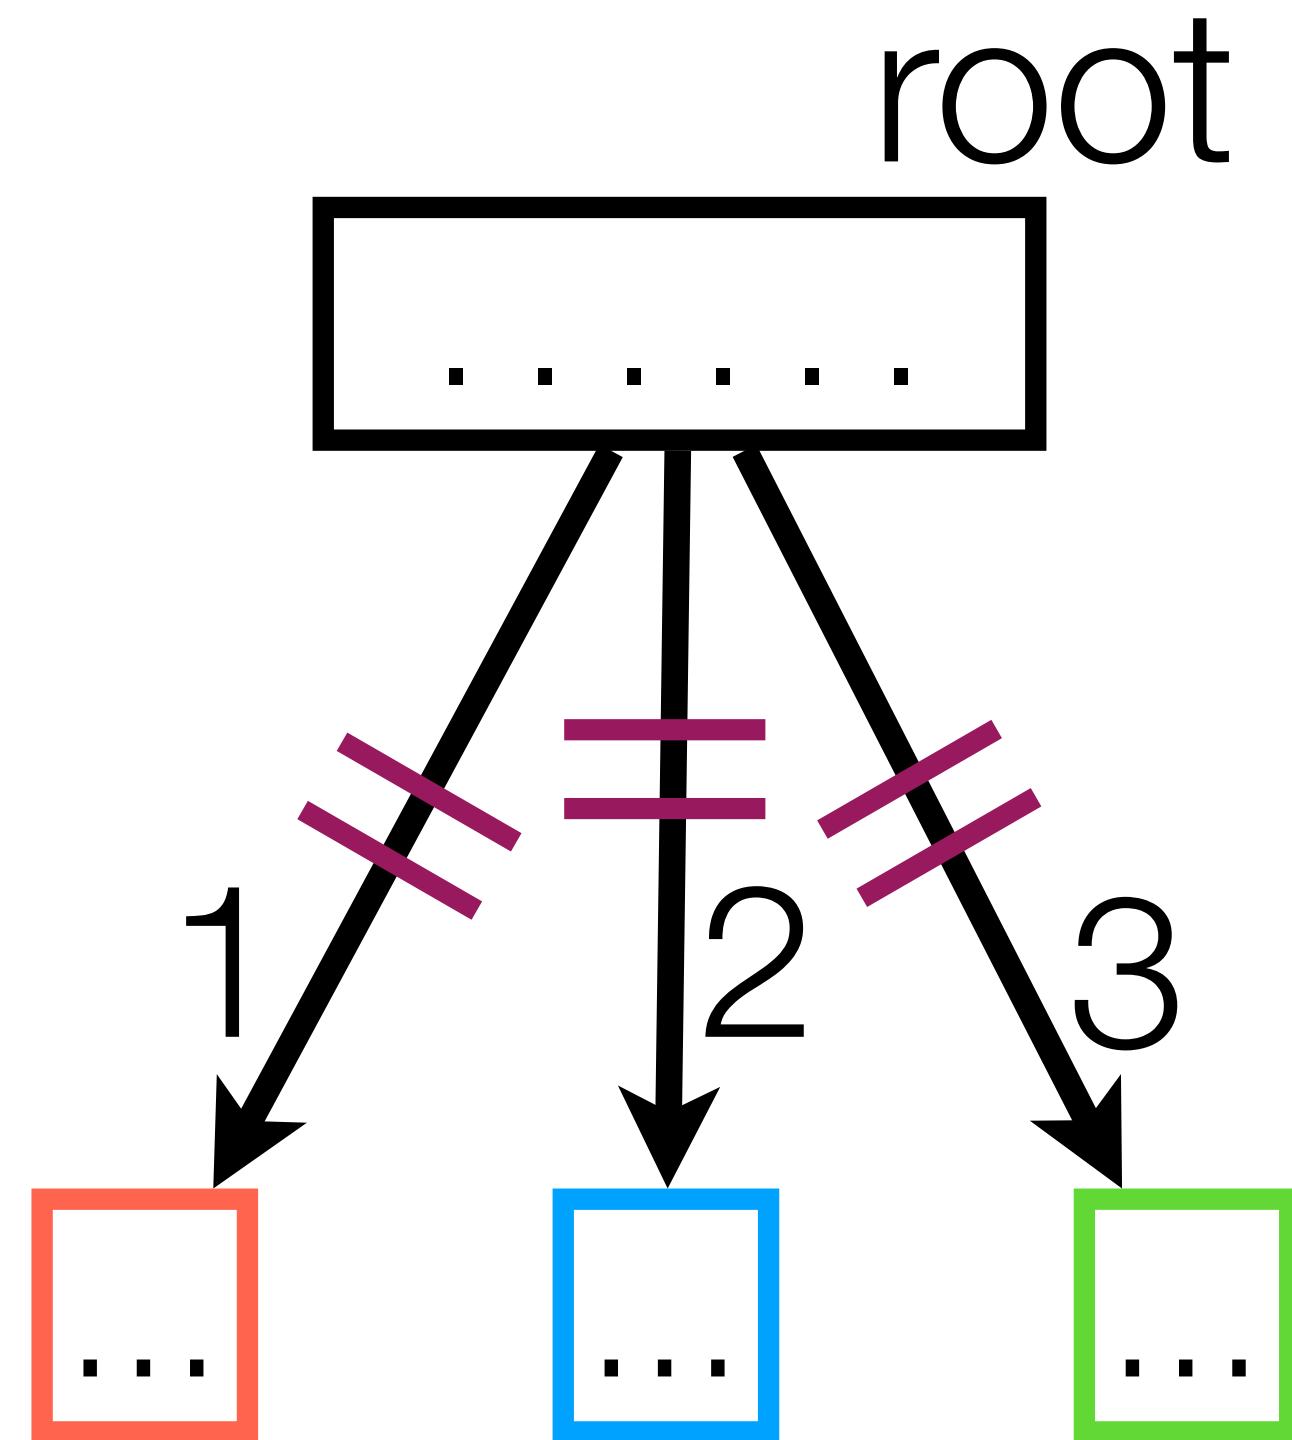
Compiling programs



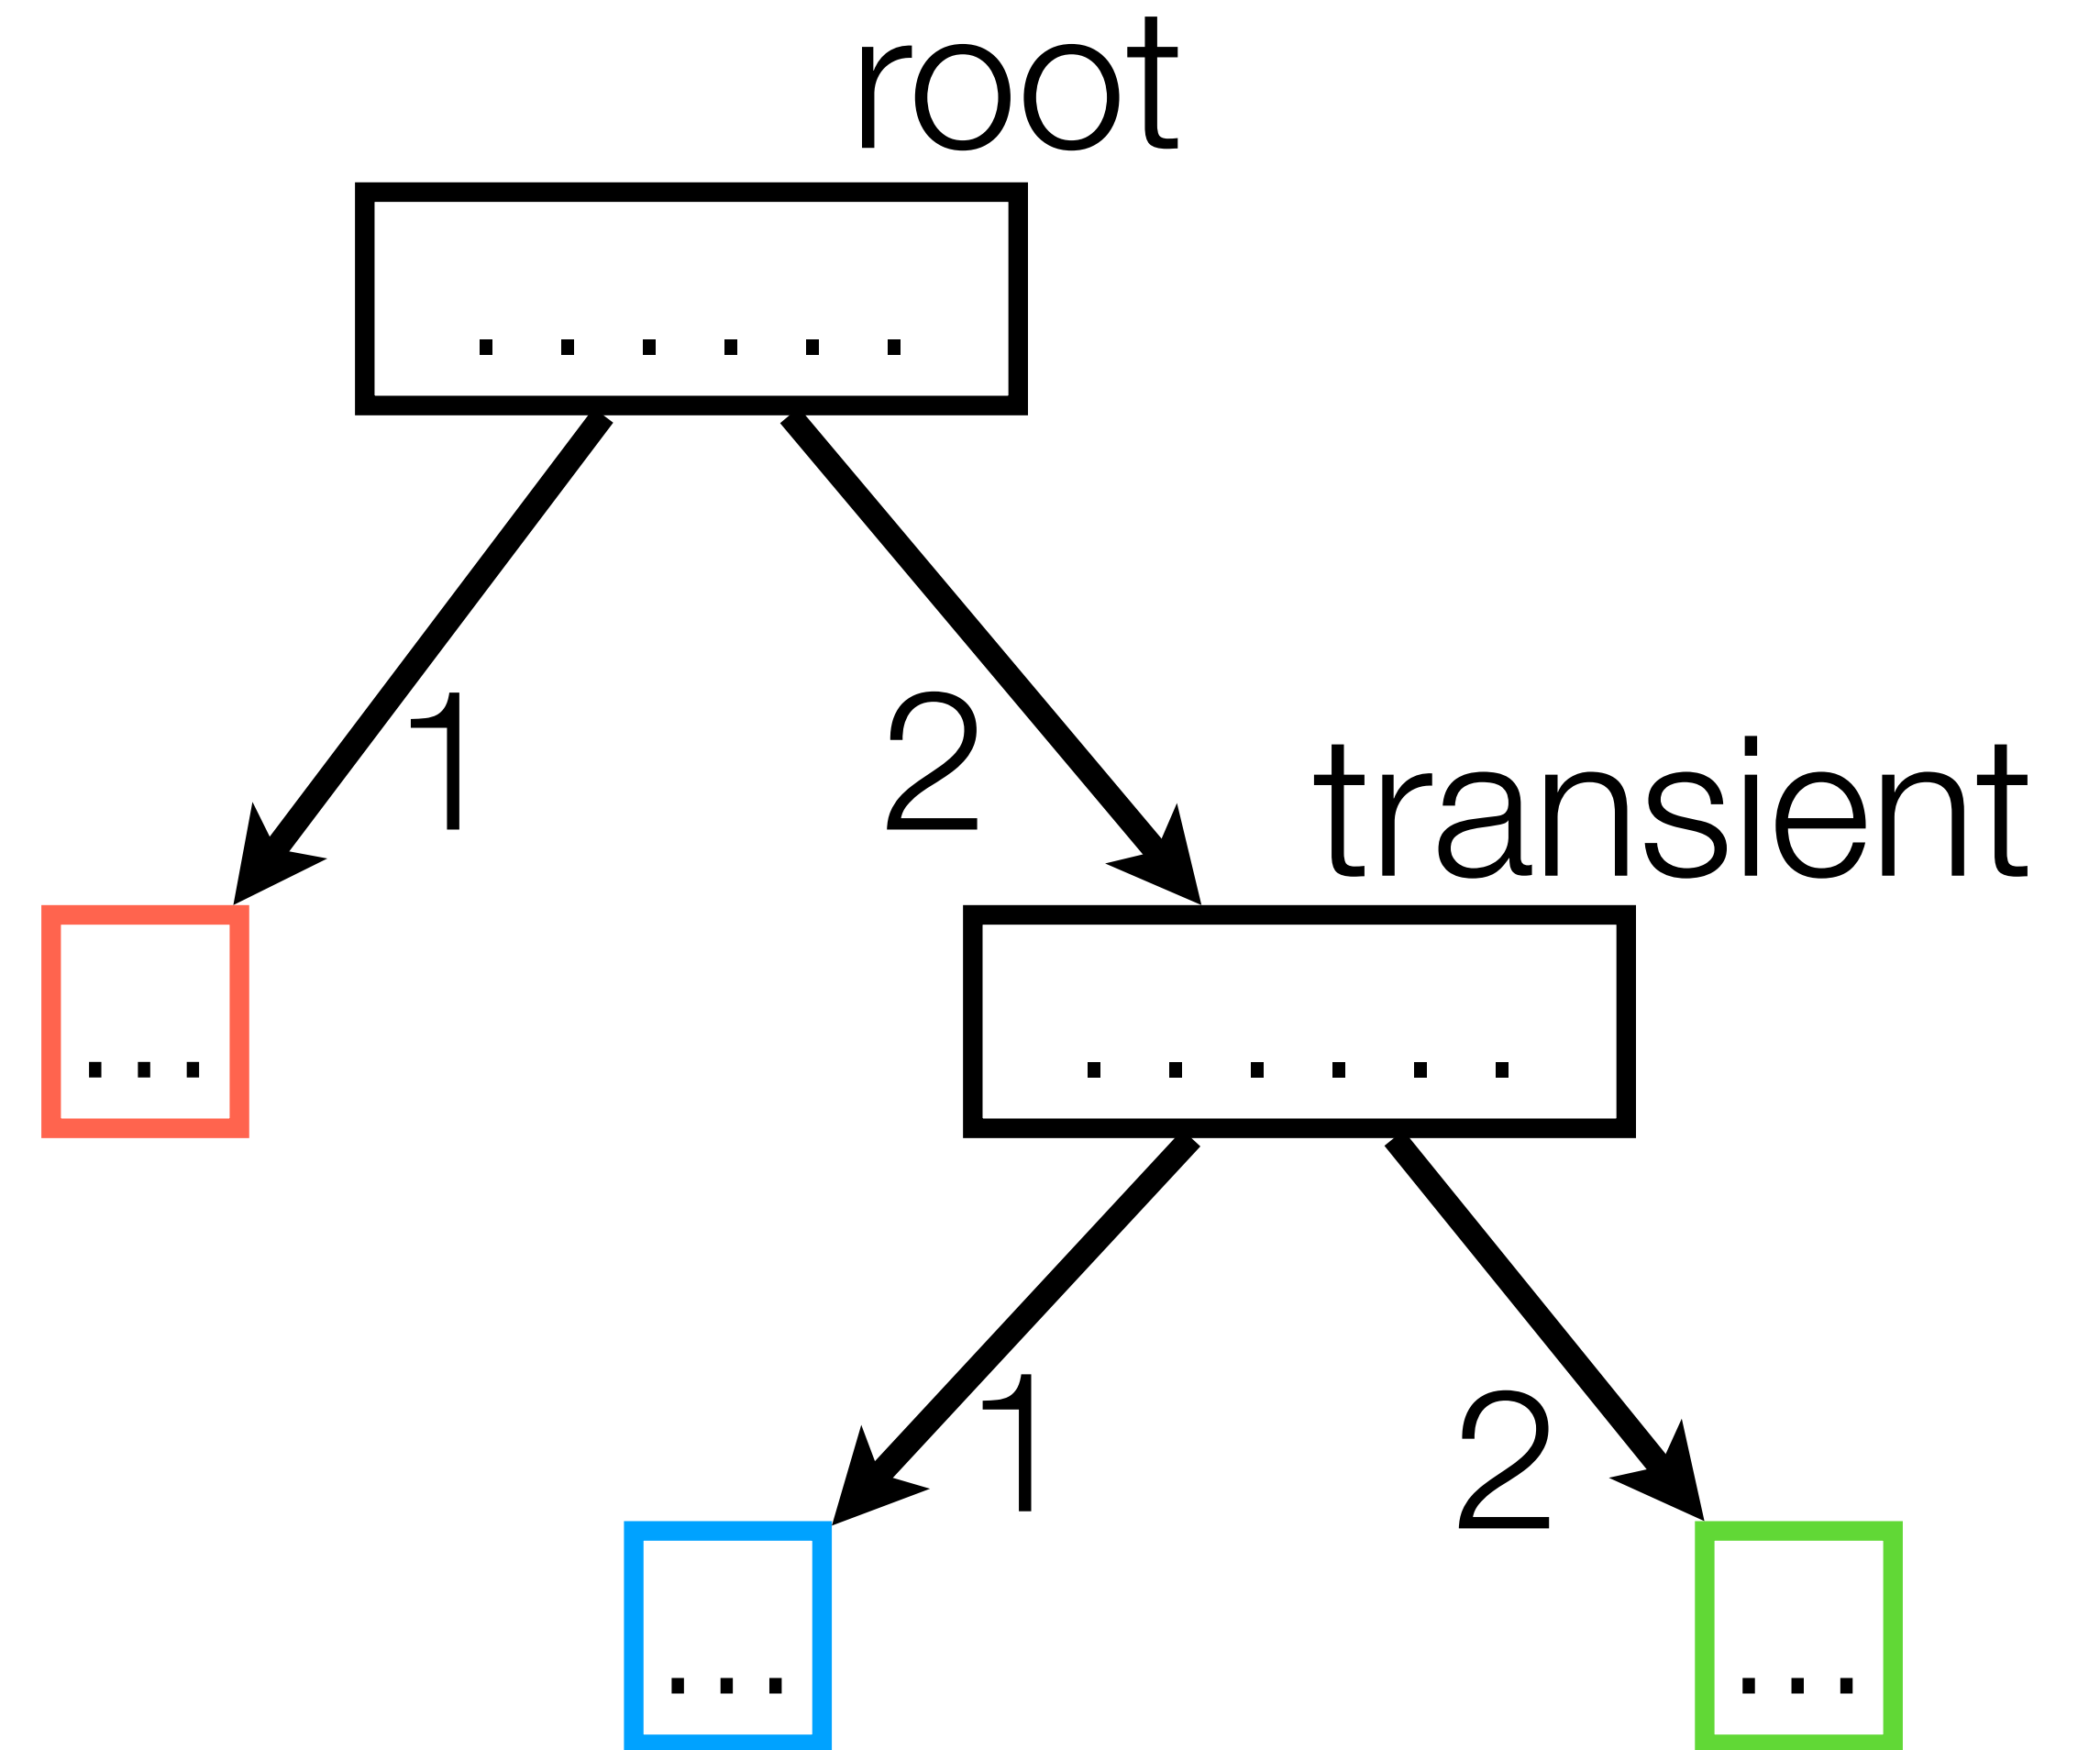
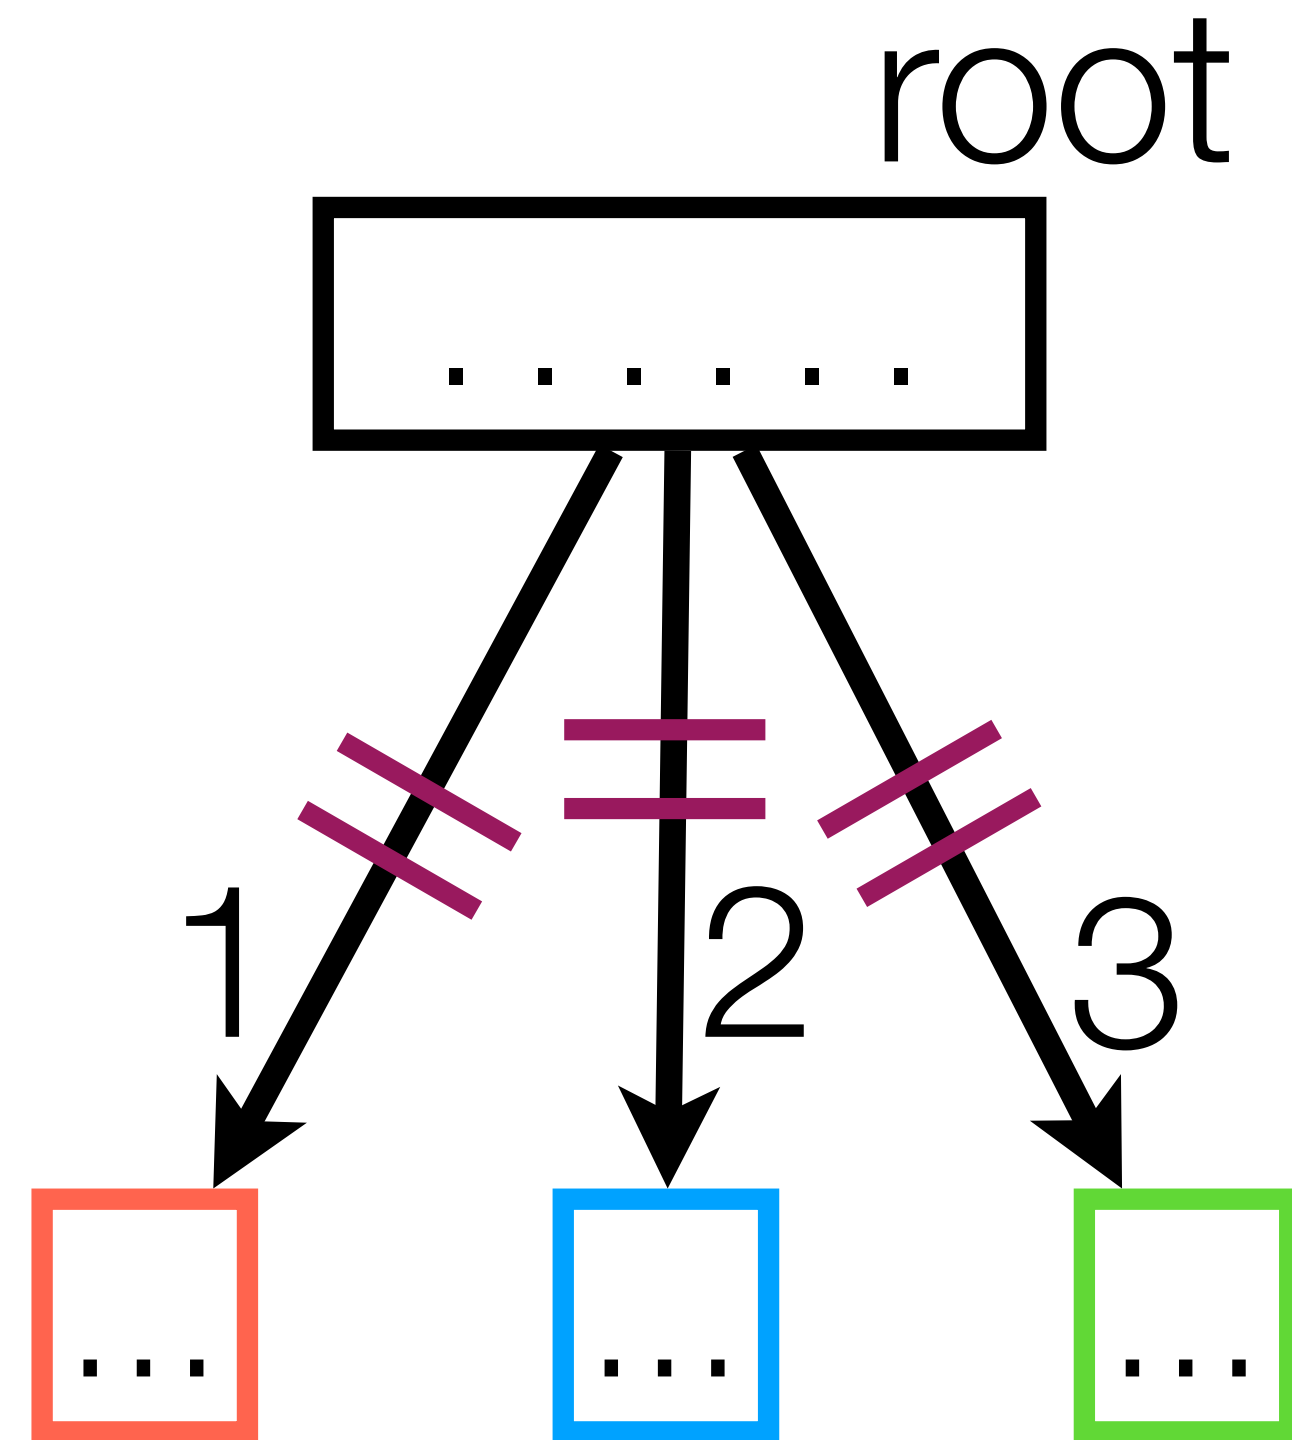
Compiling programs



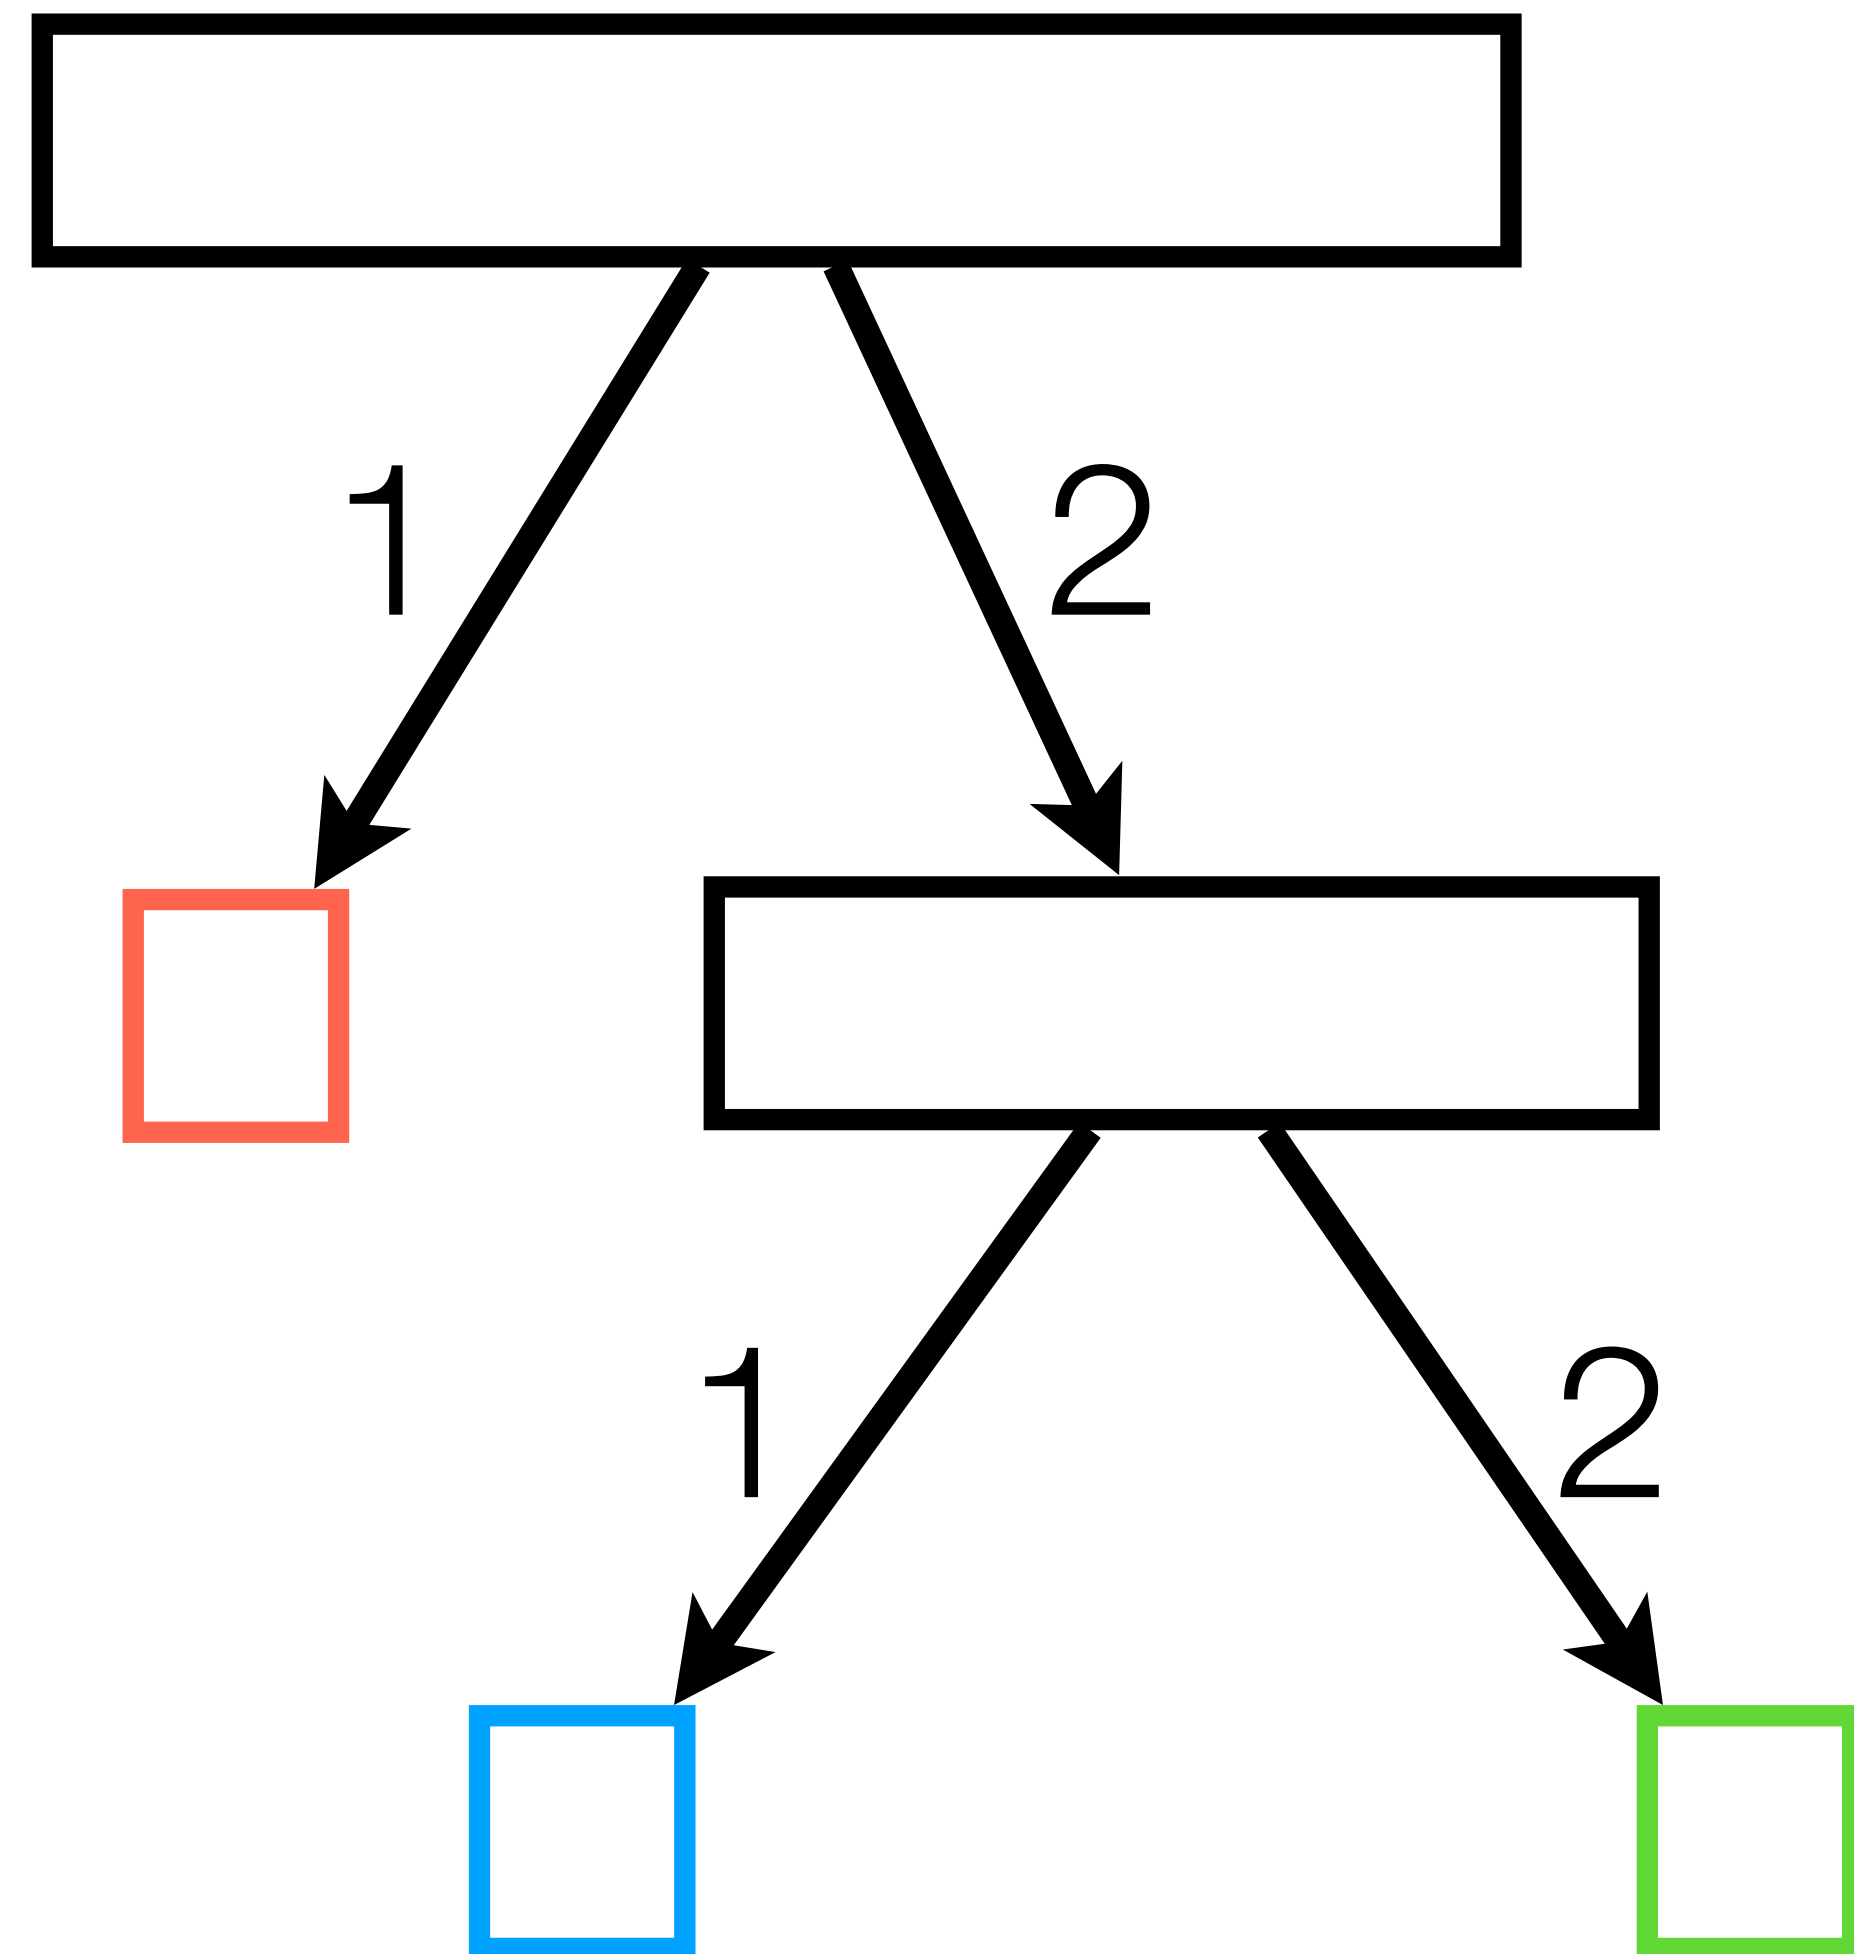
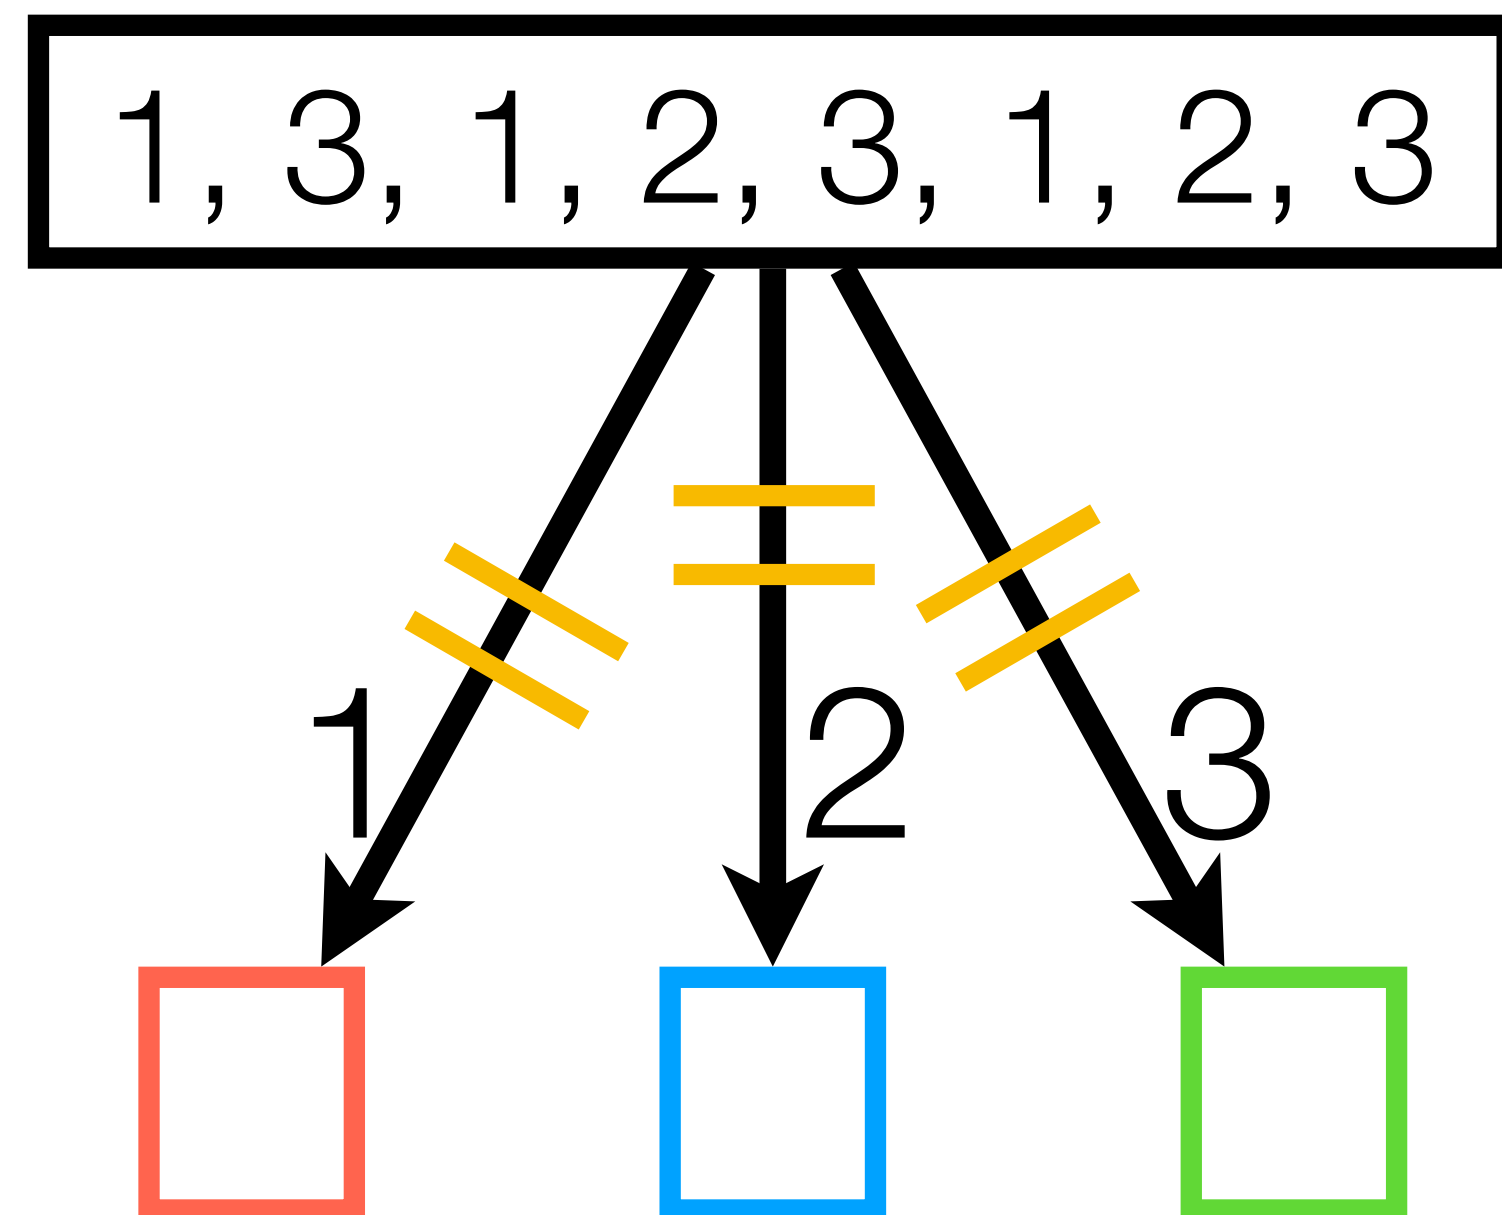
Compiling programs



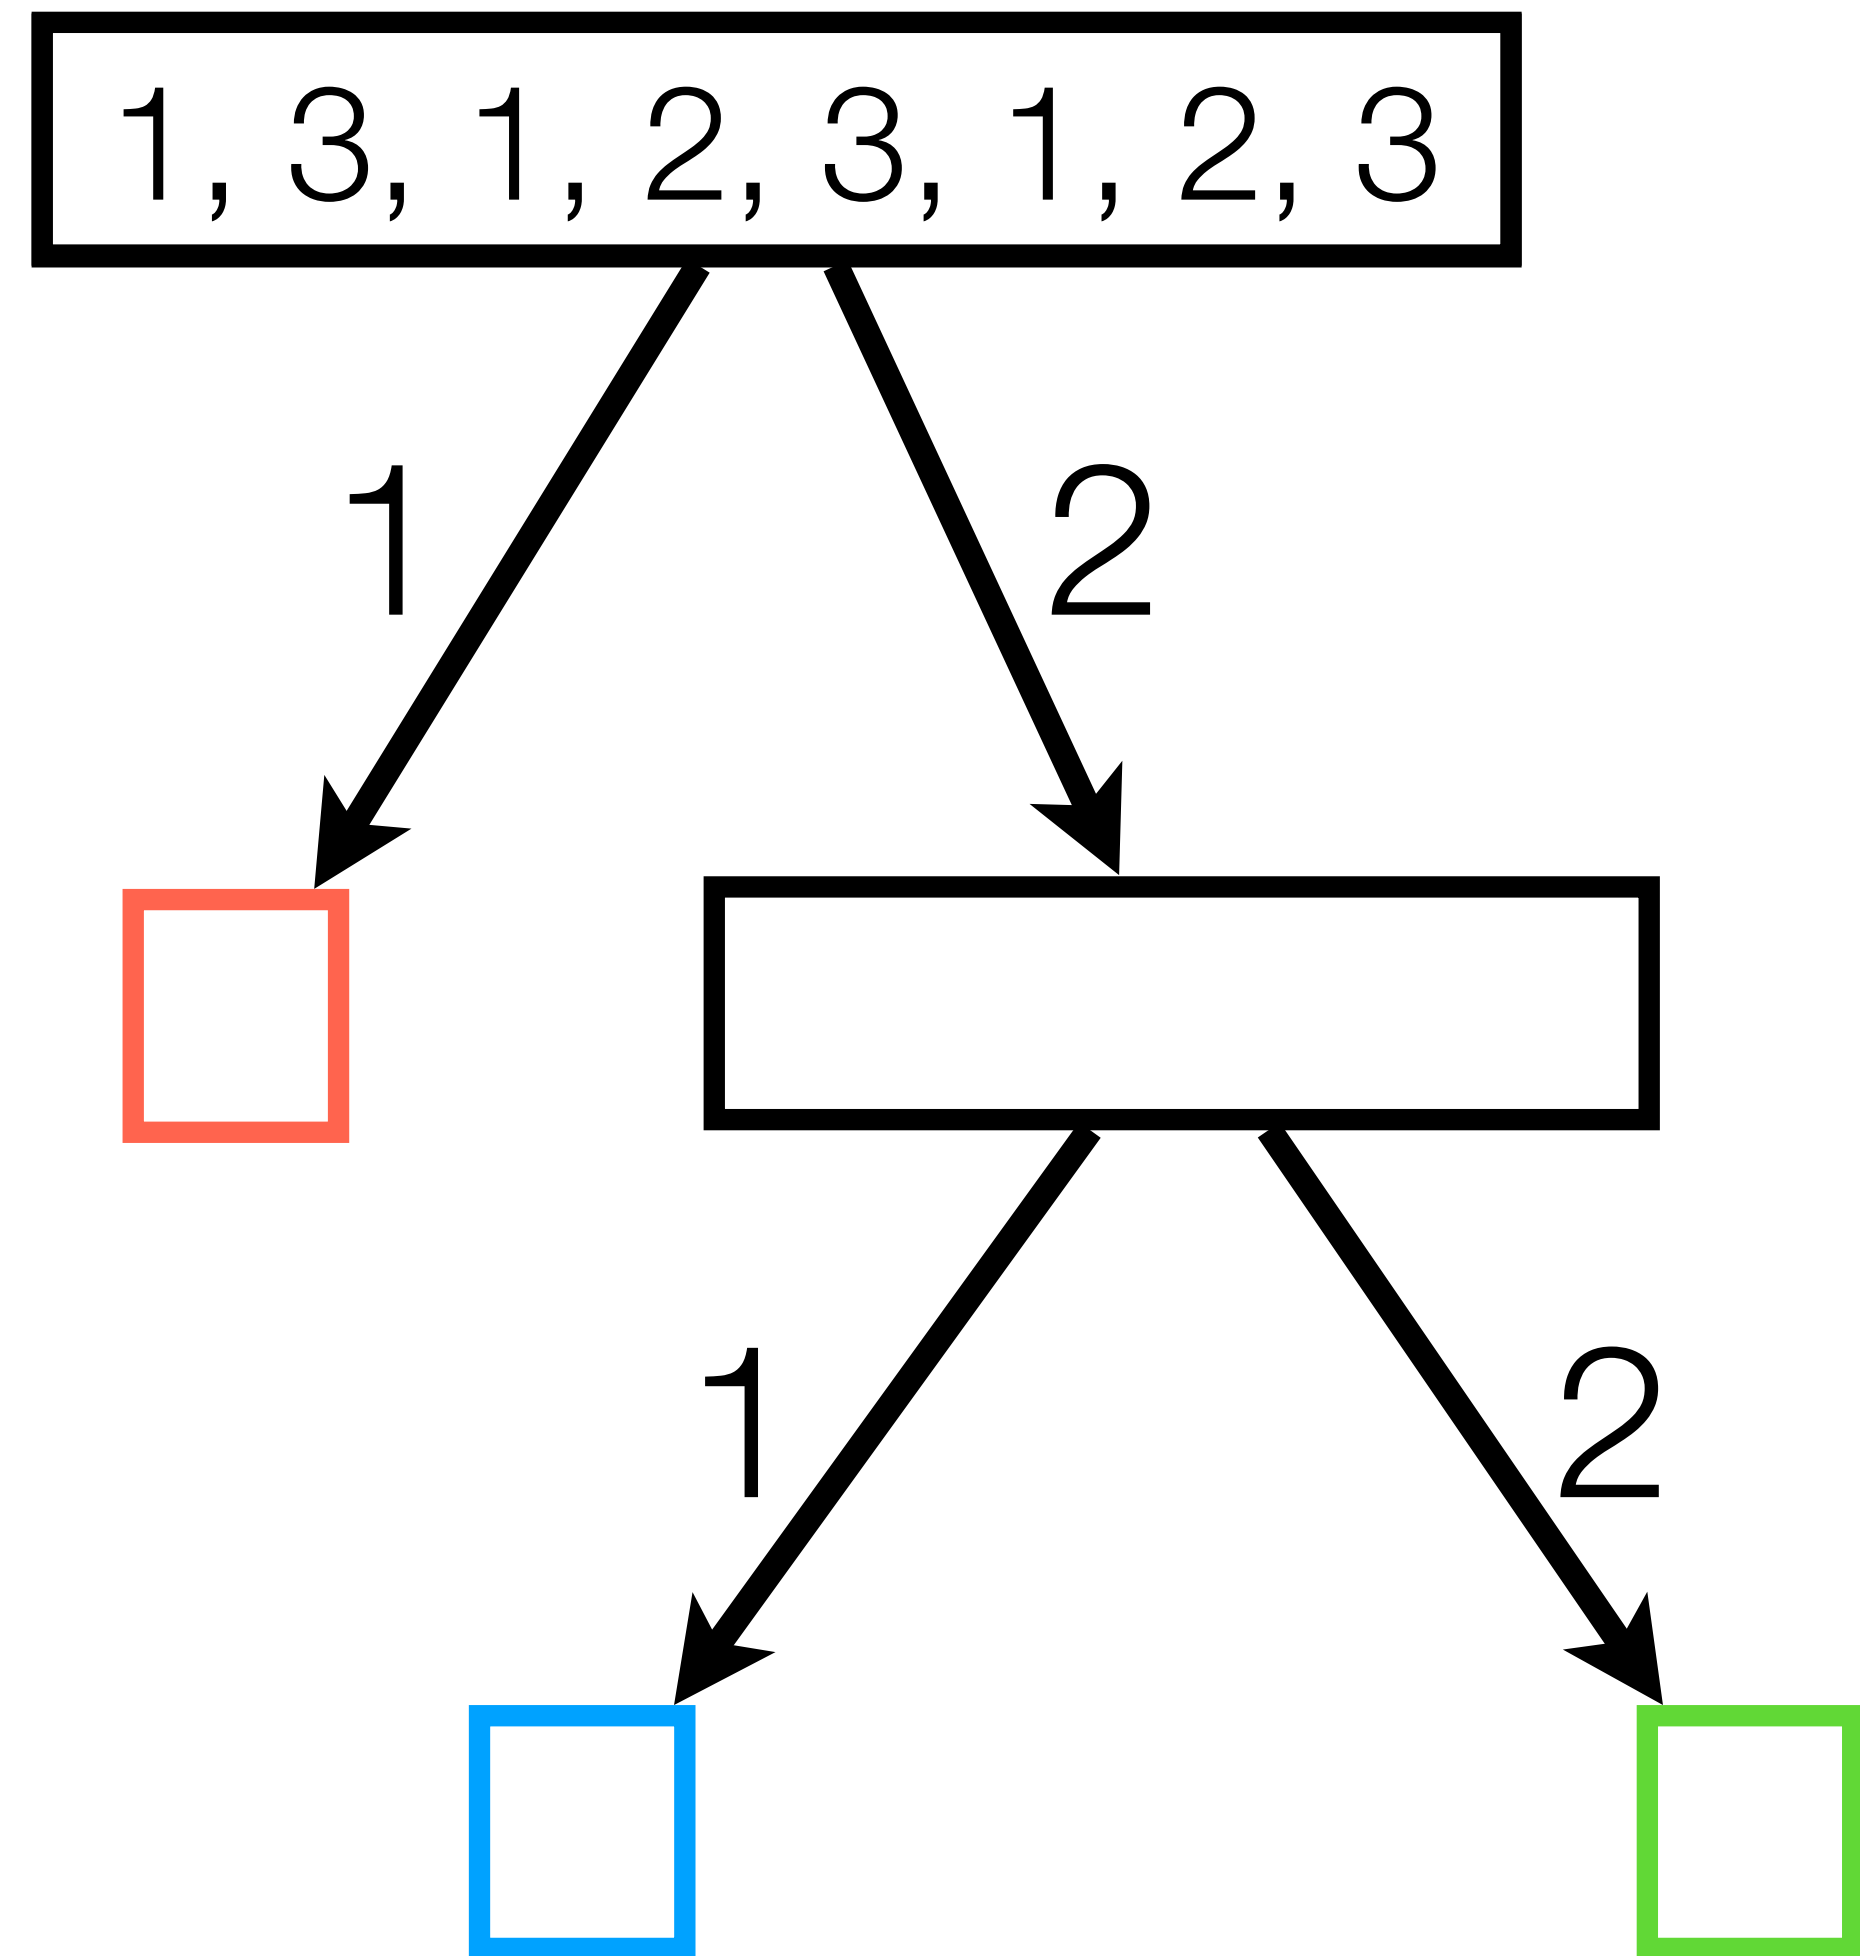
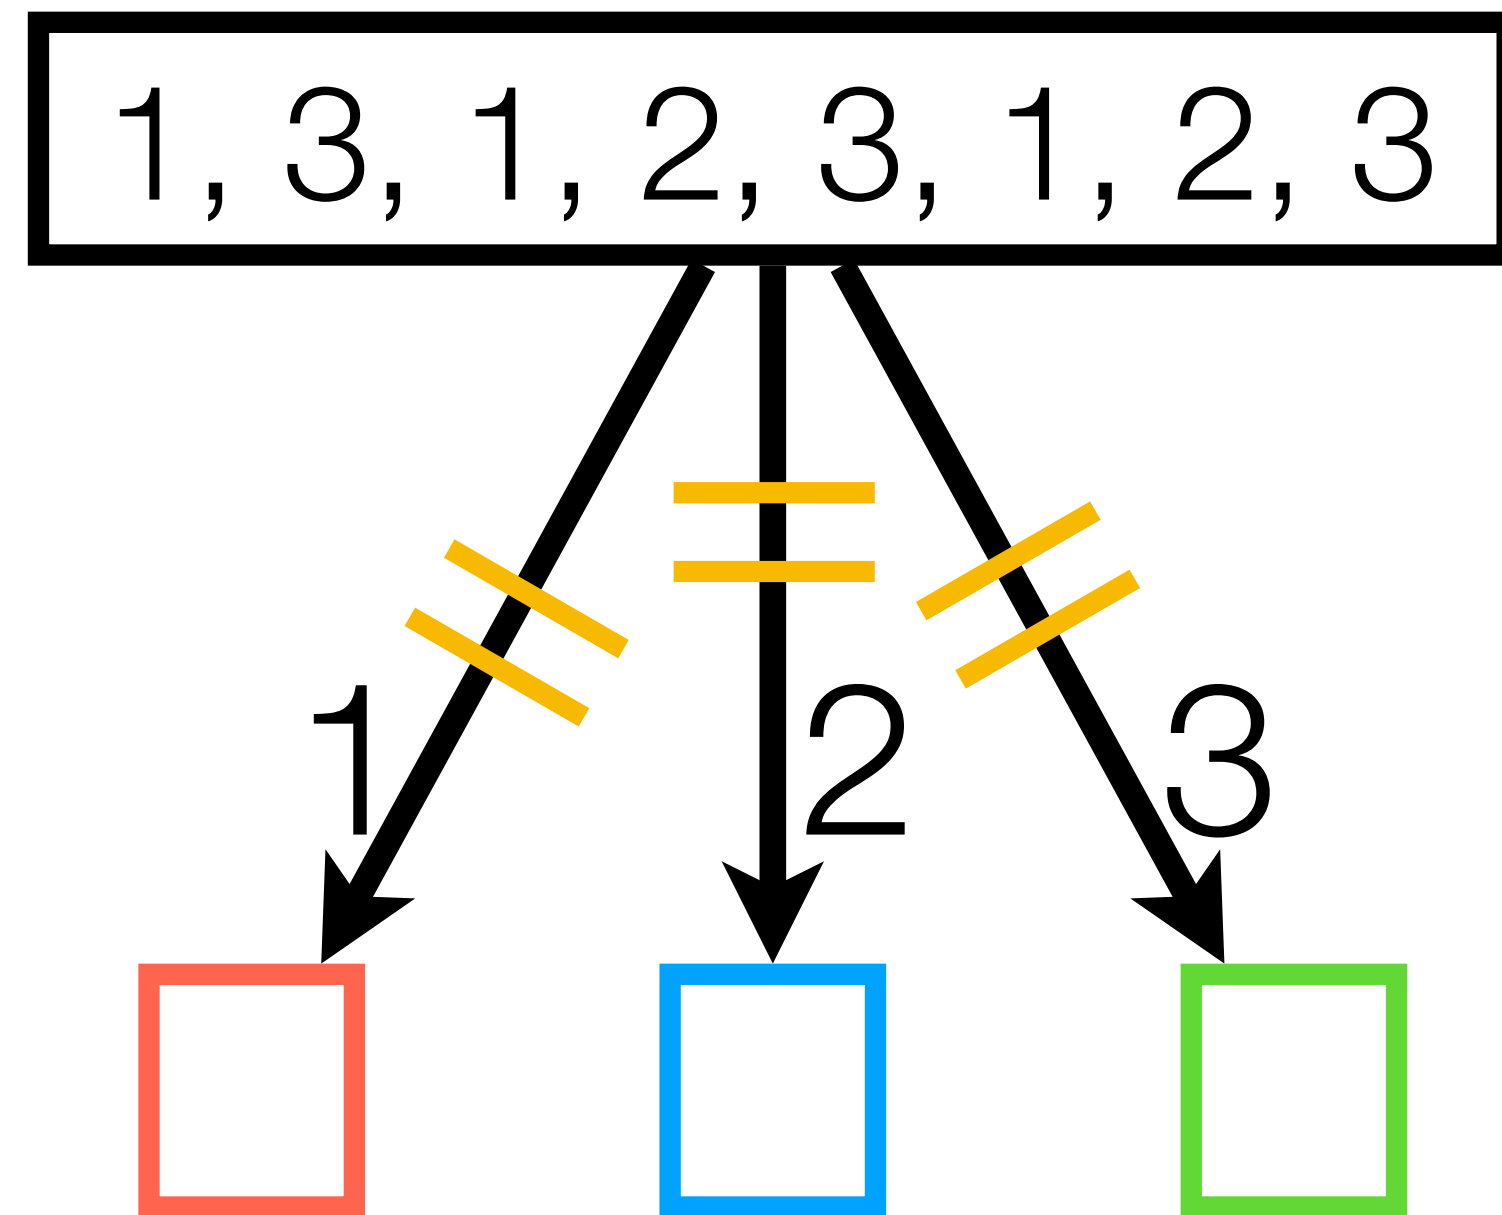
Compiling programs



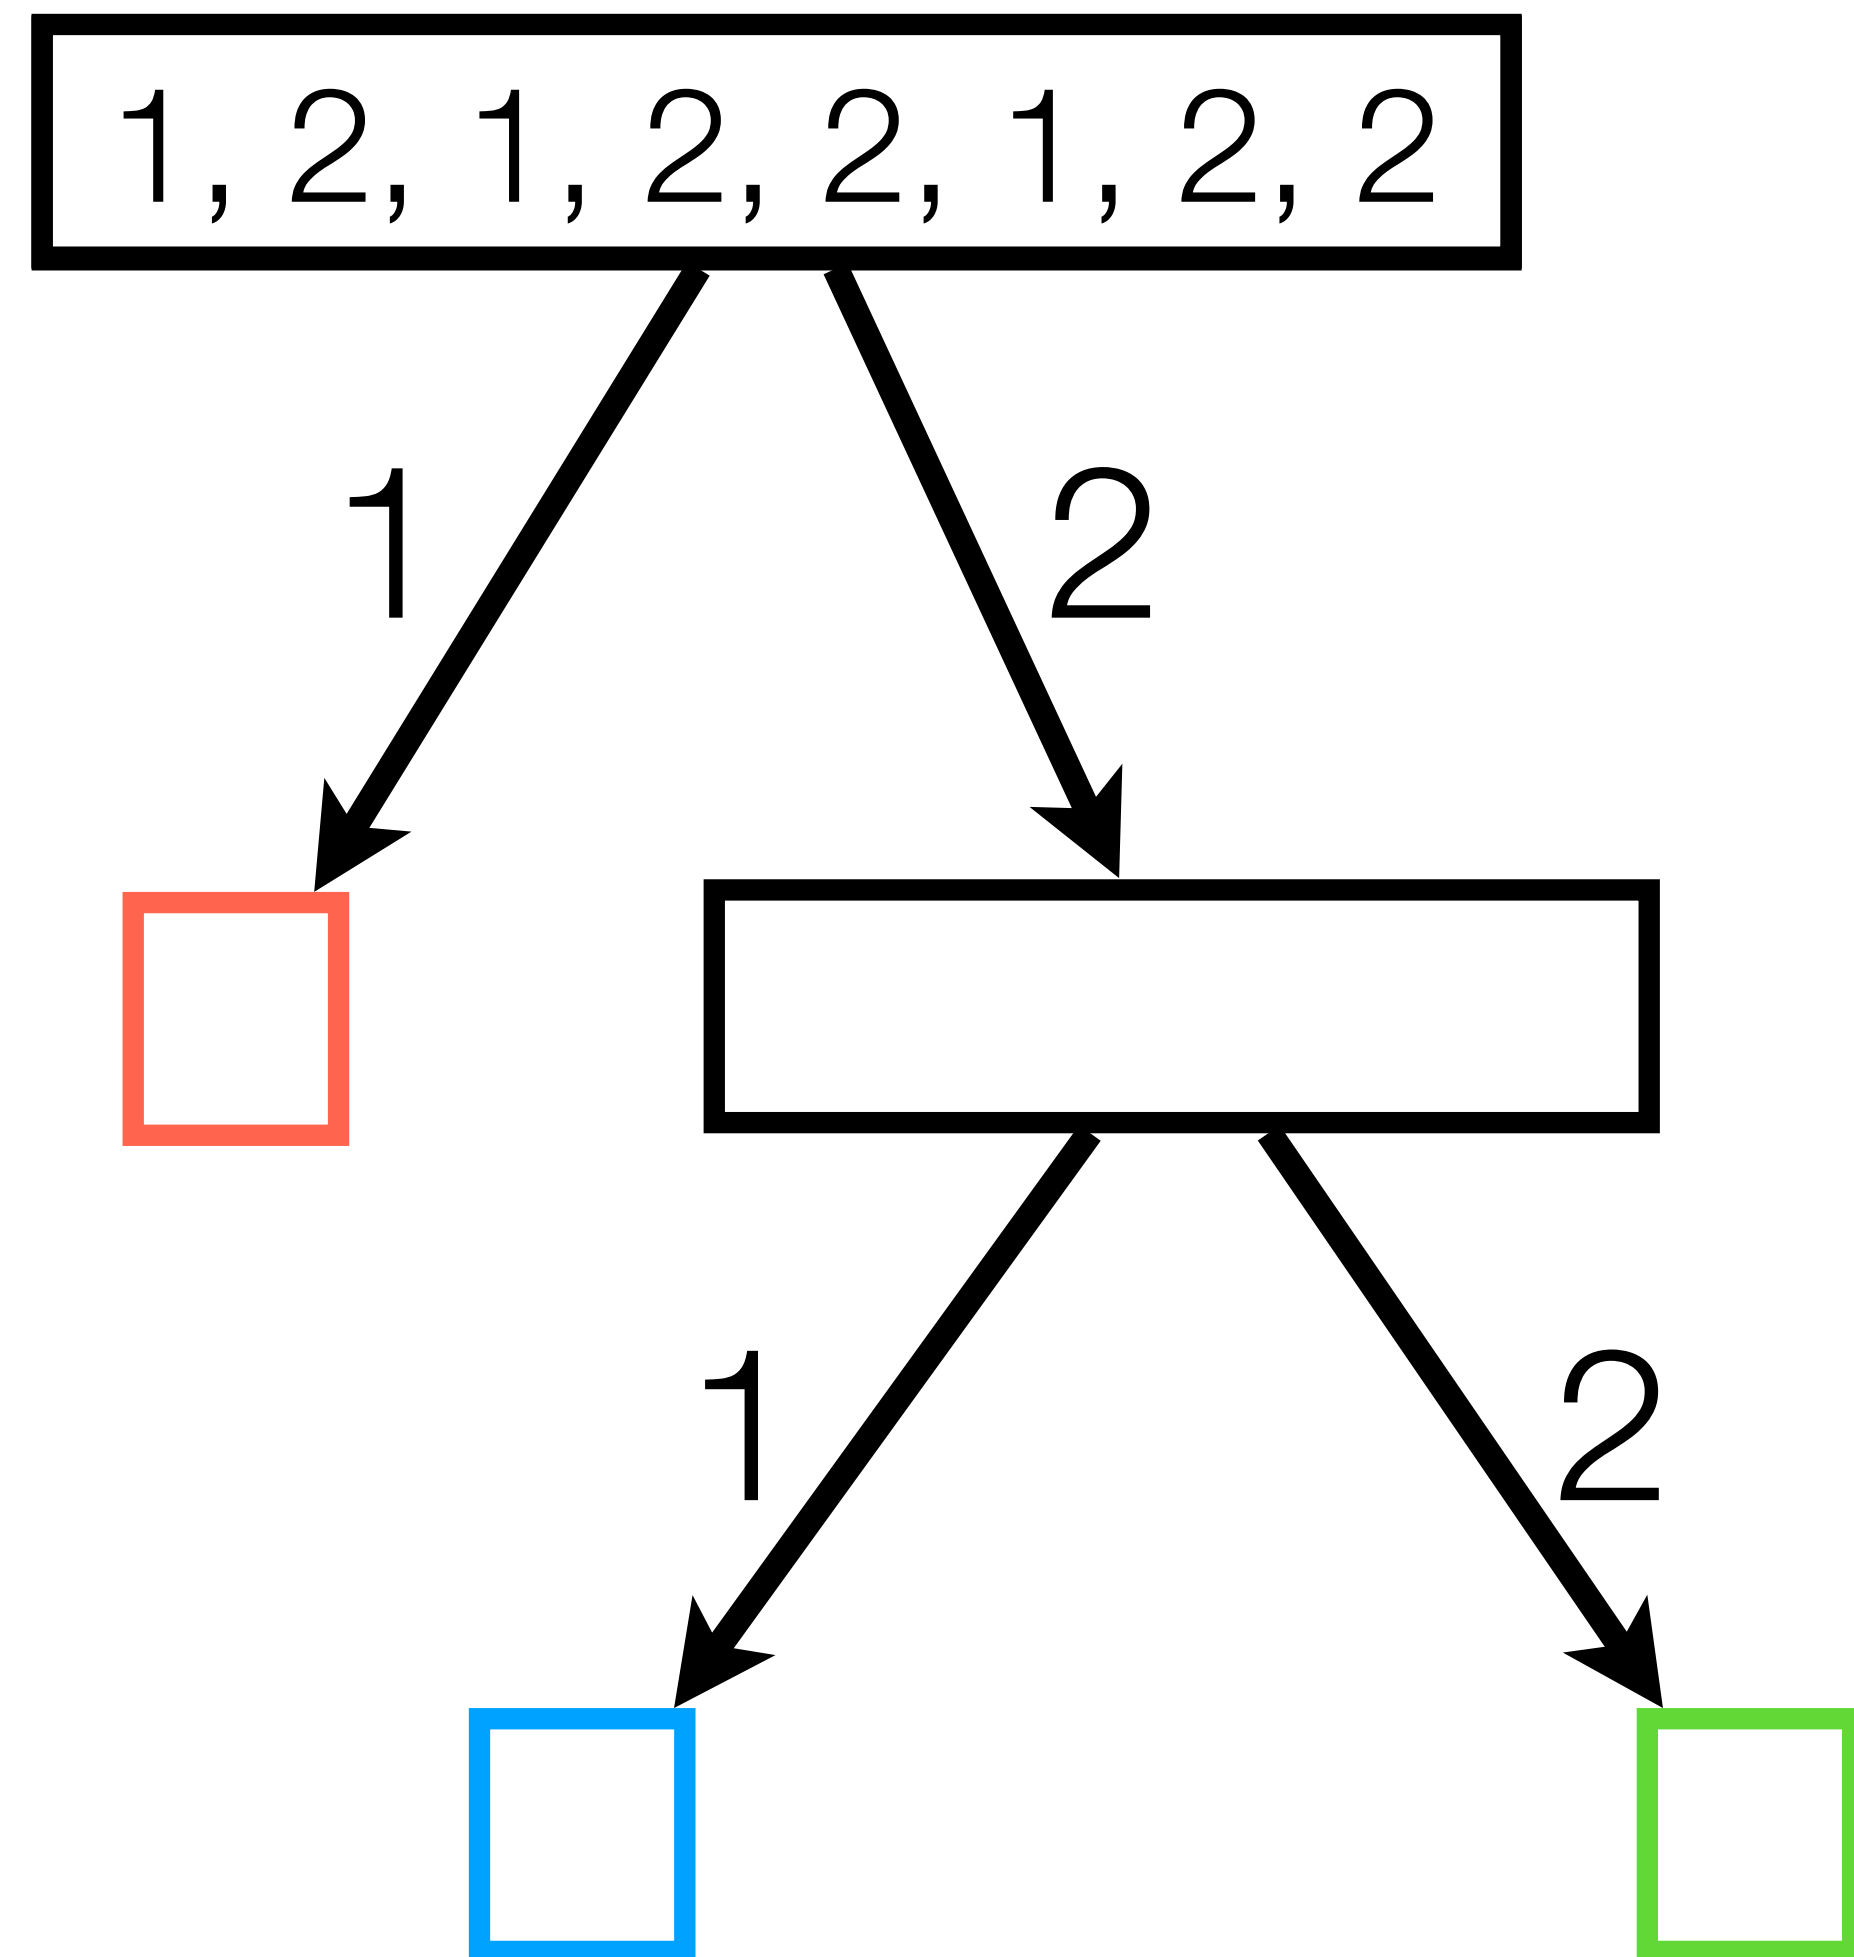
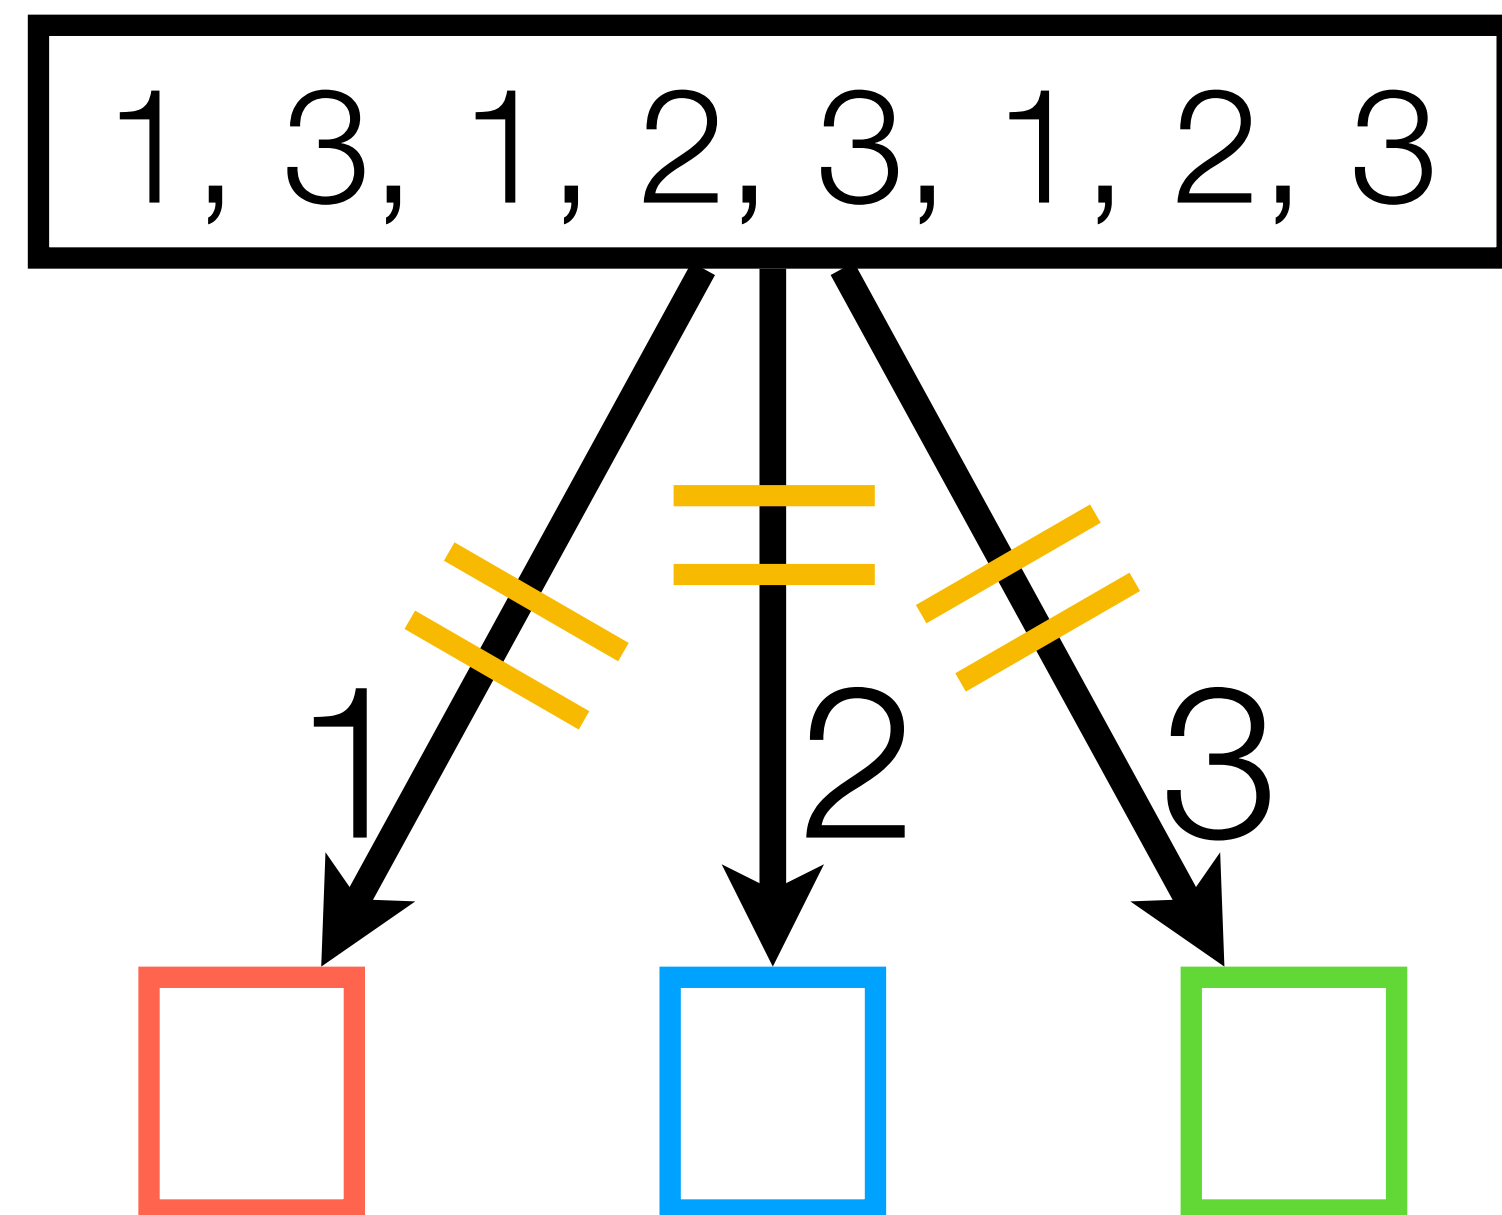
Compiling programs



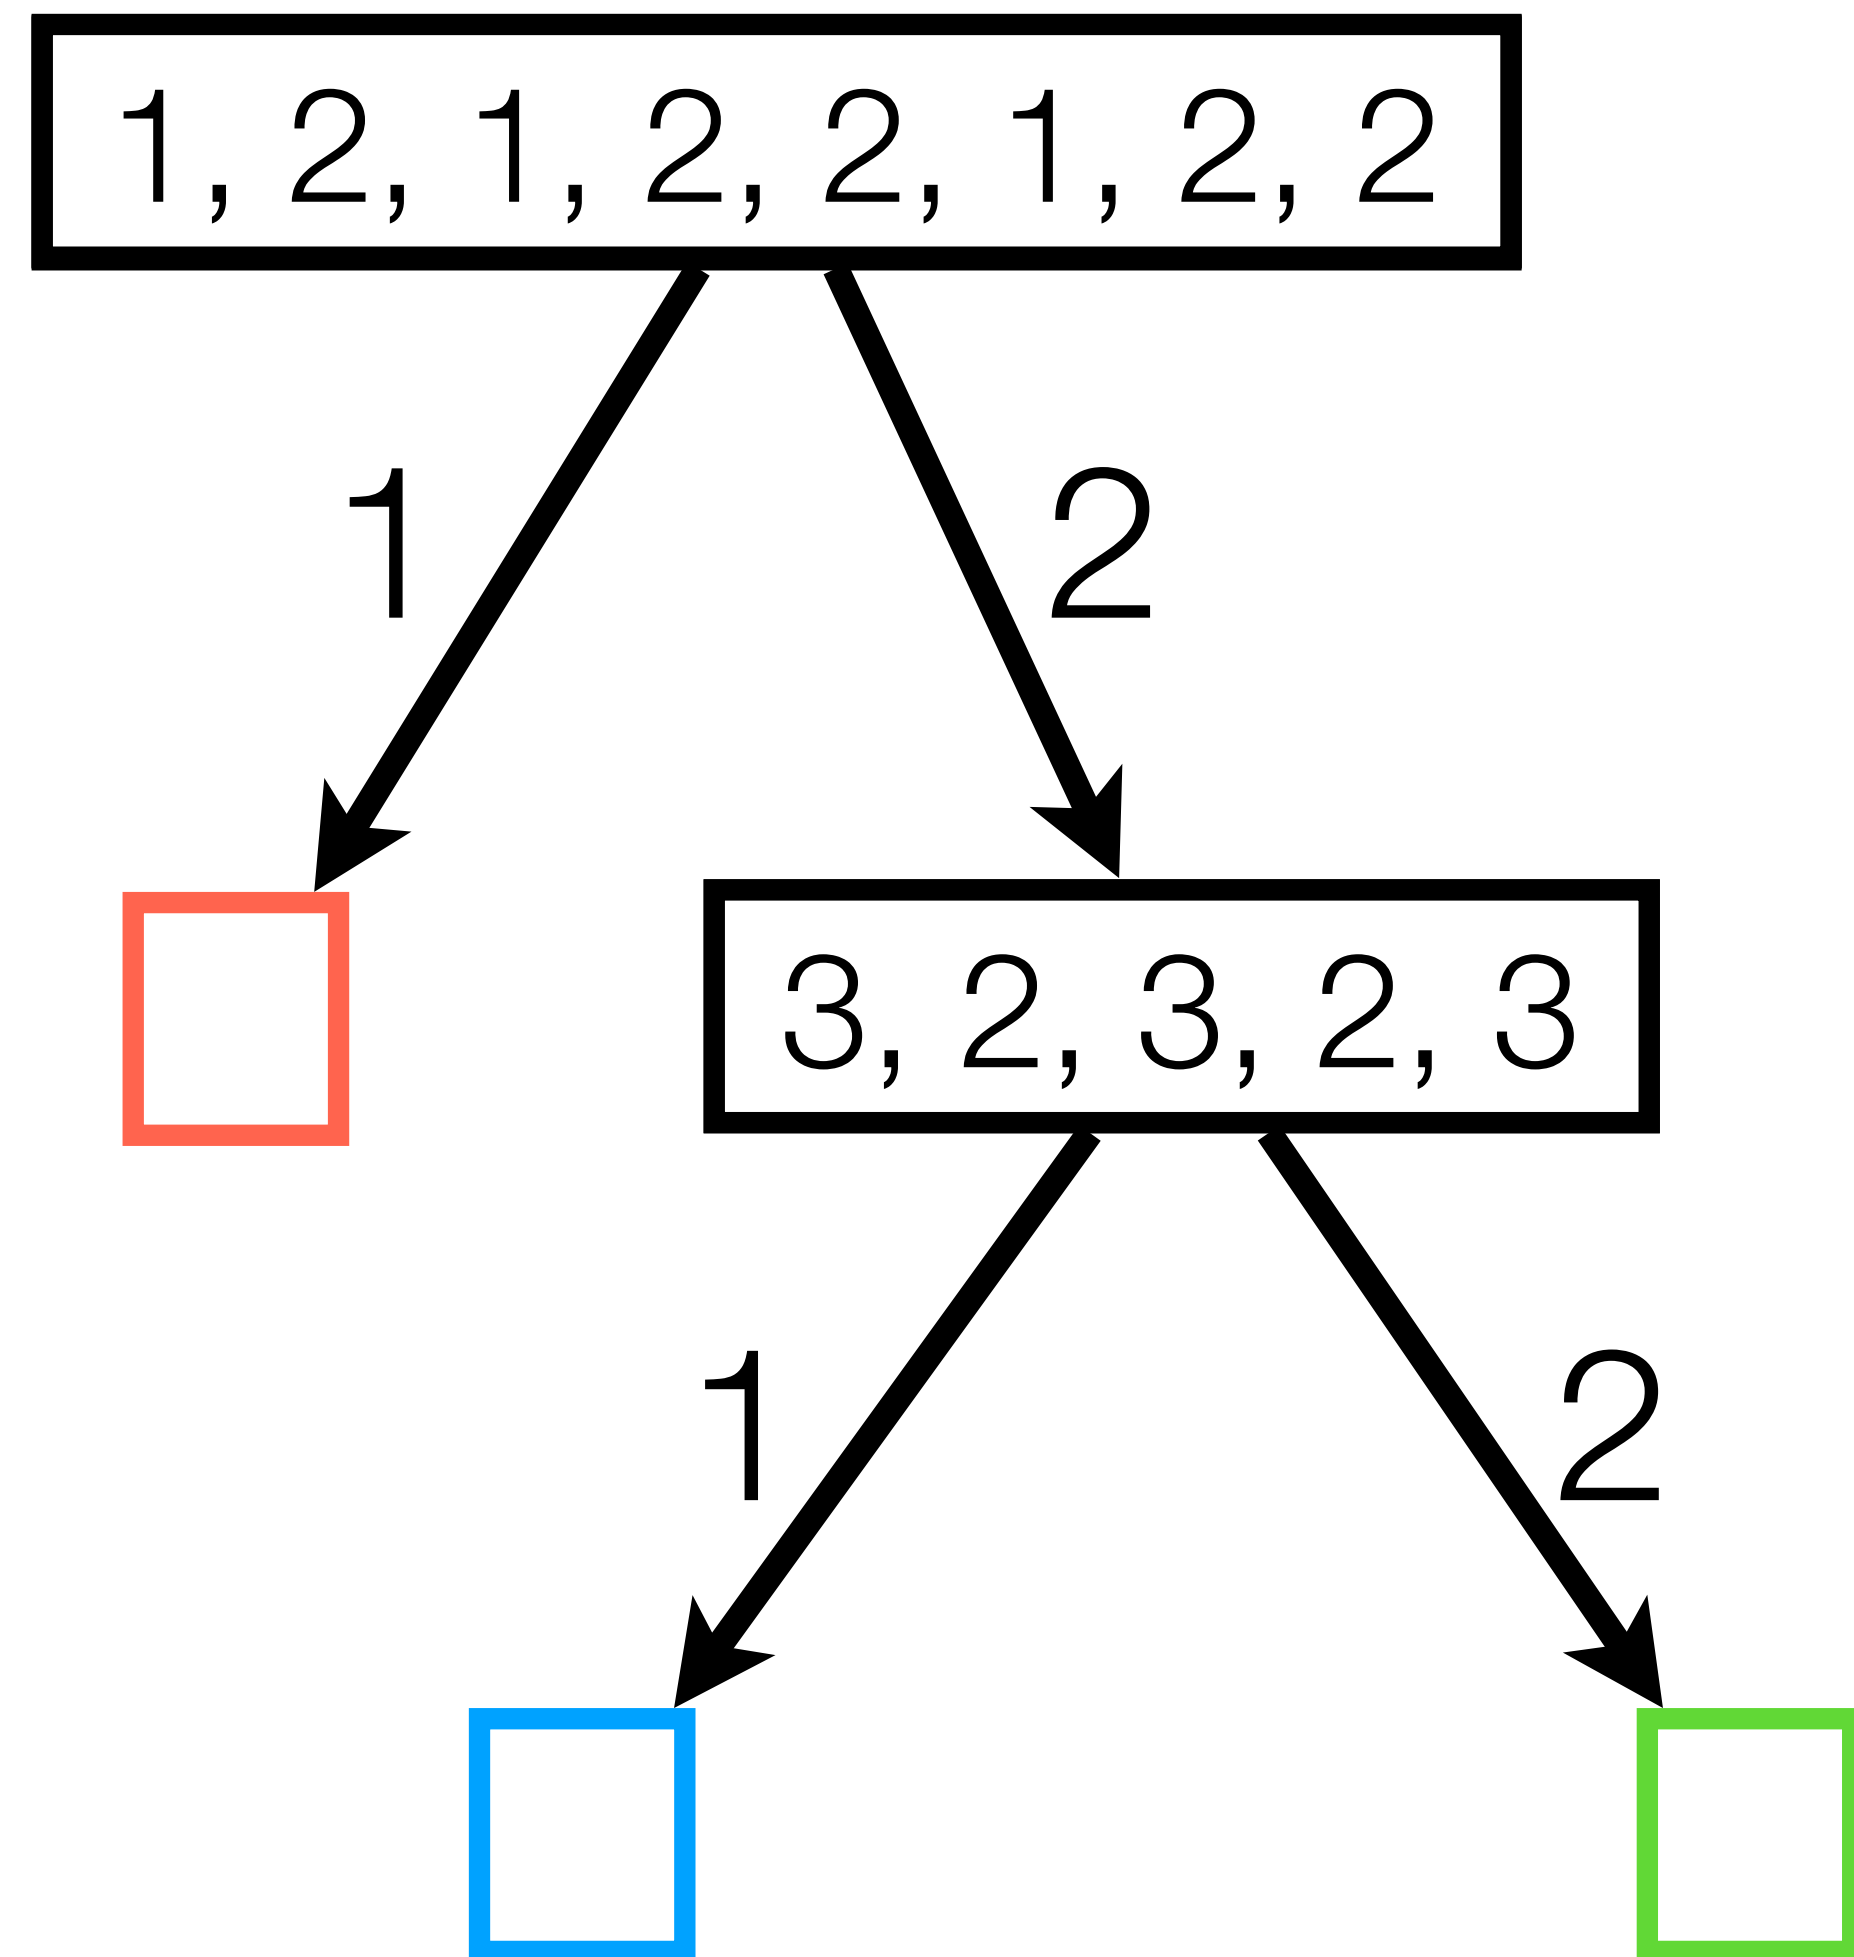
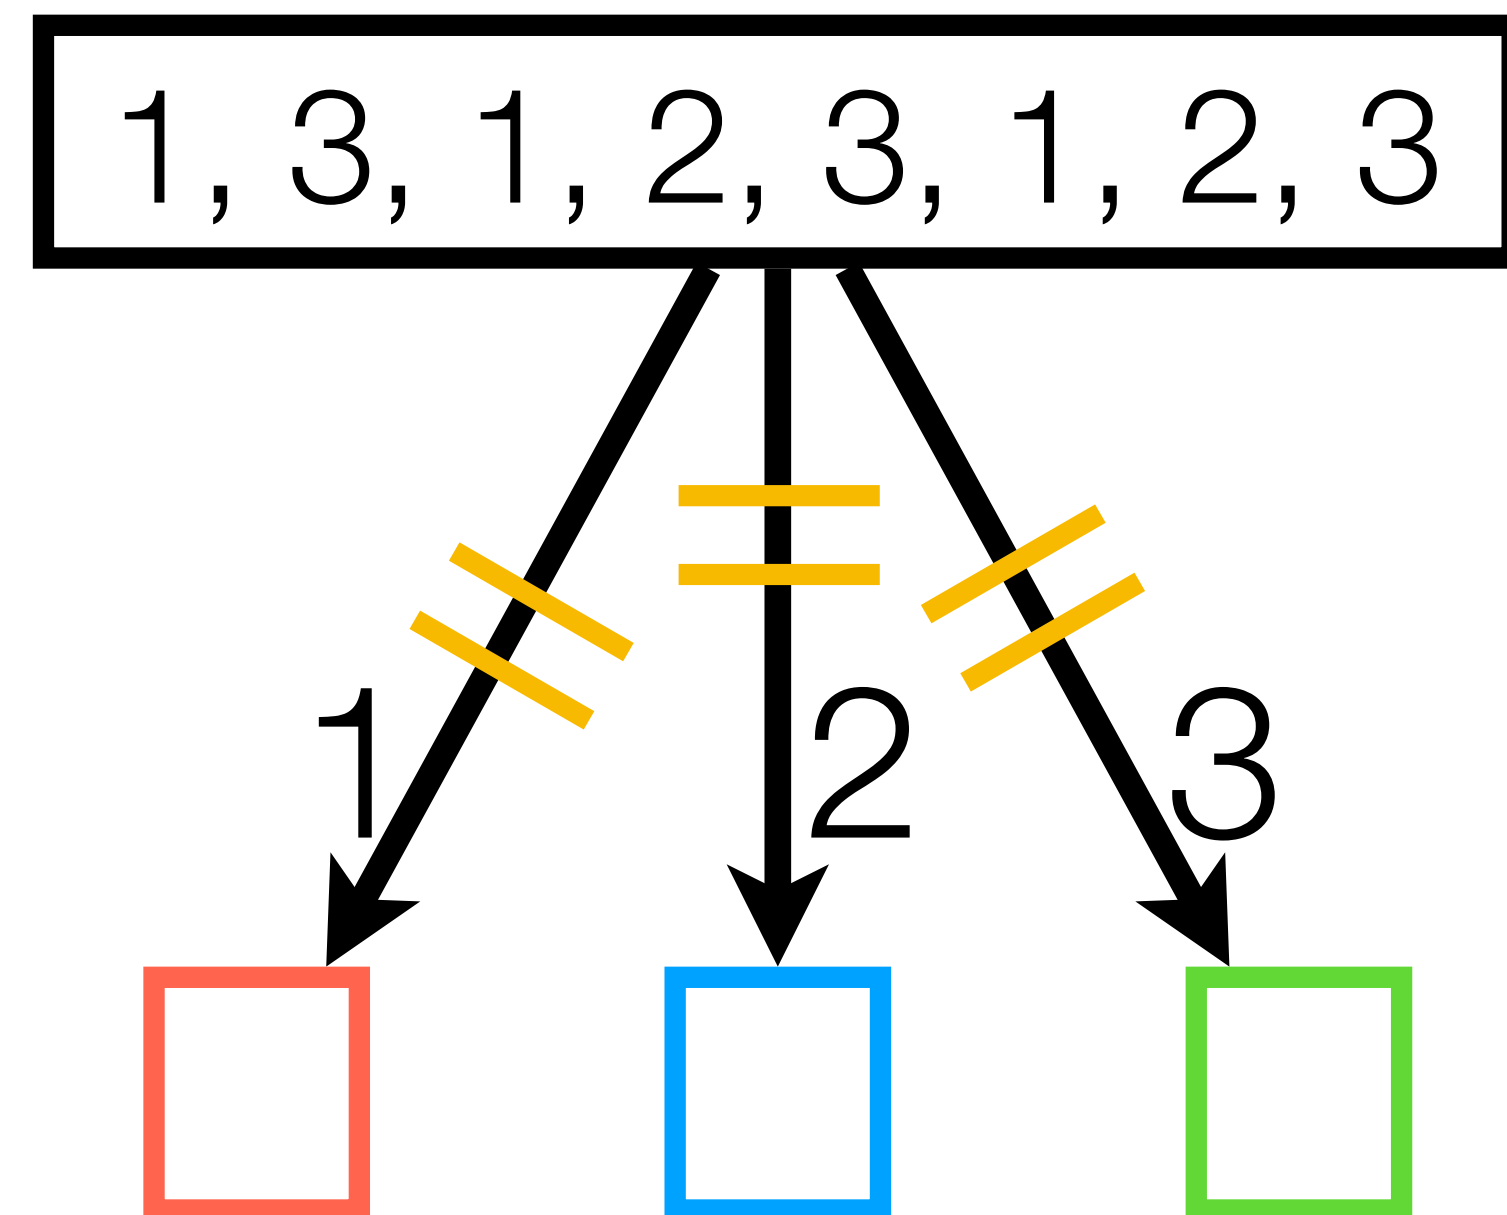
Compiling programs



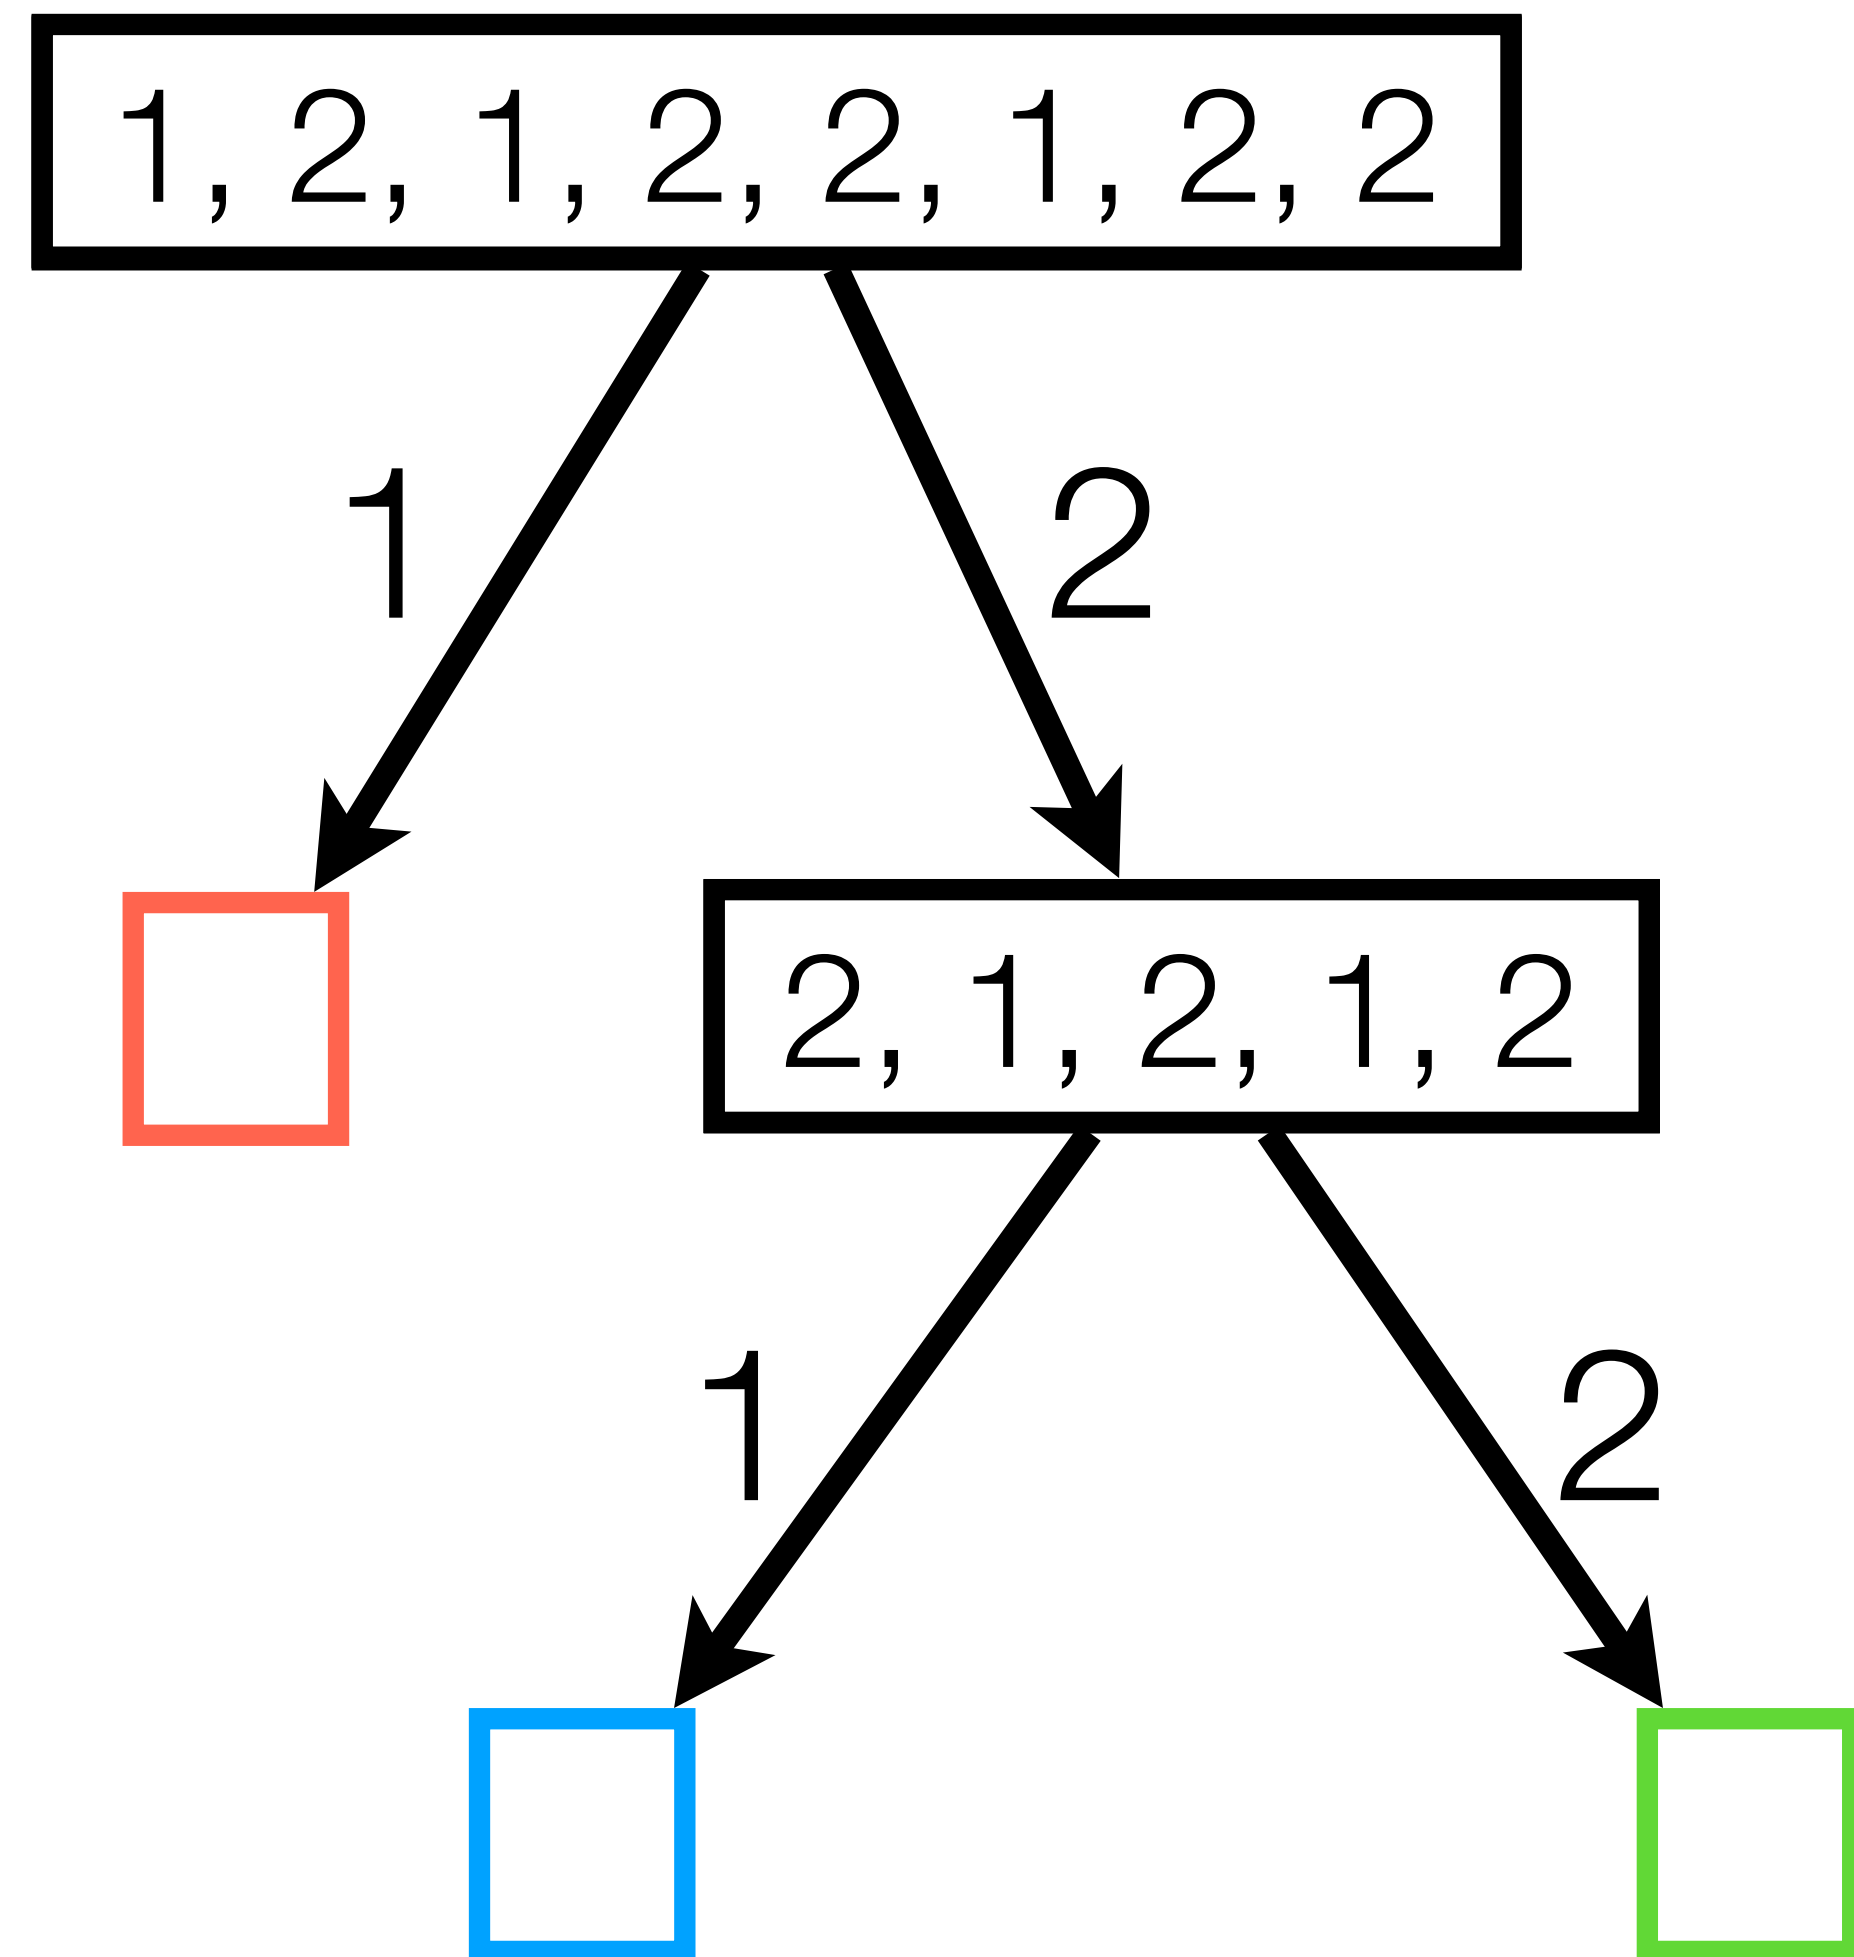
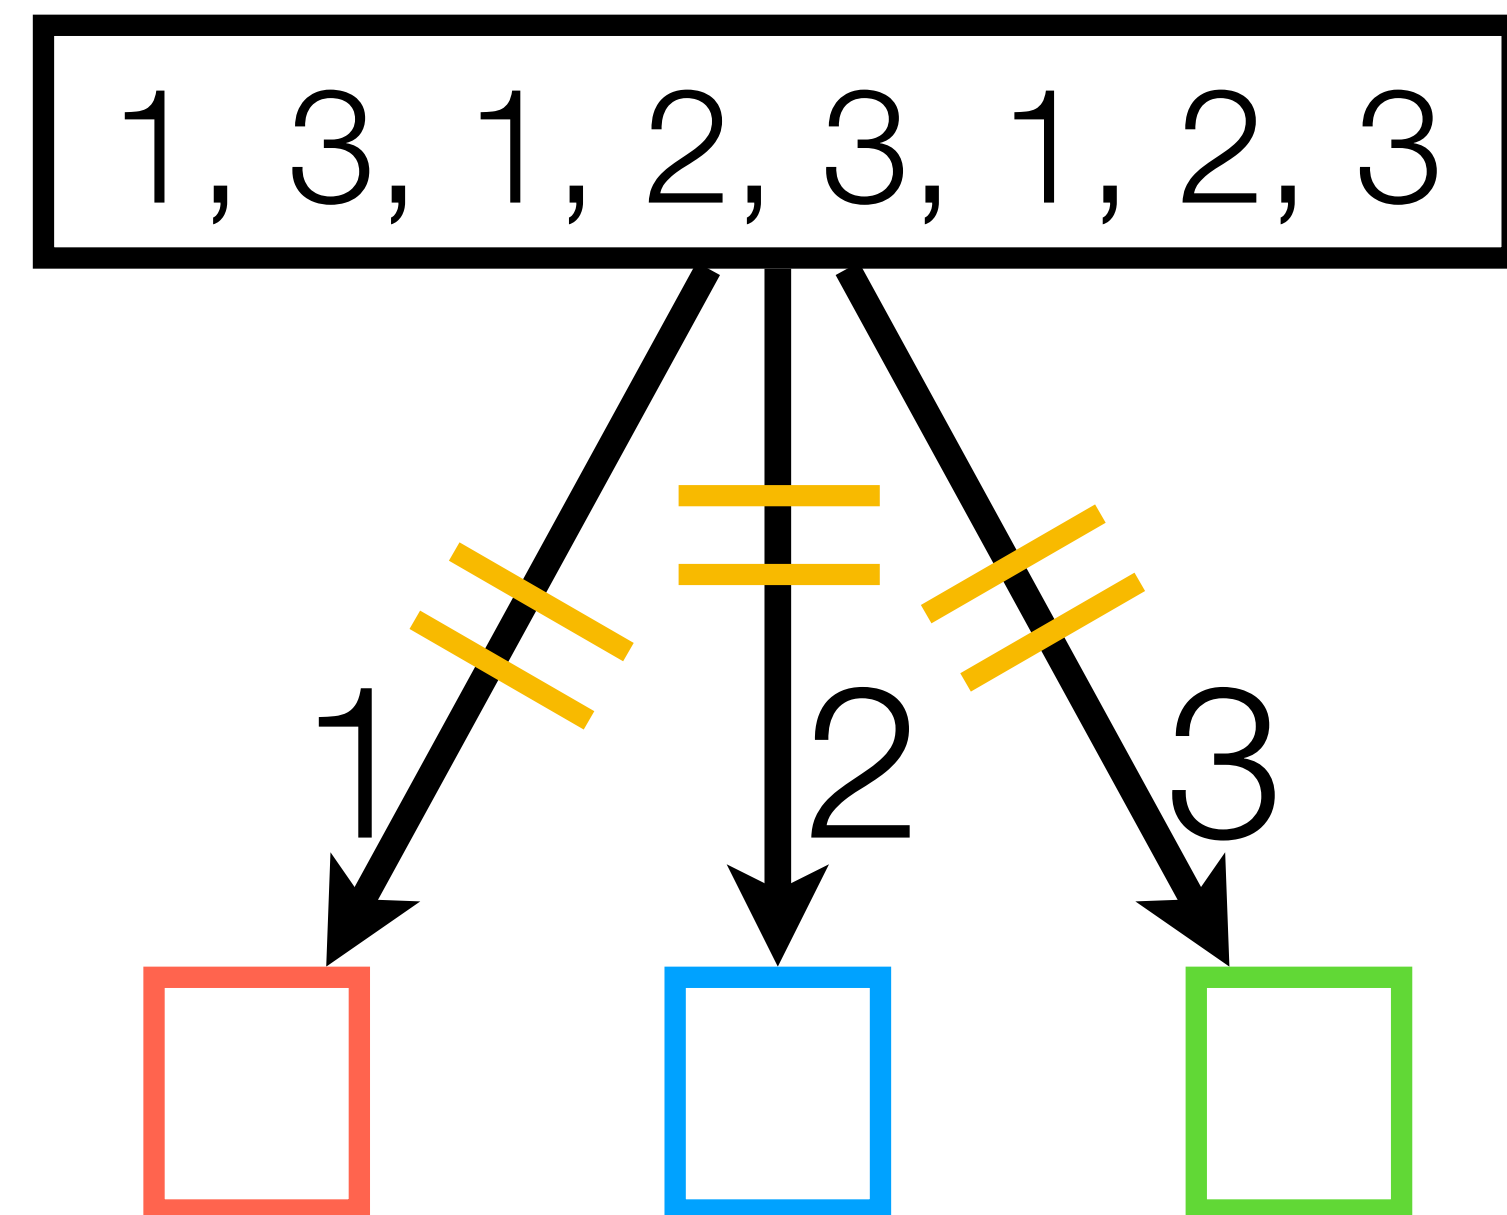
Compiling programs



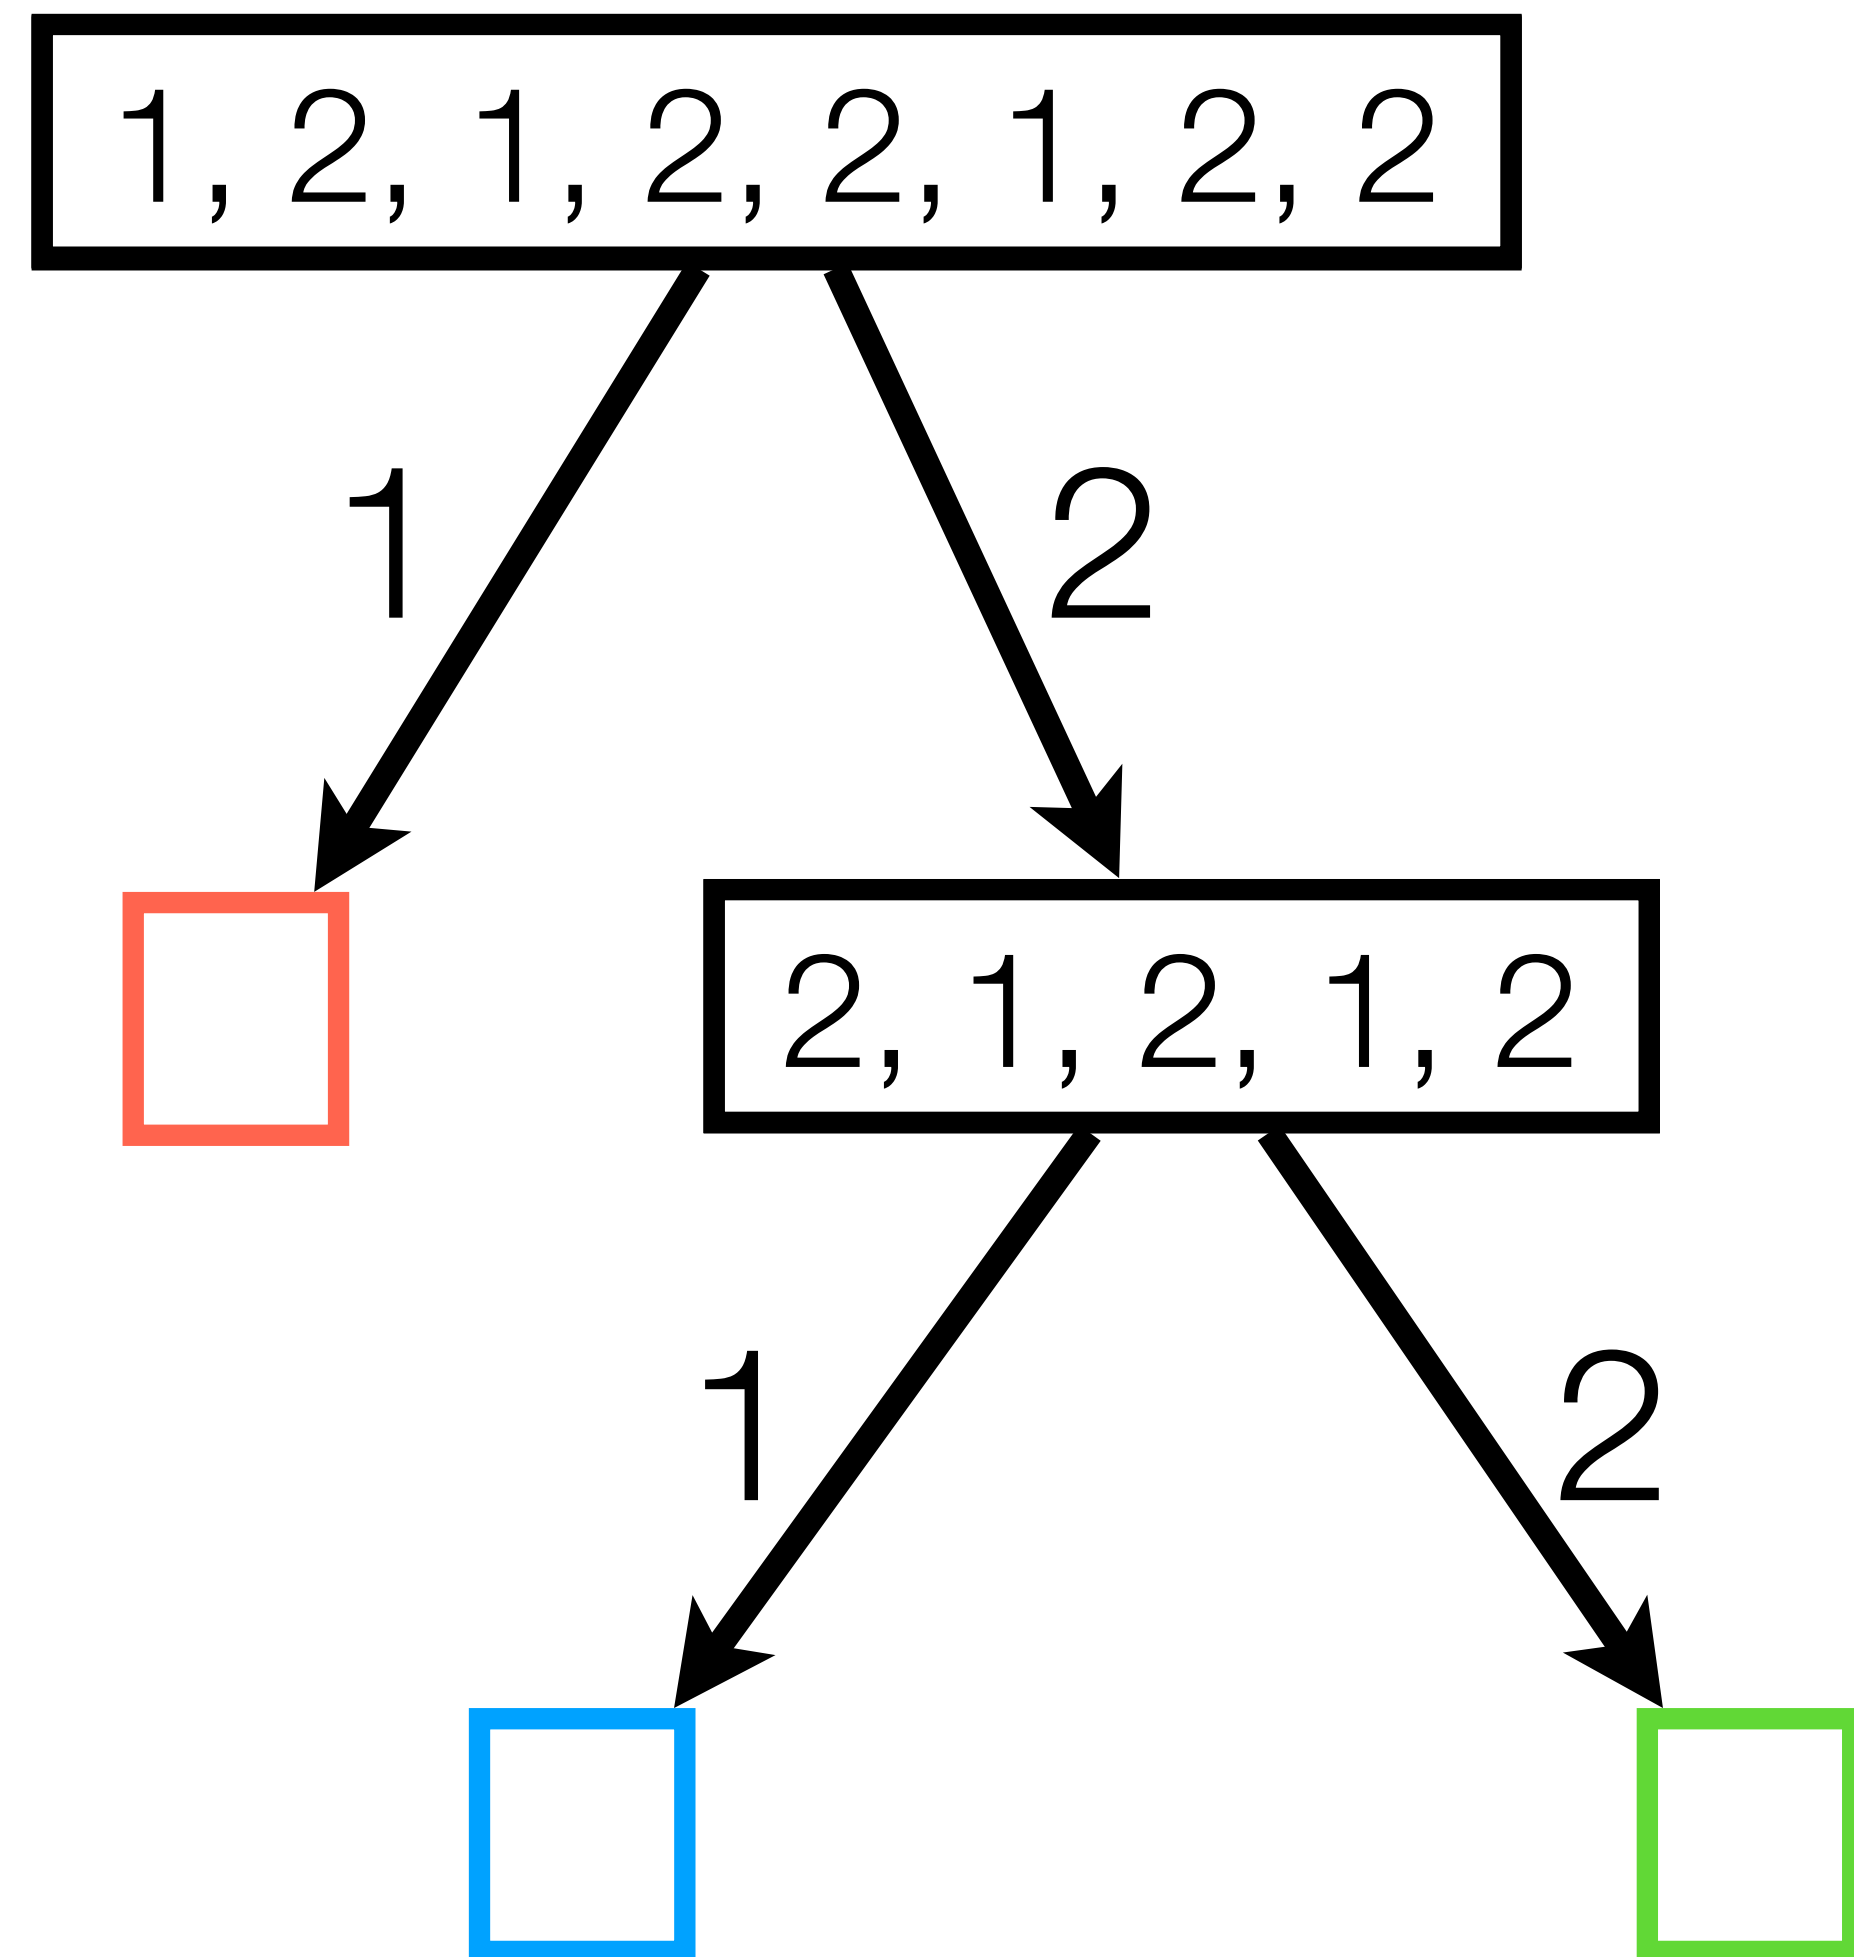
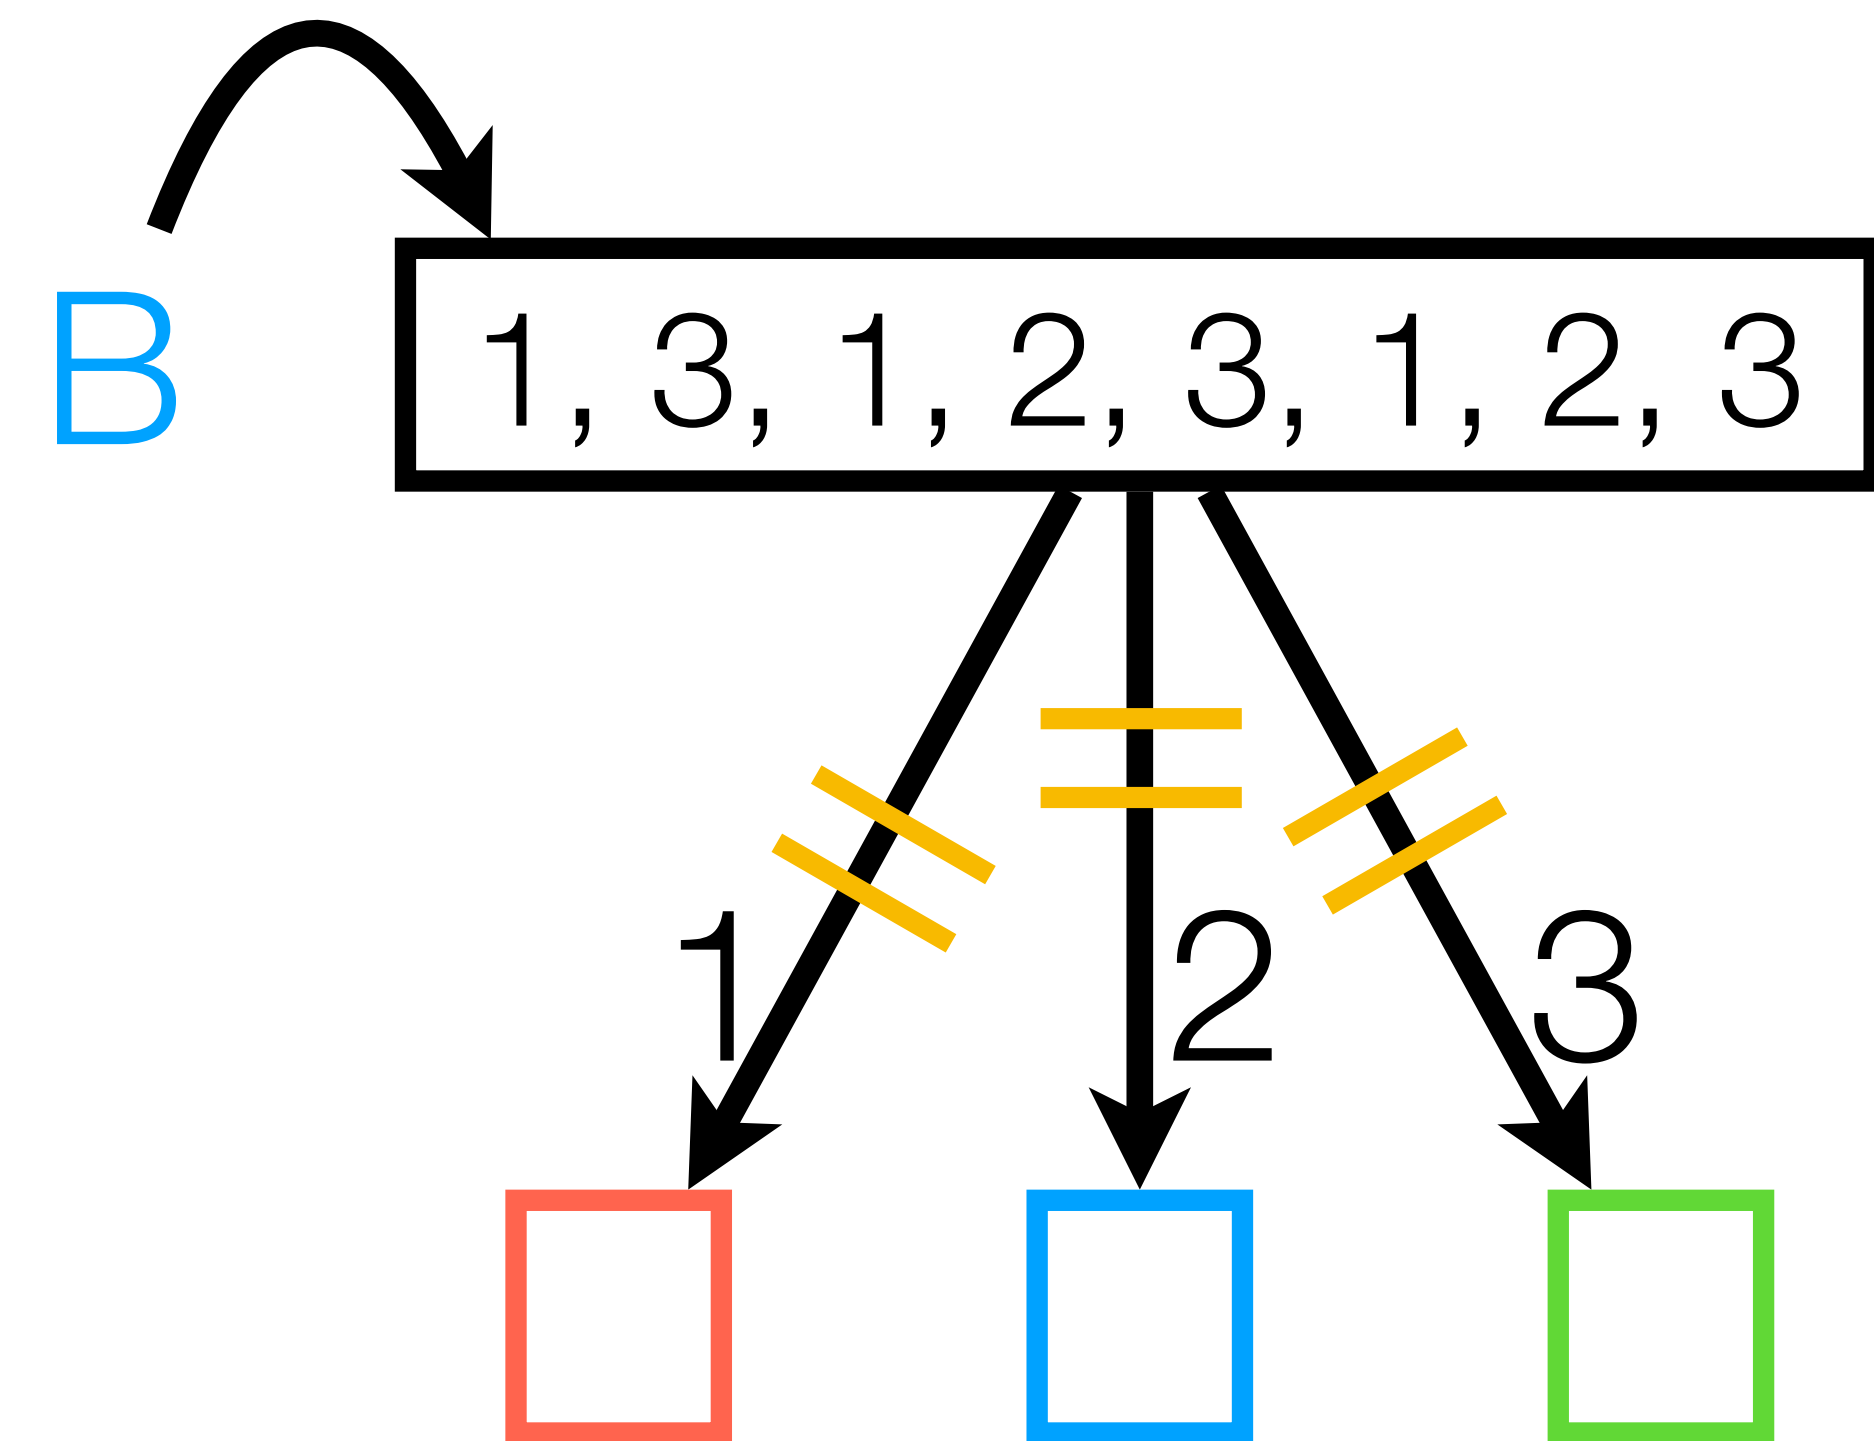
Compiling programs



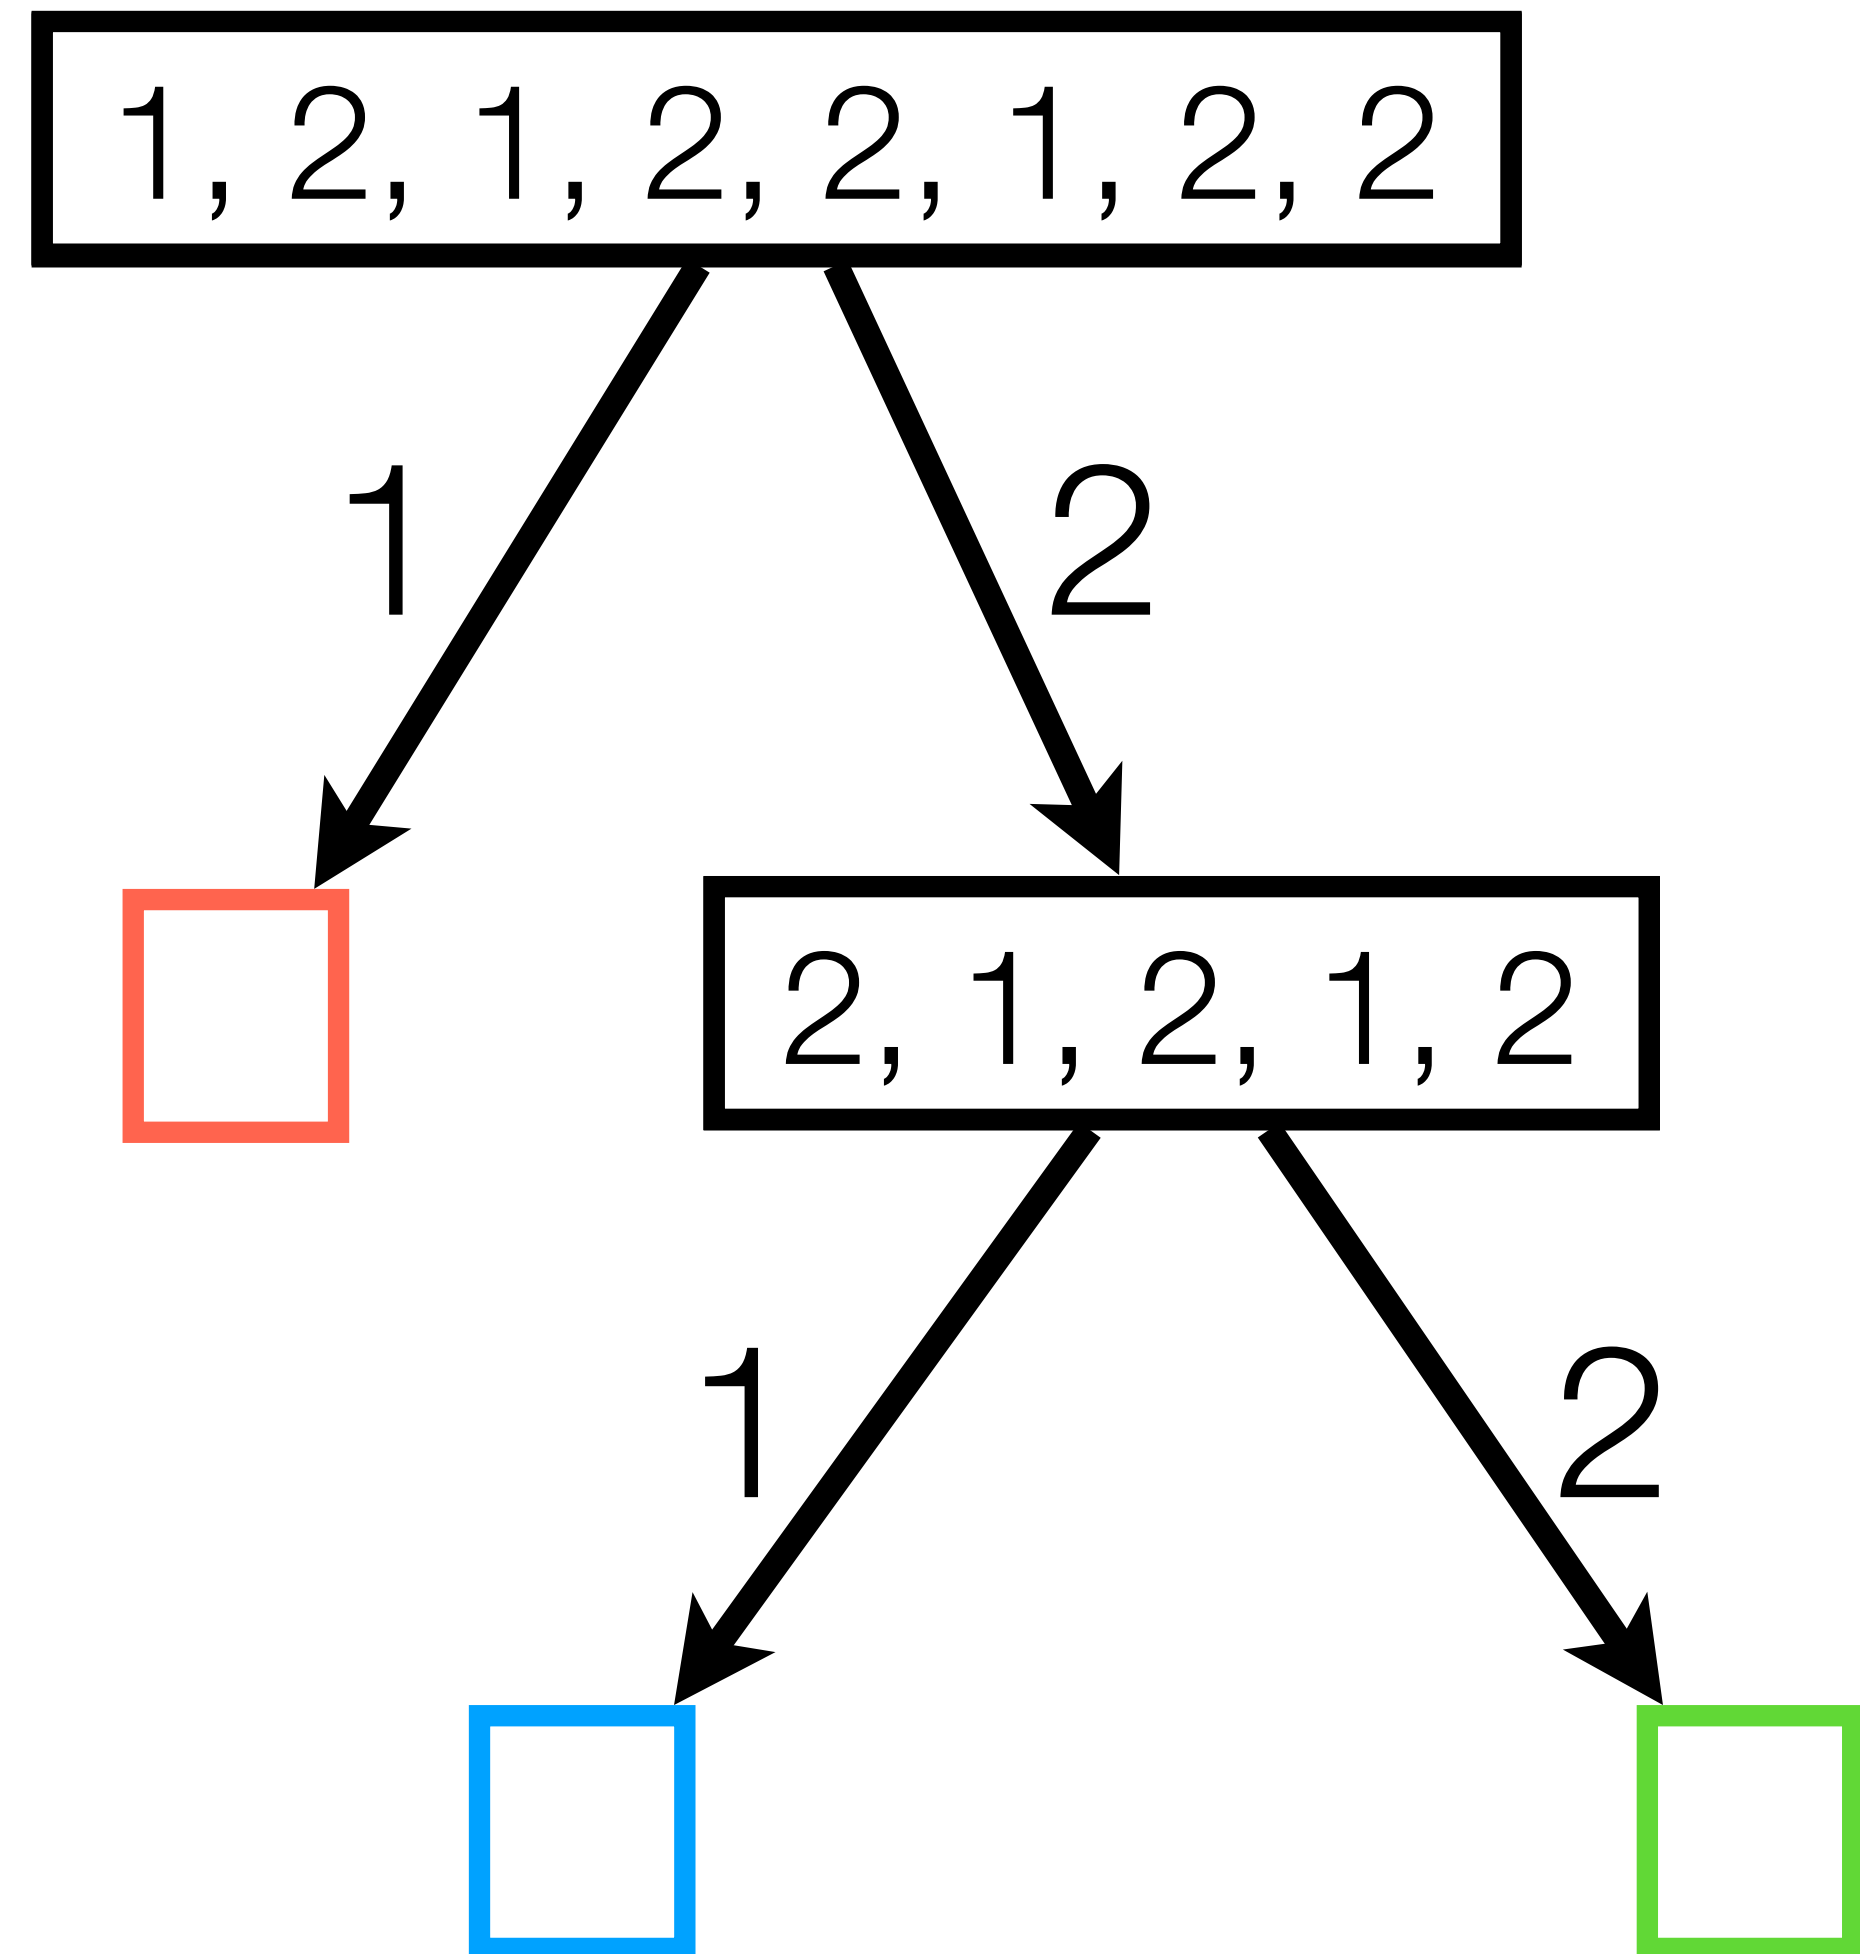
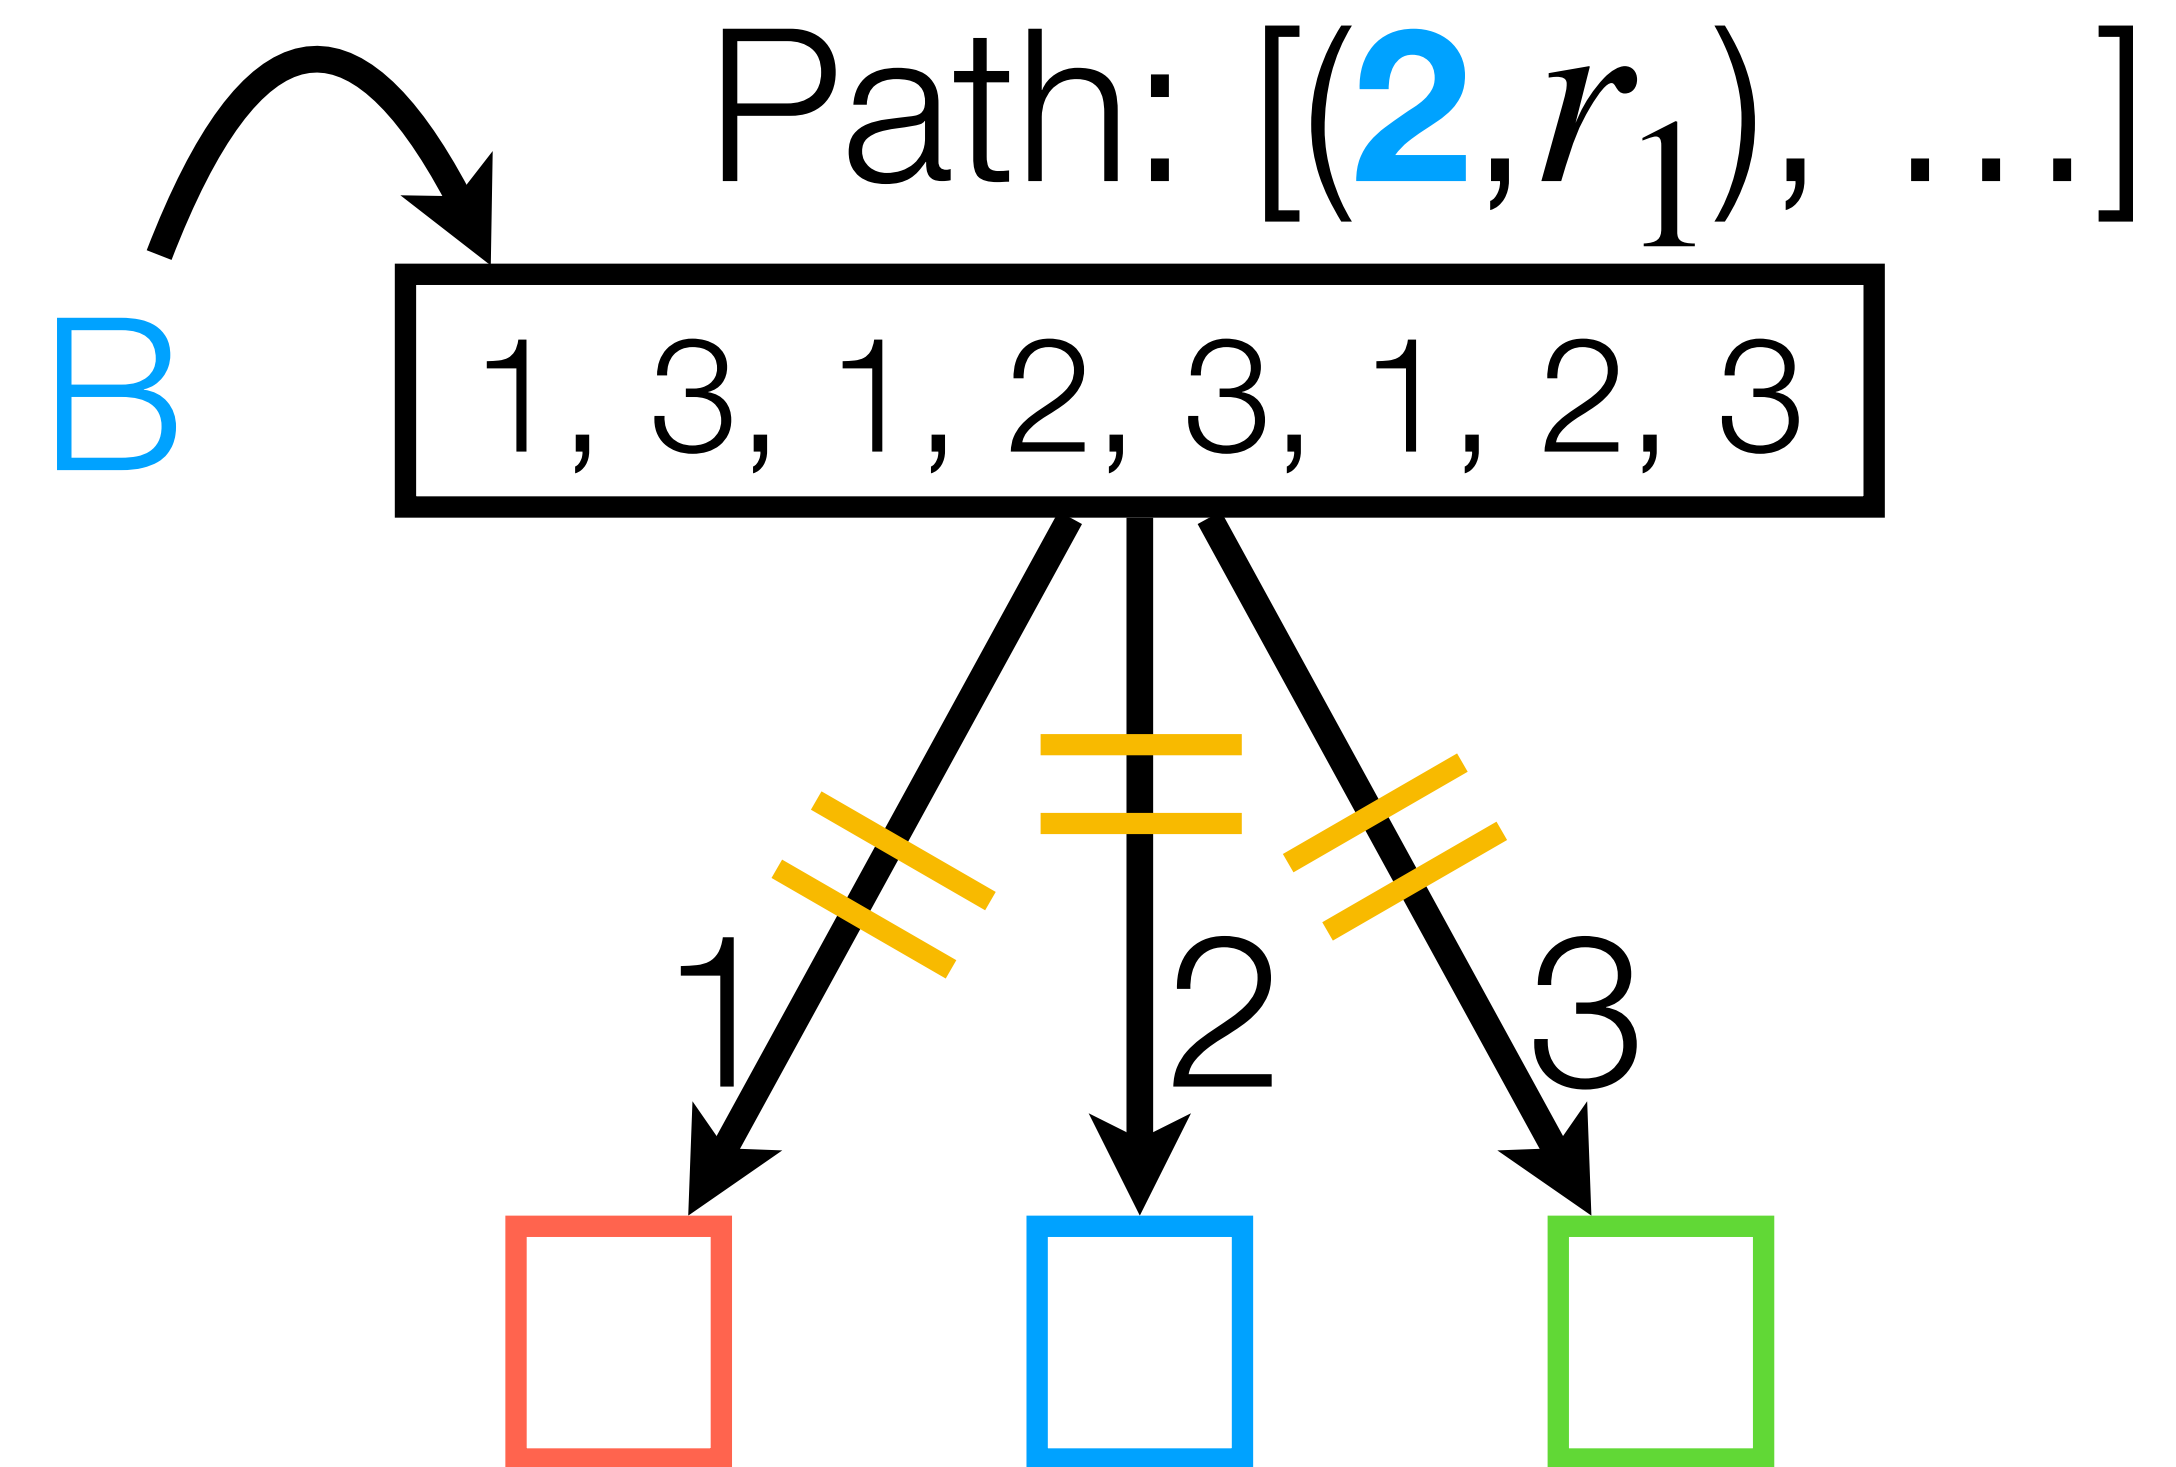
Compiling programs



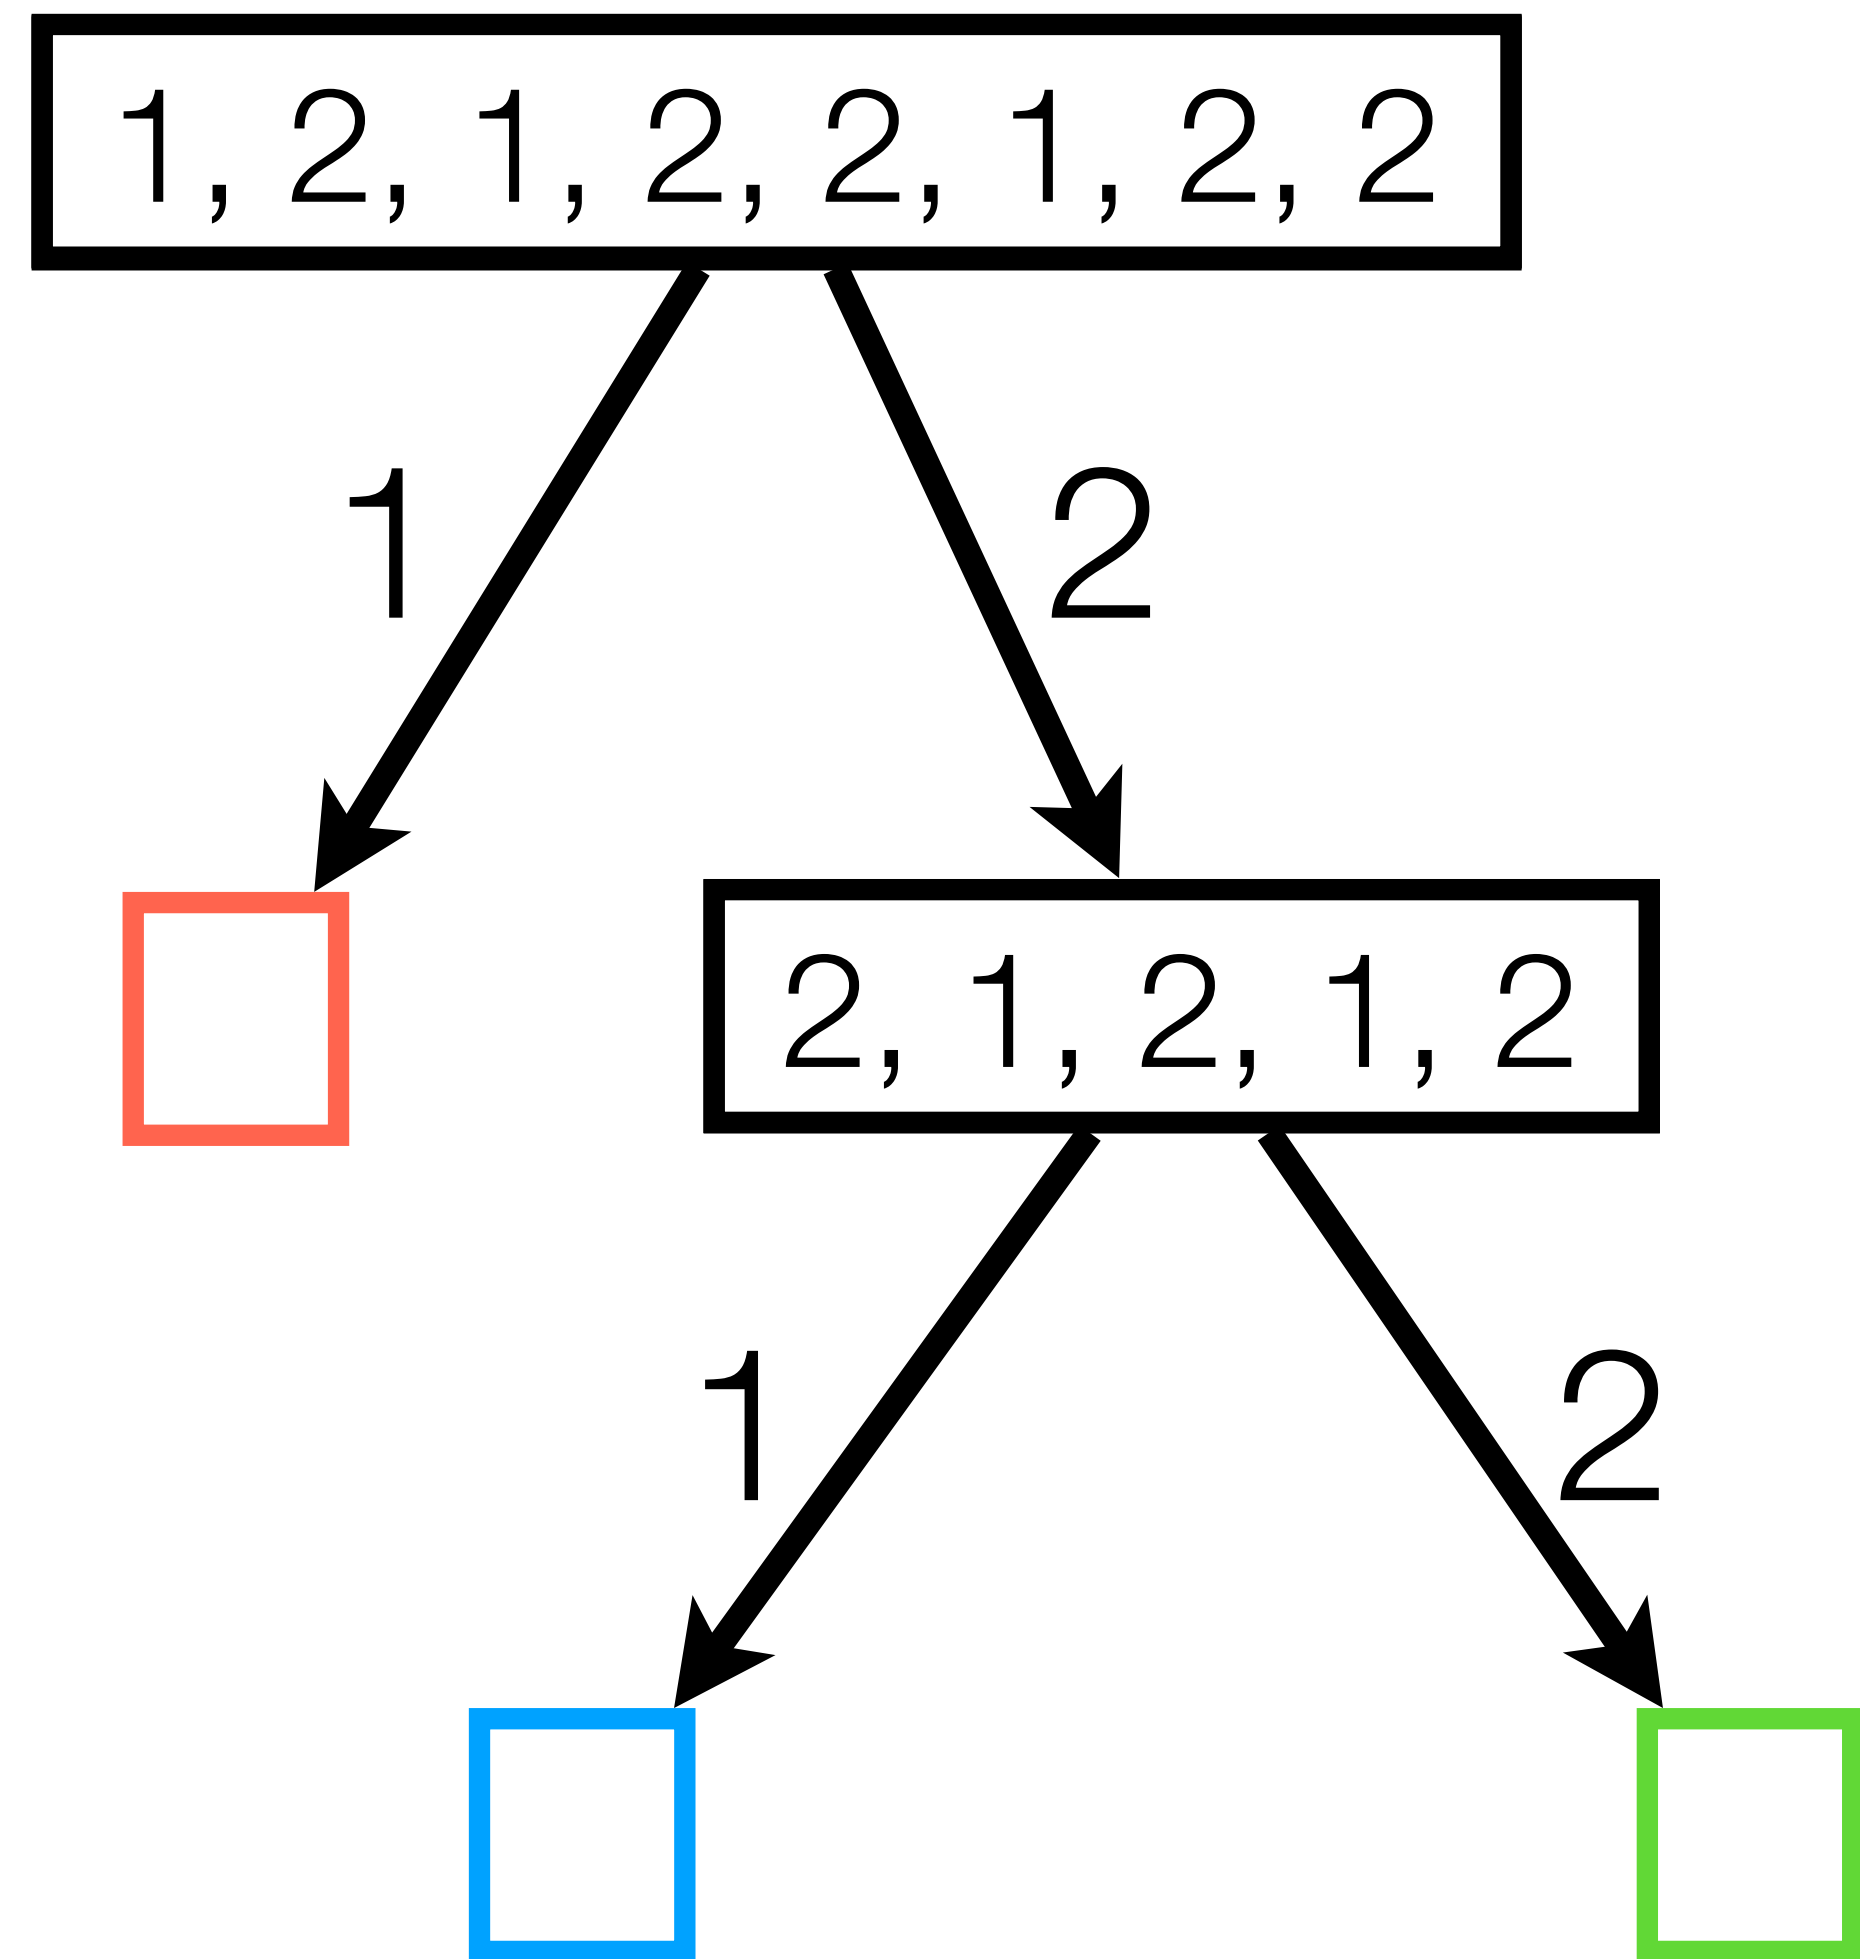
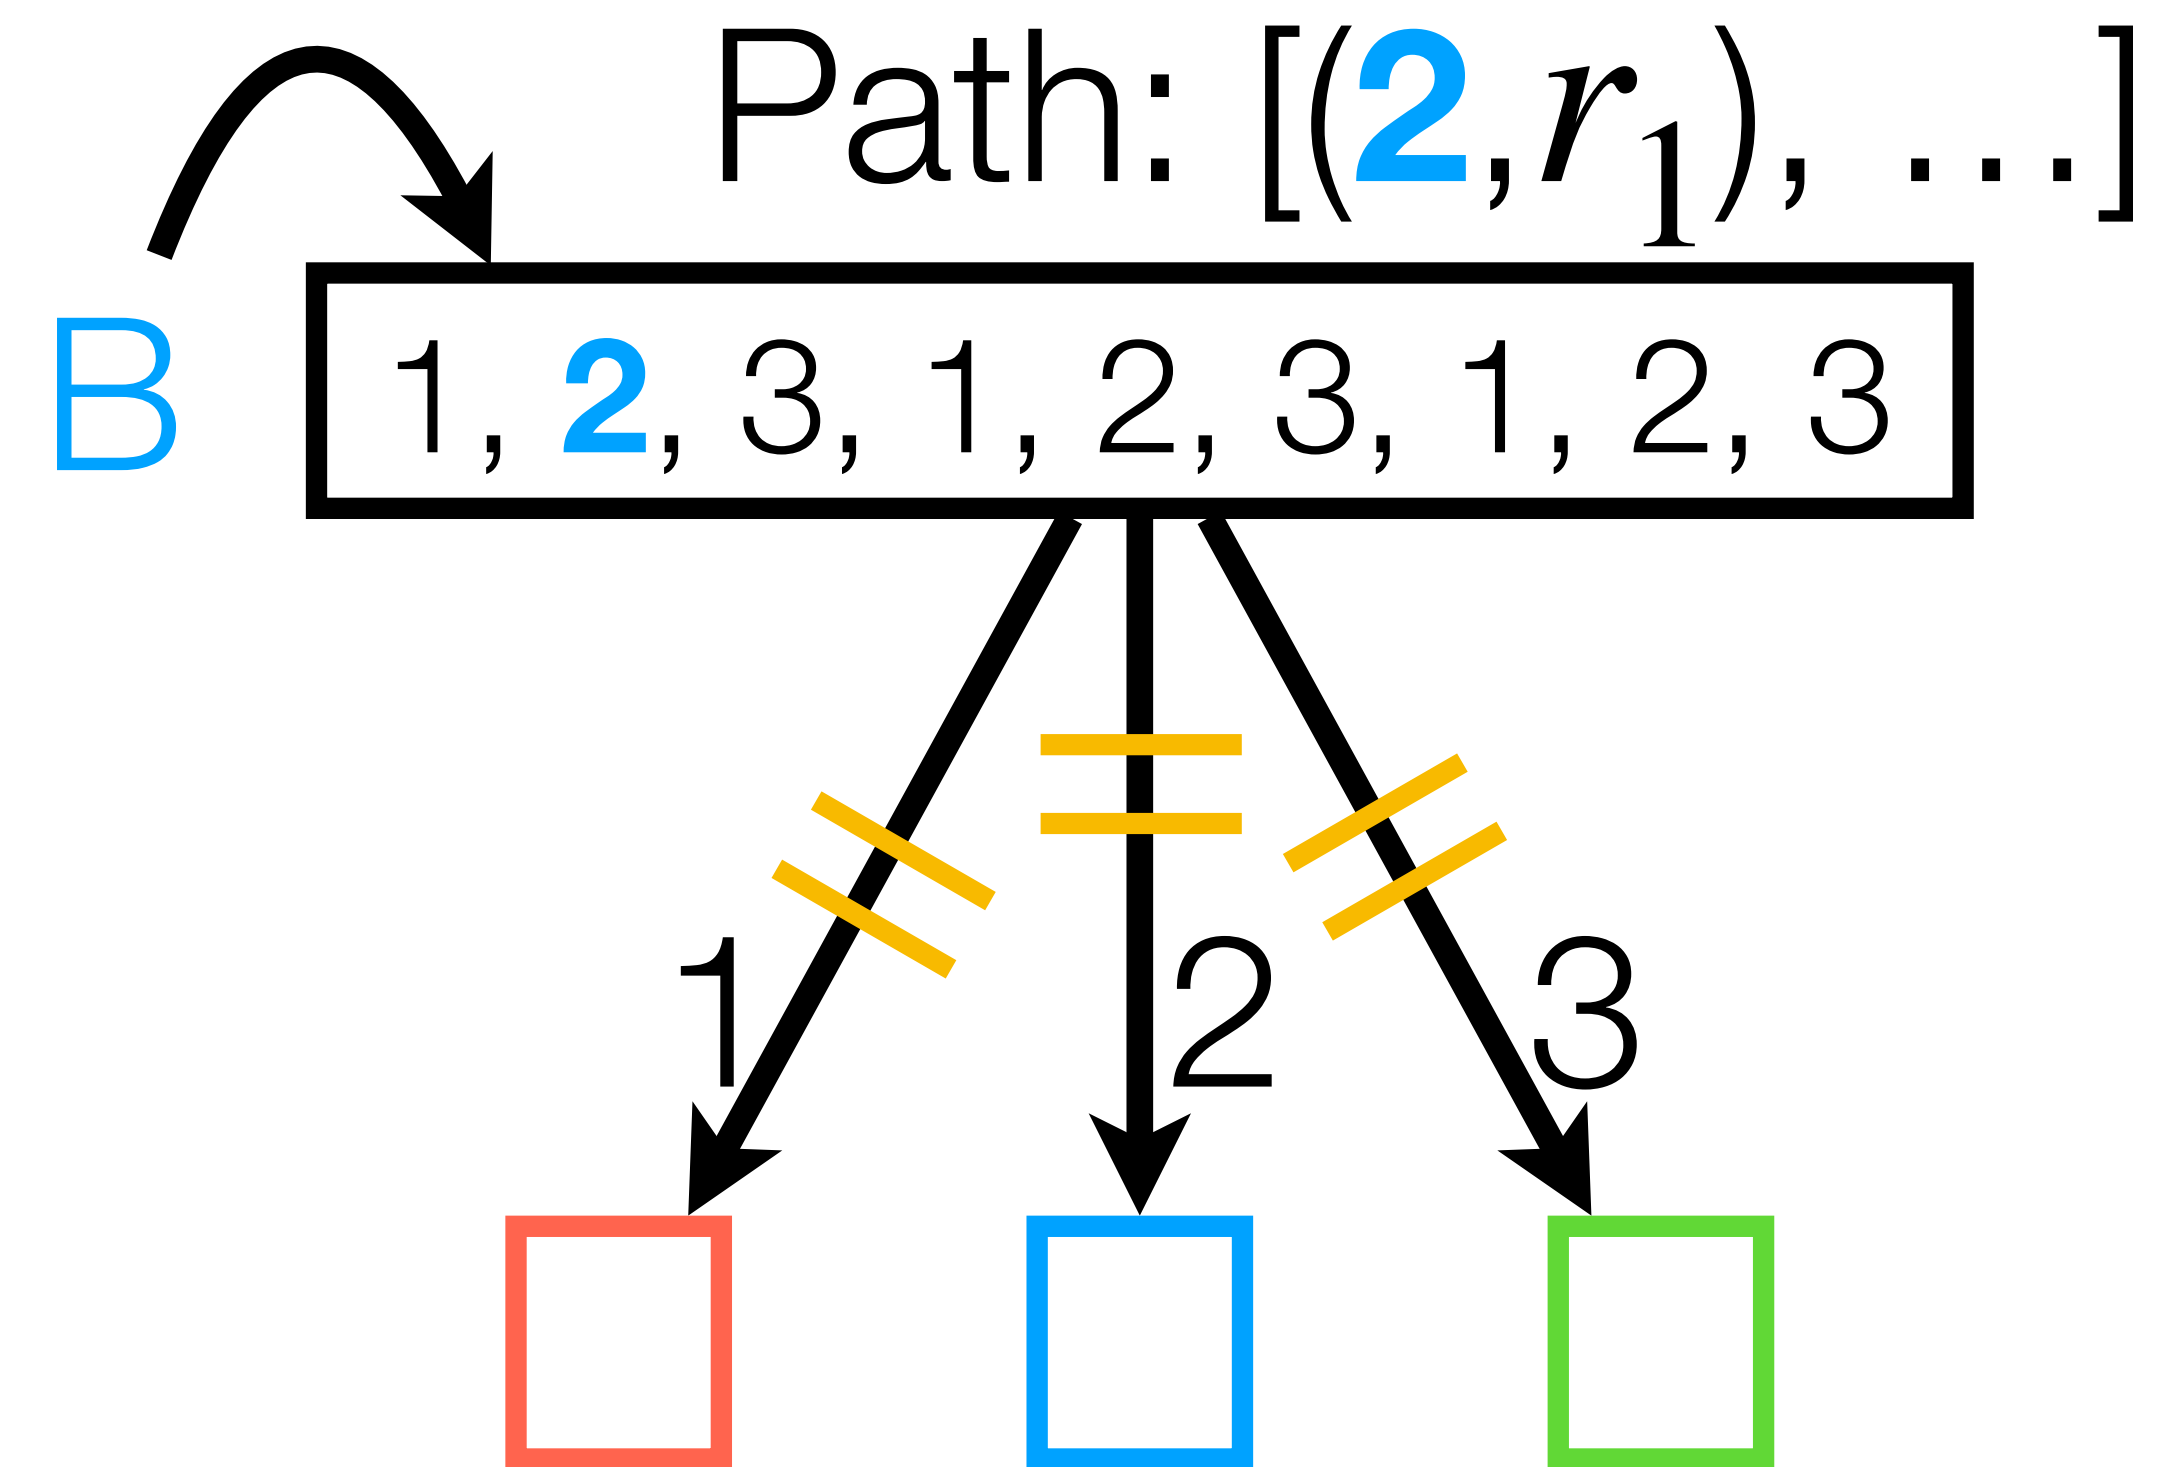
Compiling programs



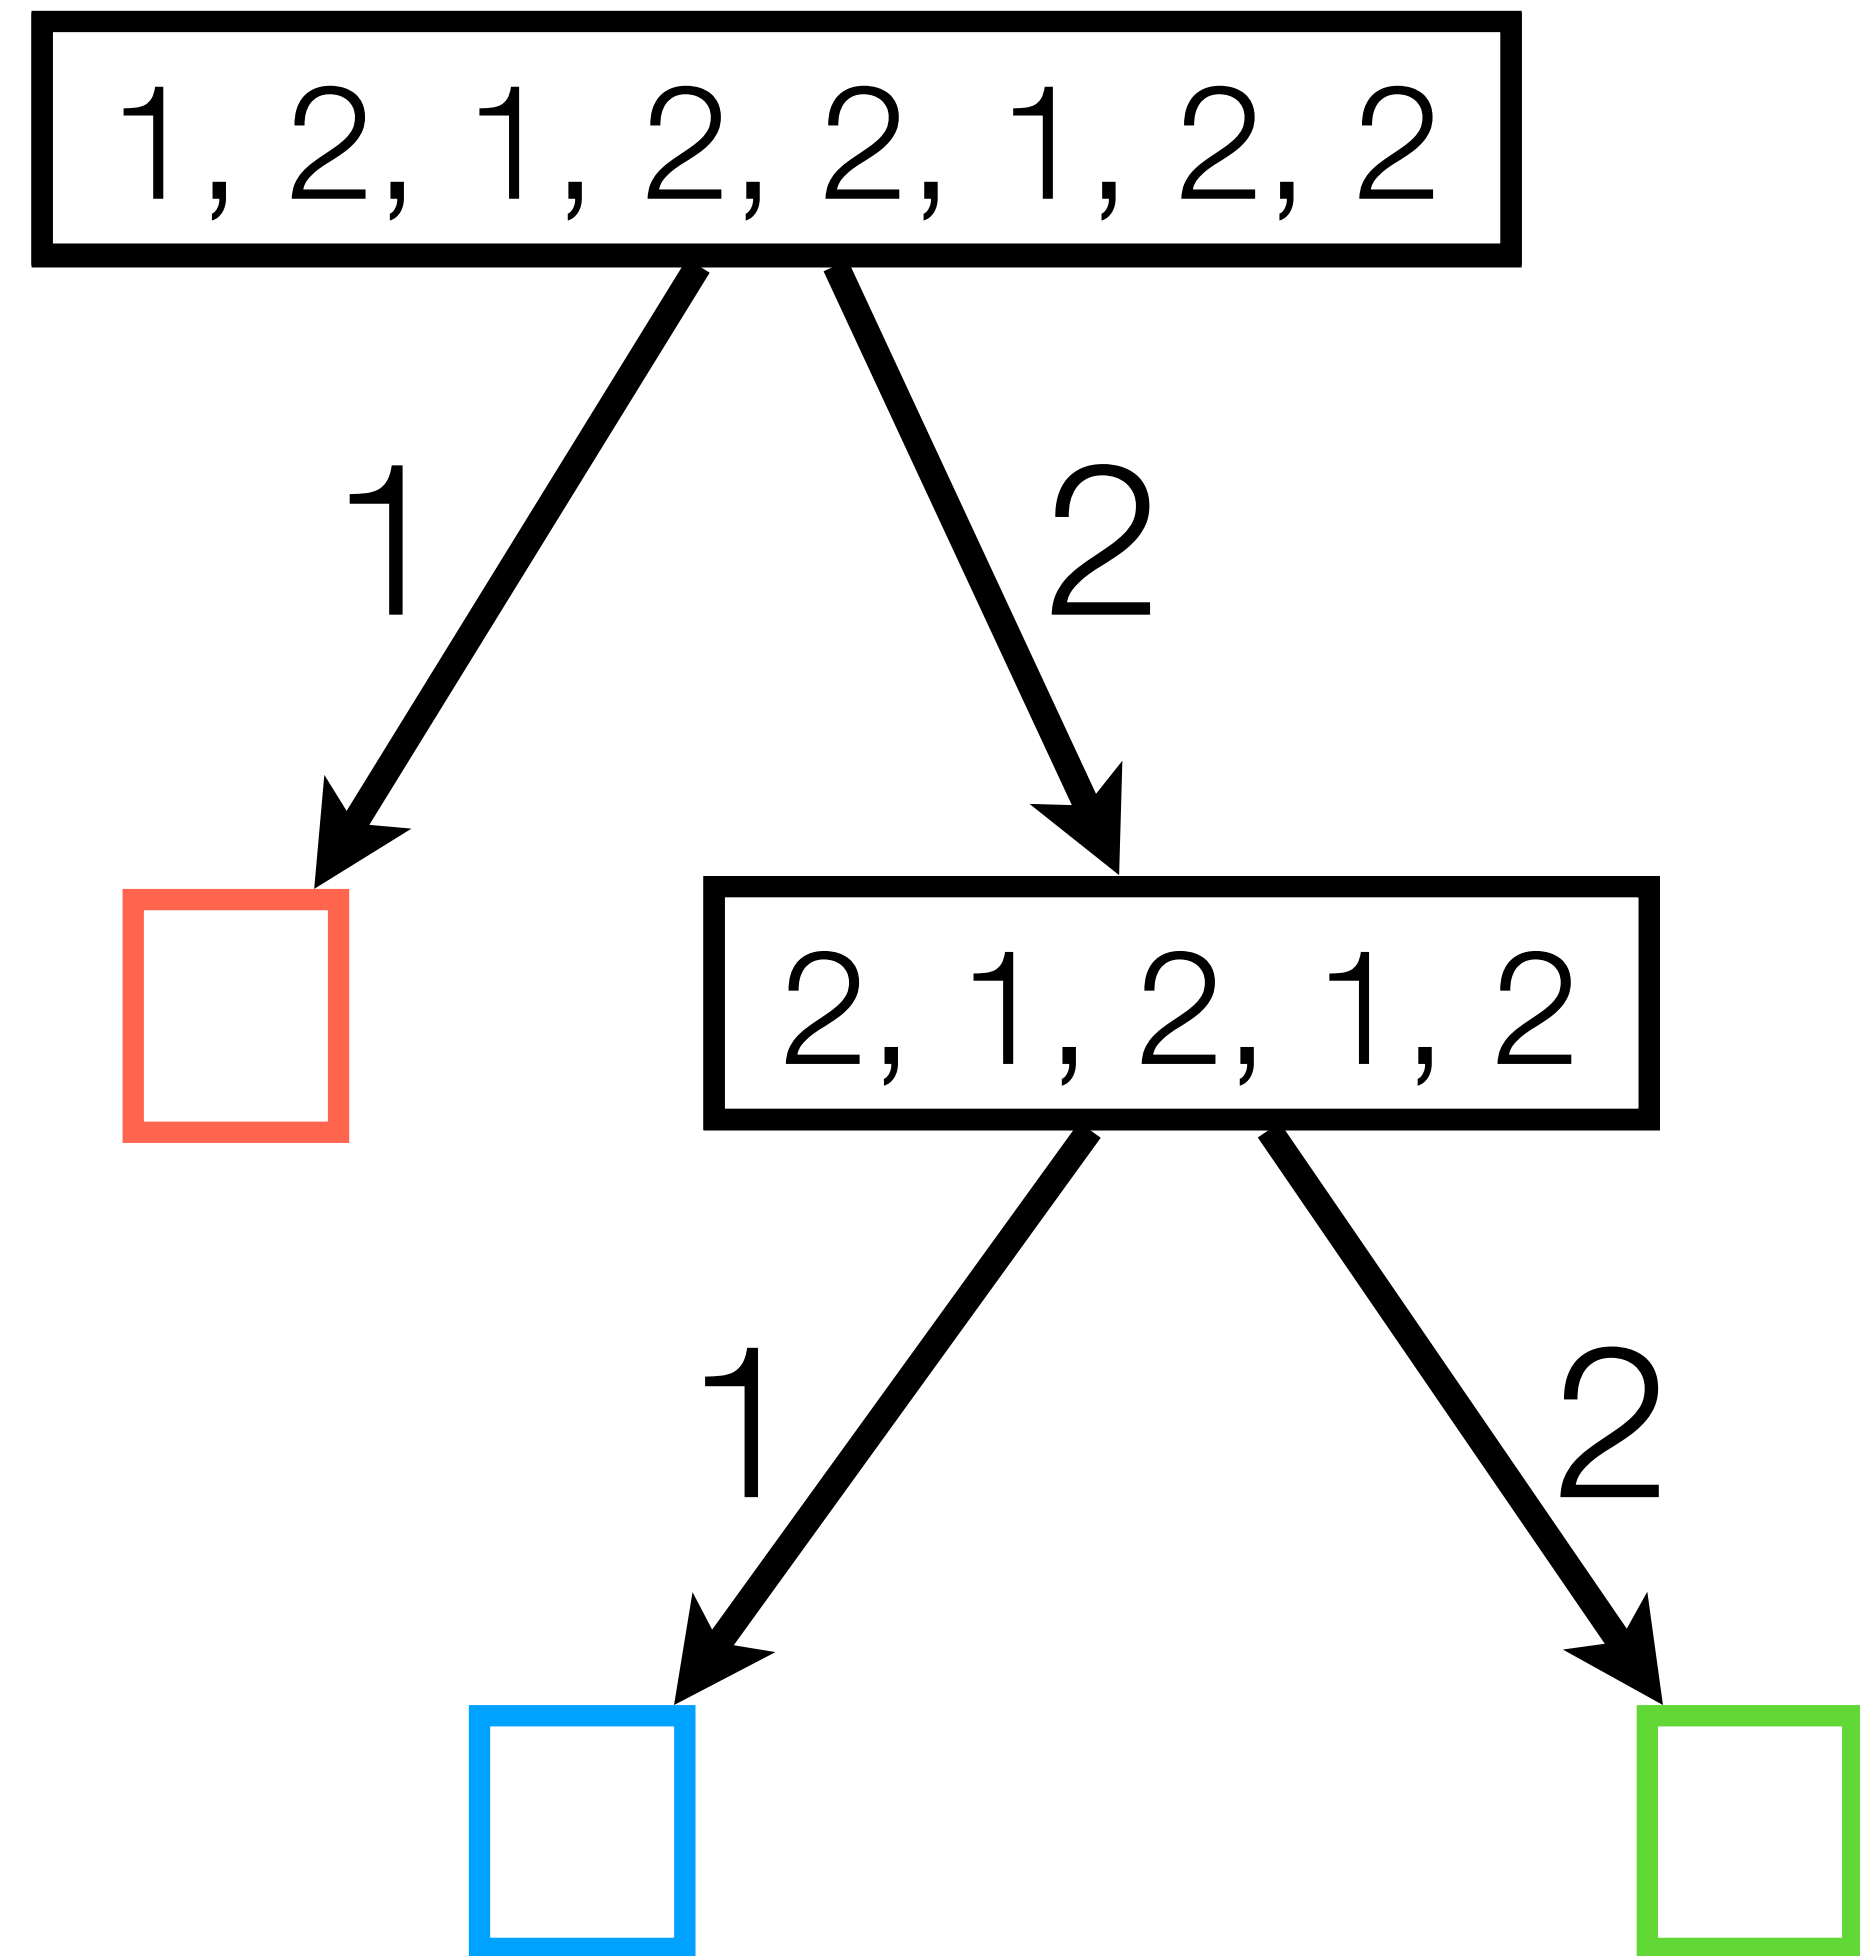
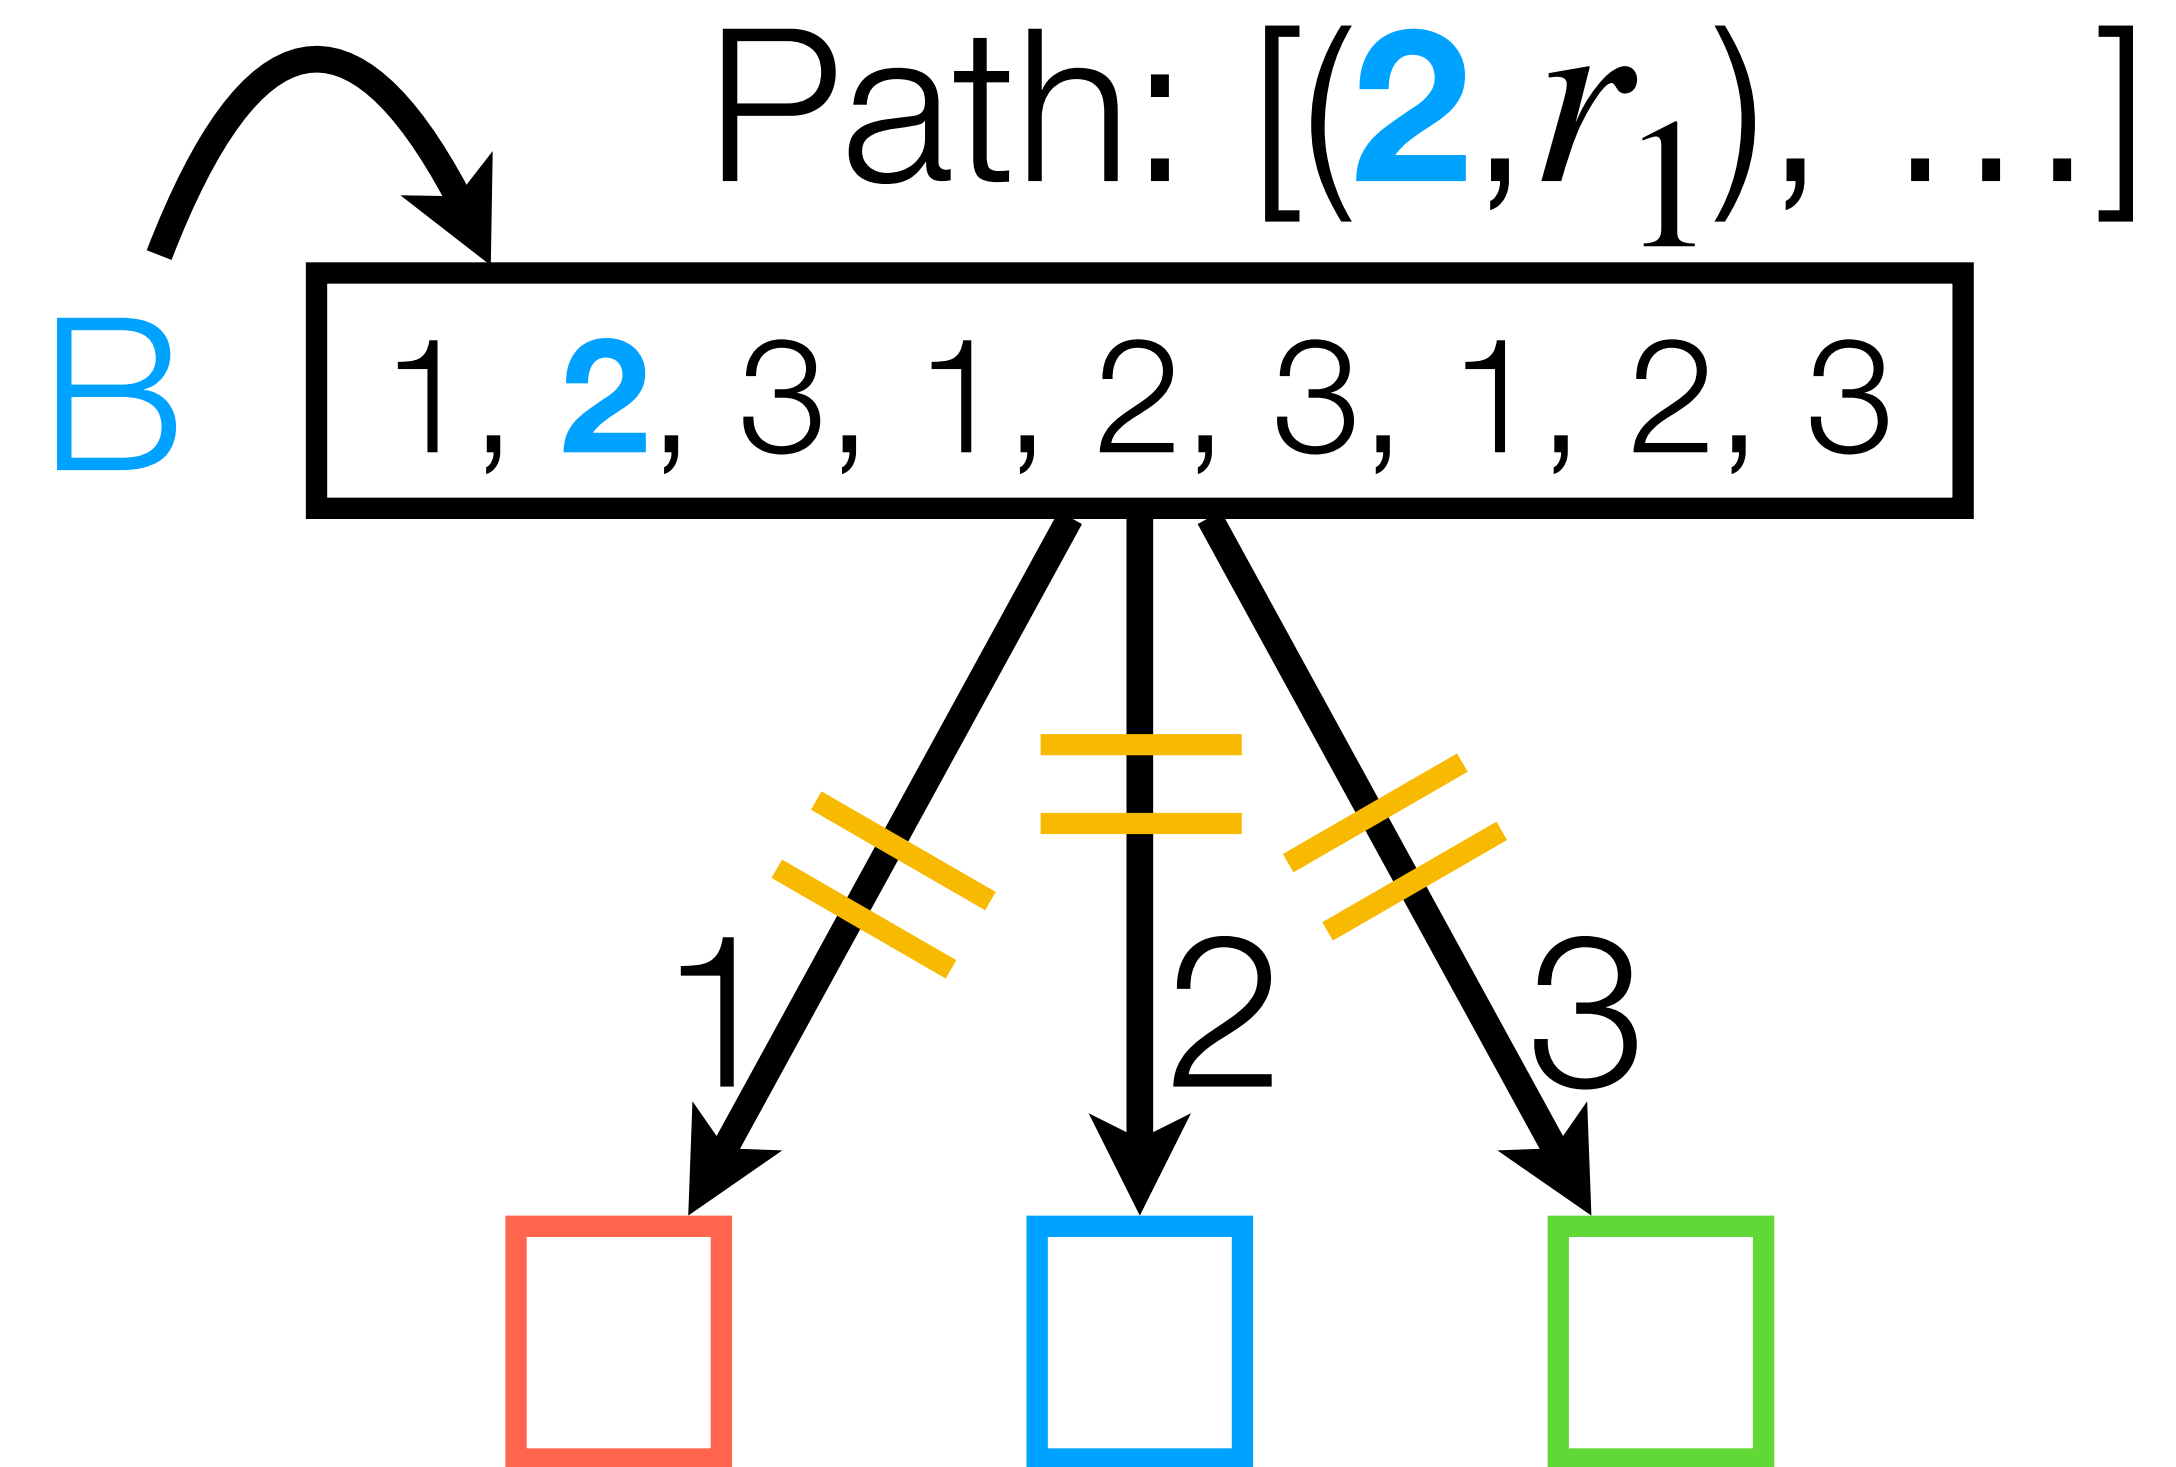
Compiling programs



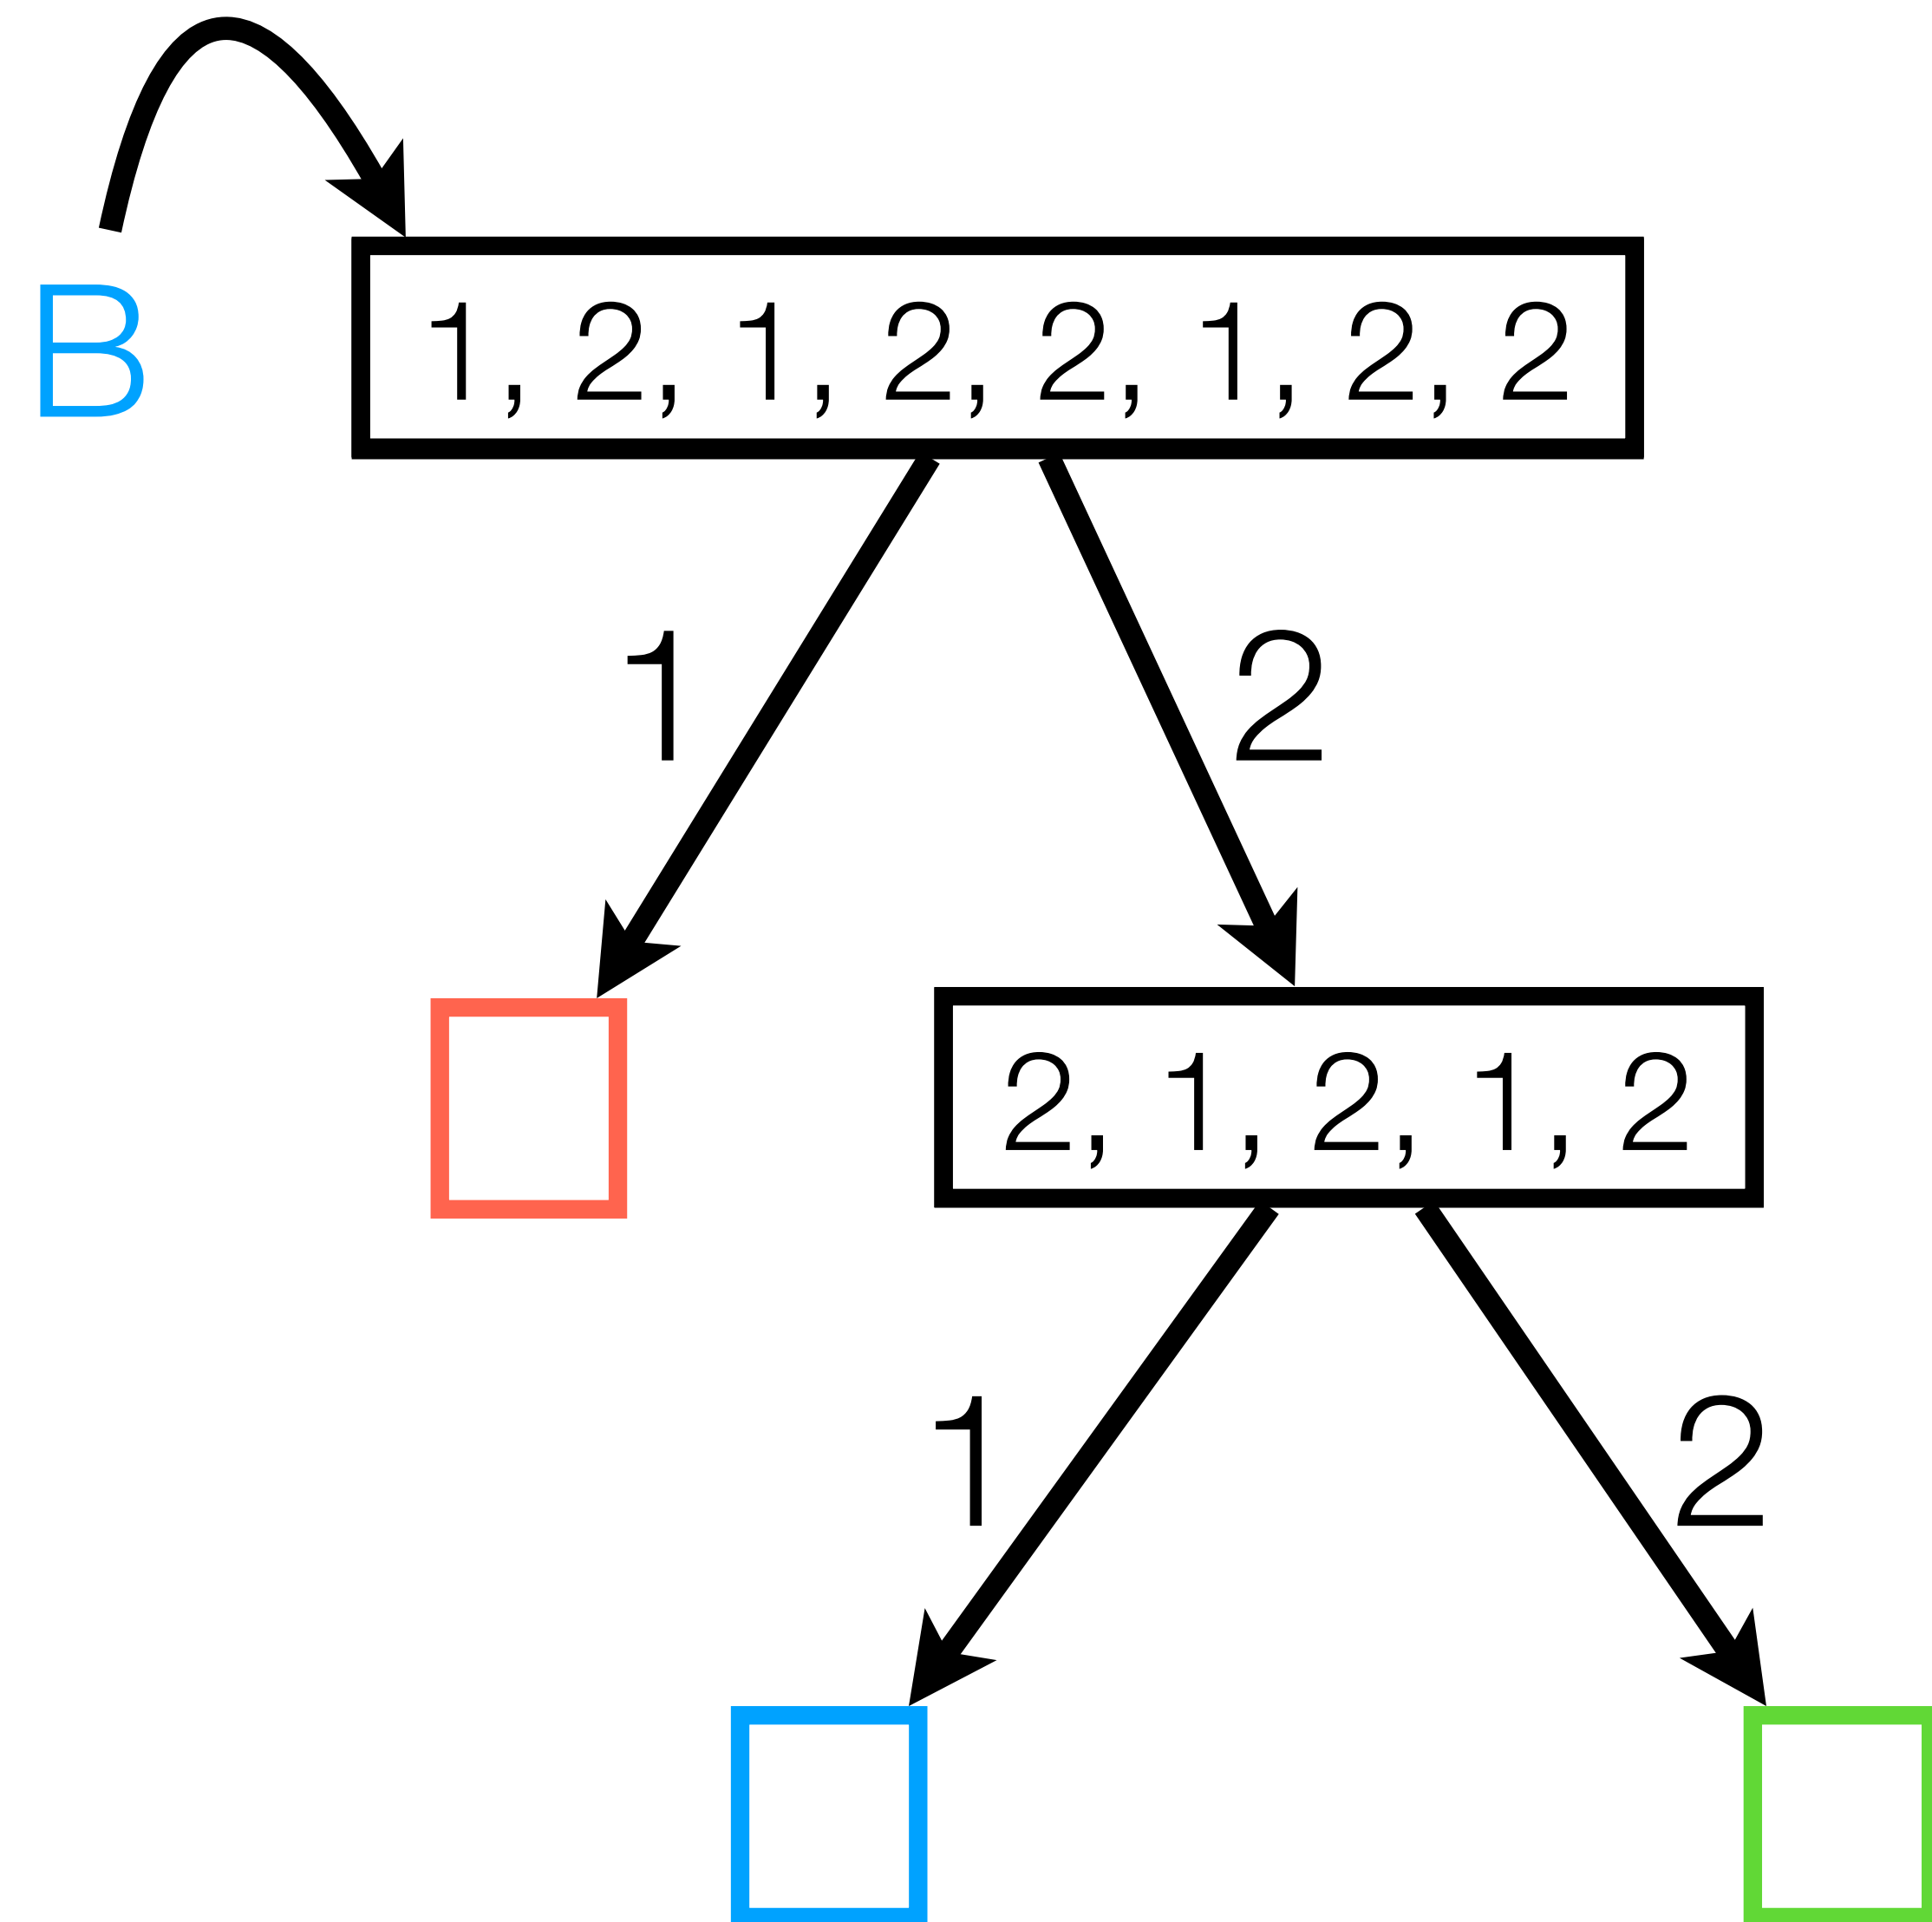
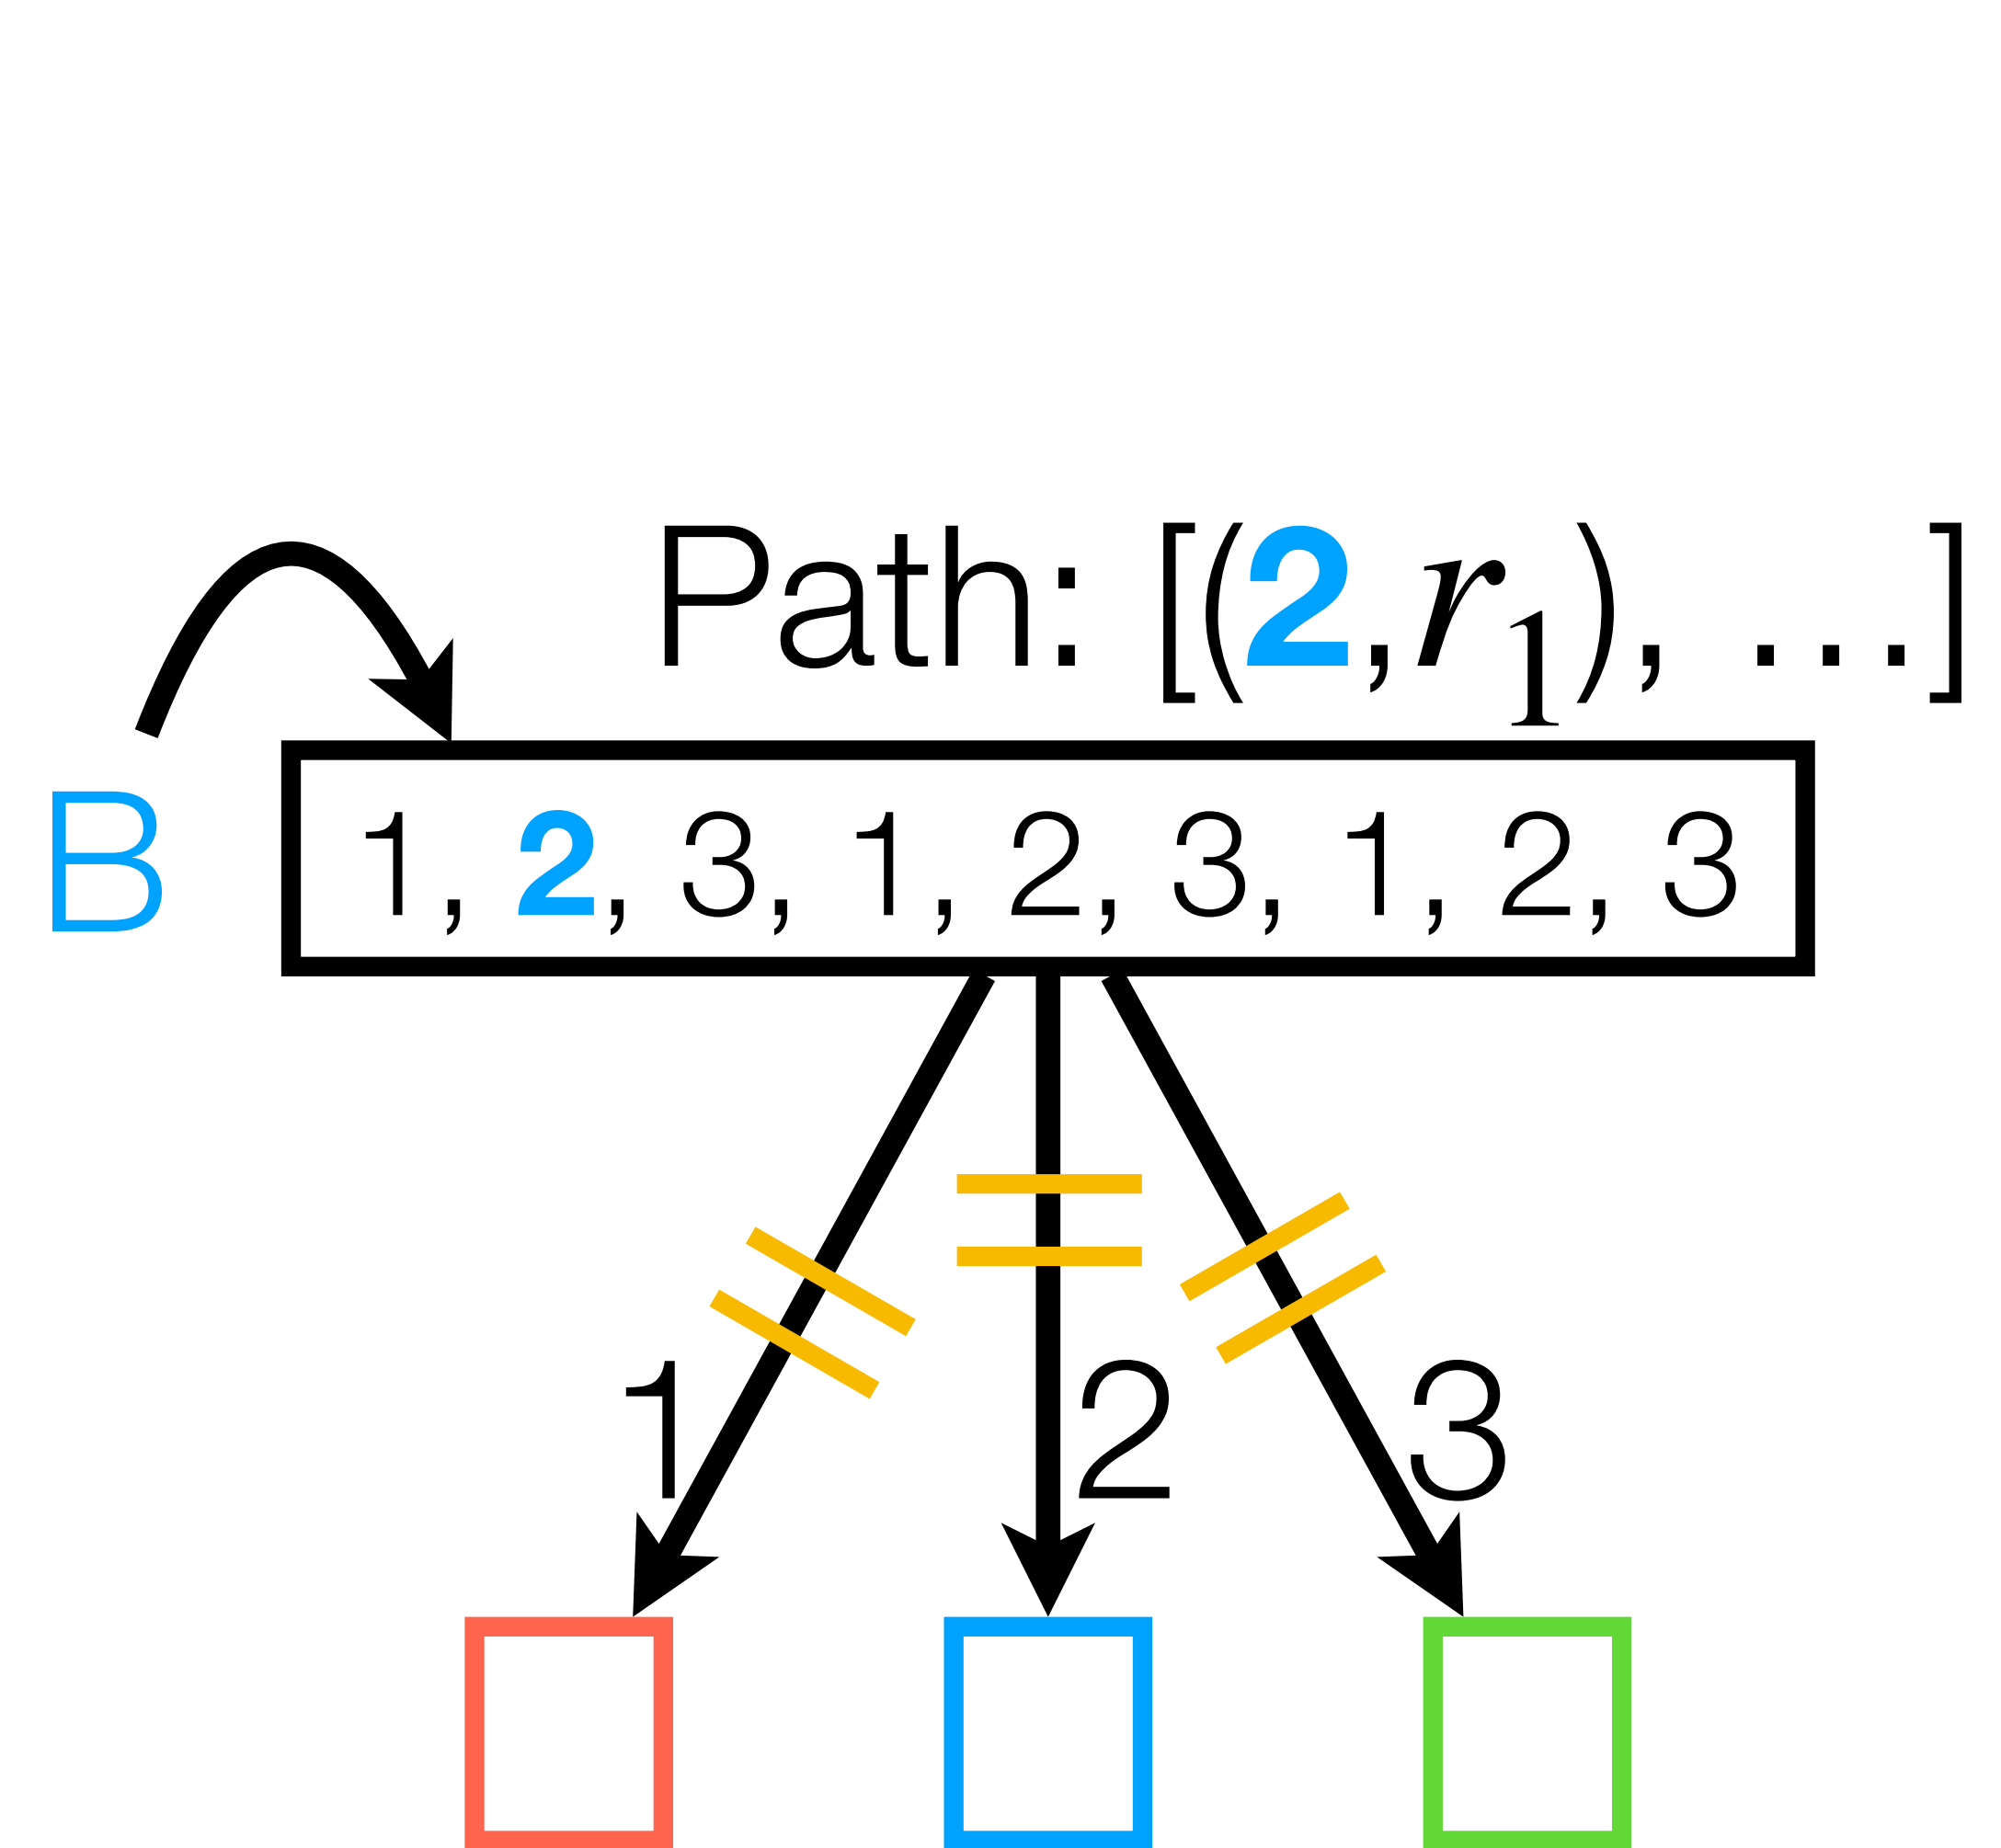
Compiling programs



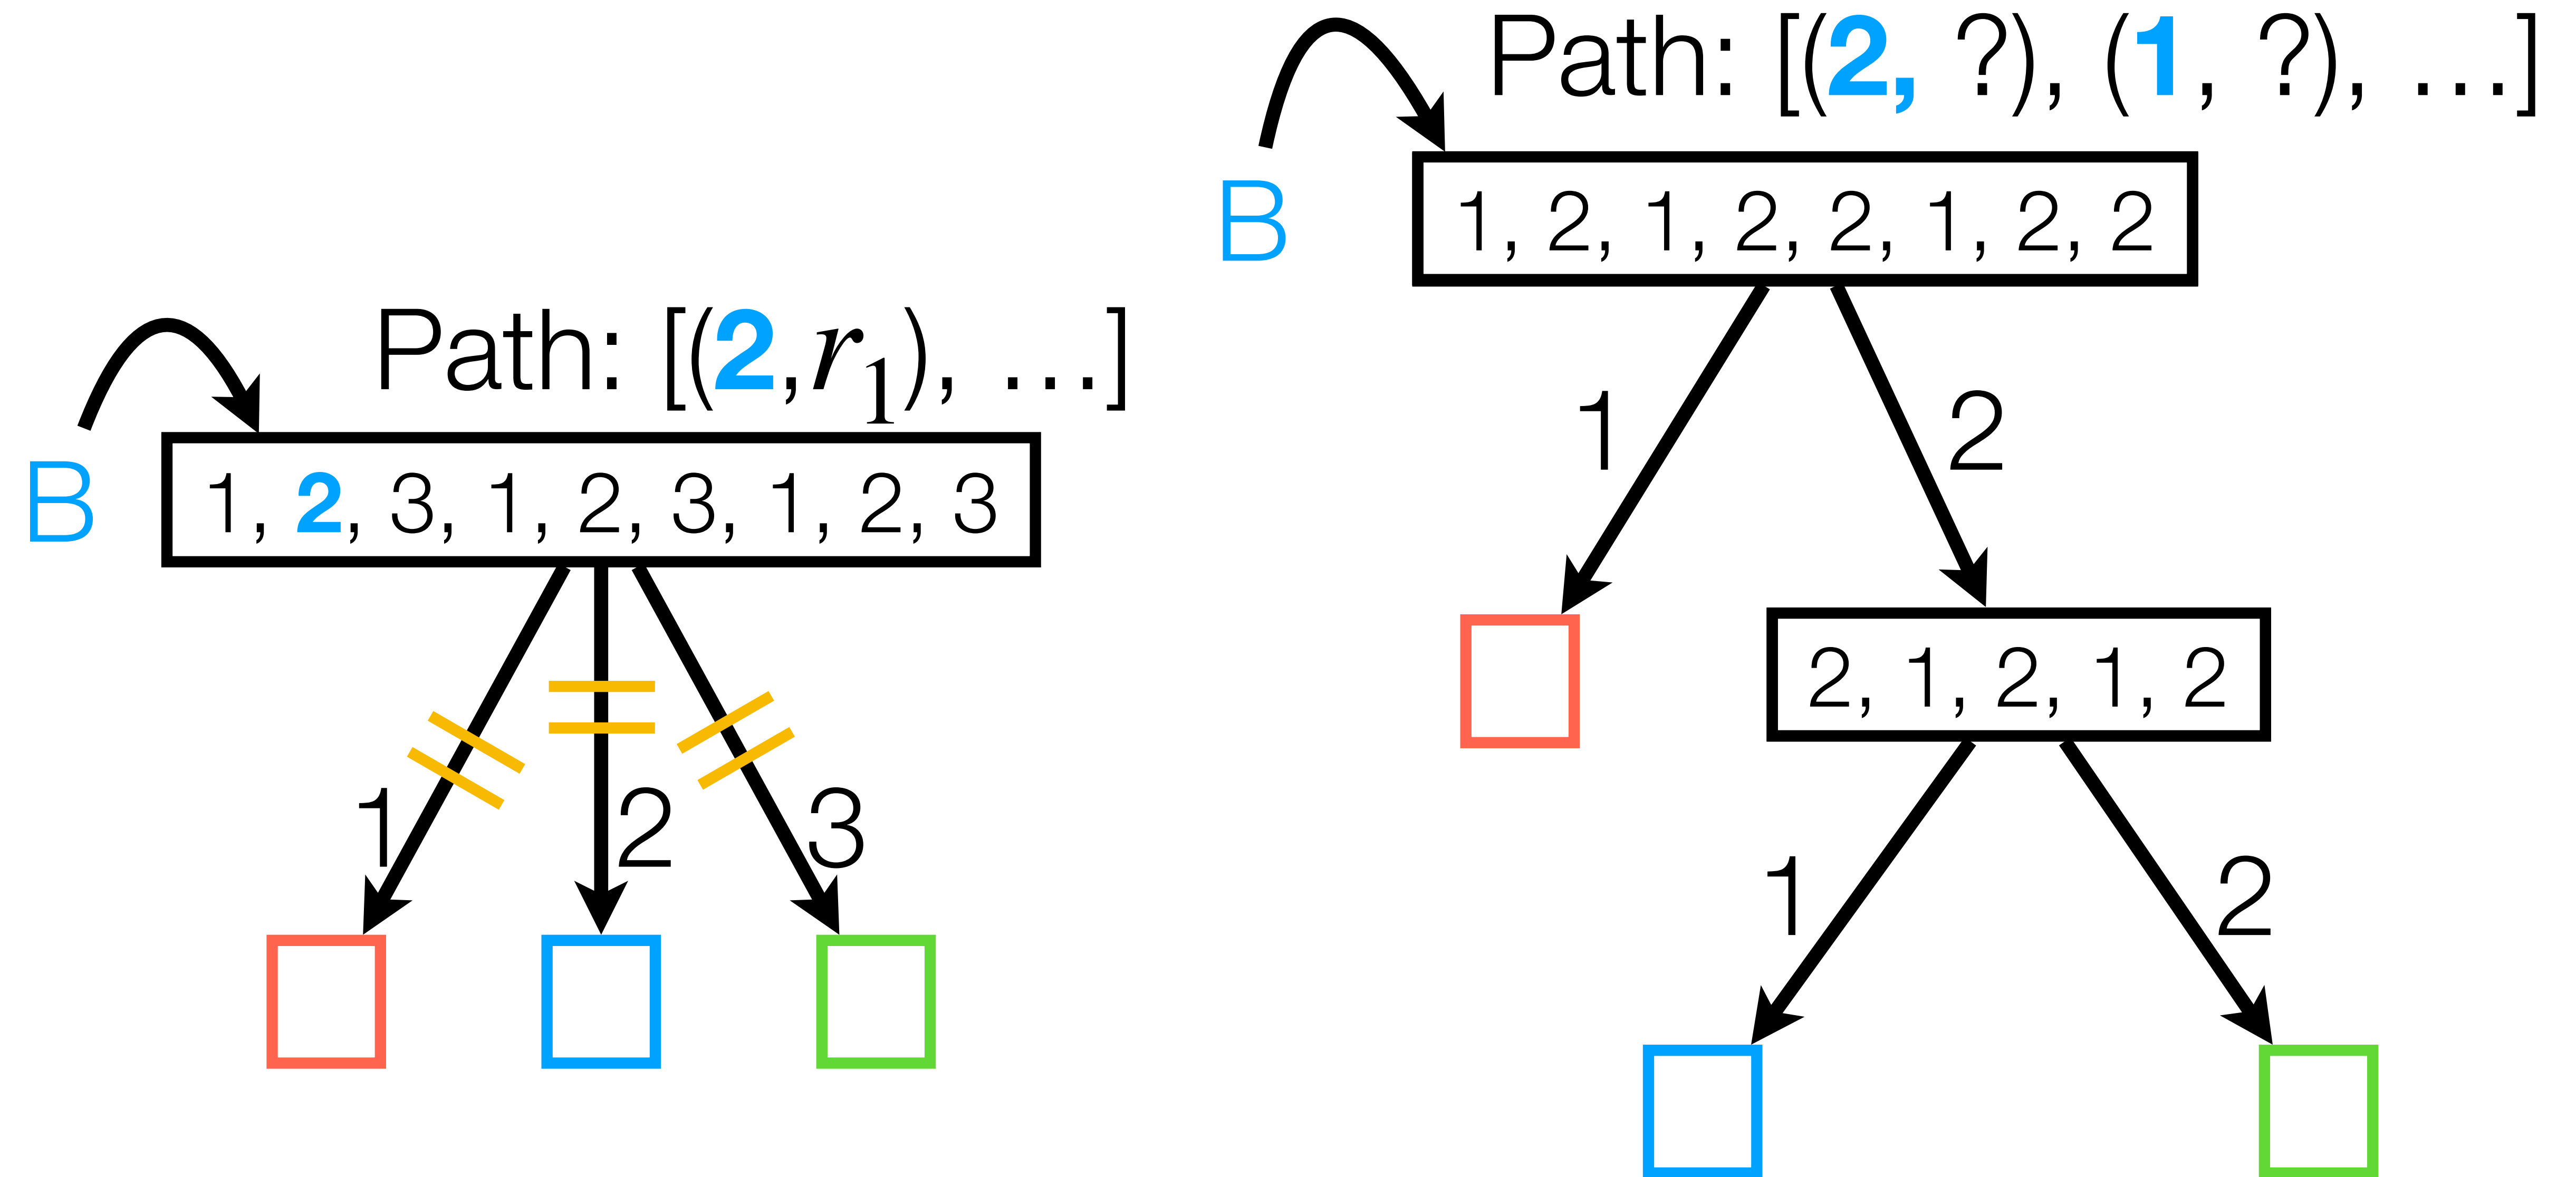
Compiling programs



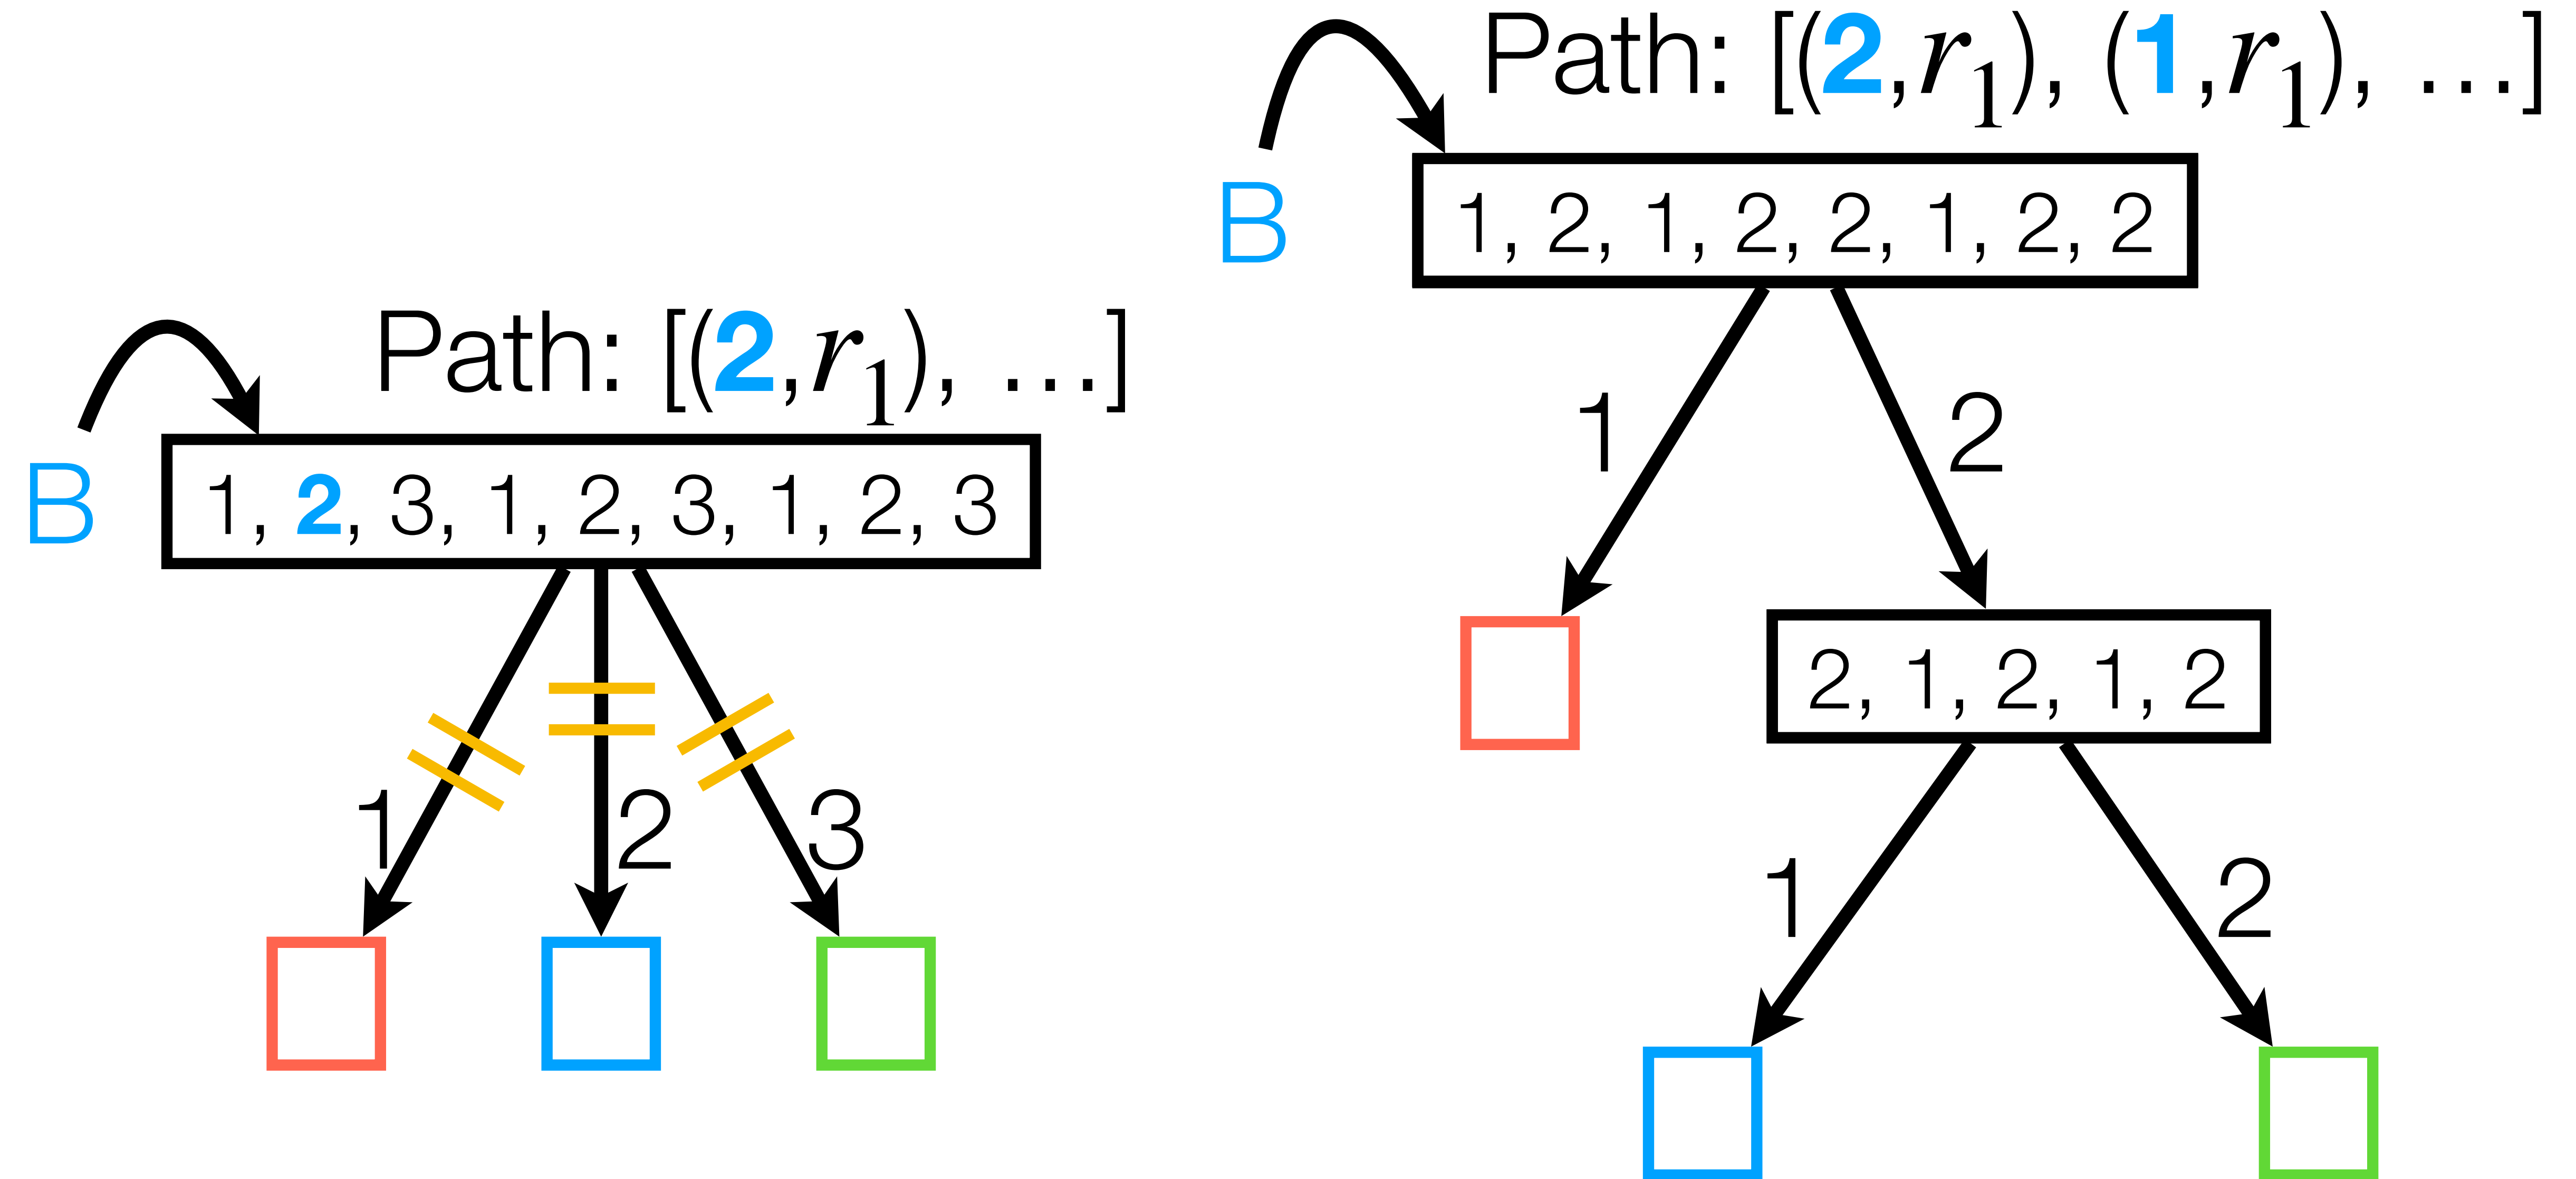
Compiling programs



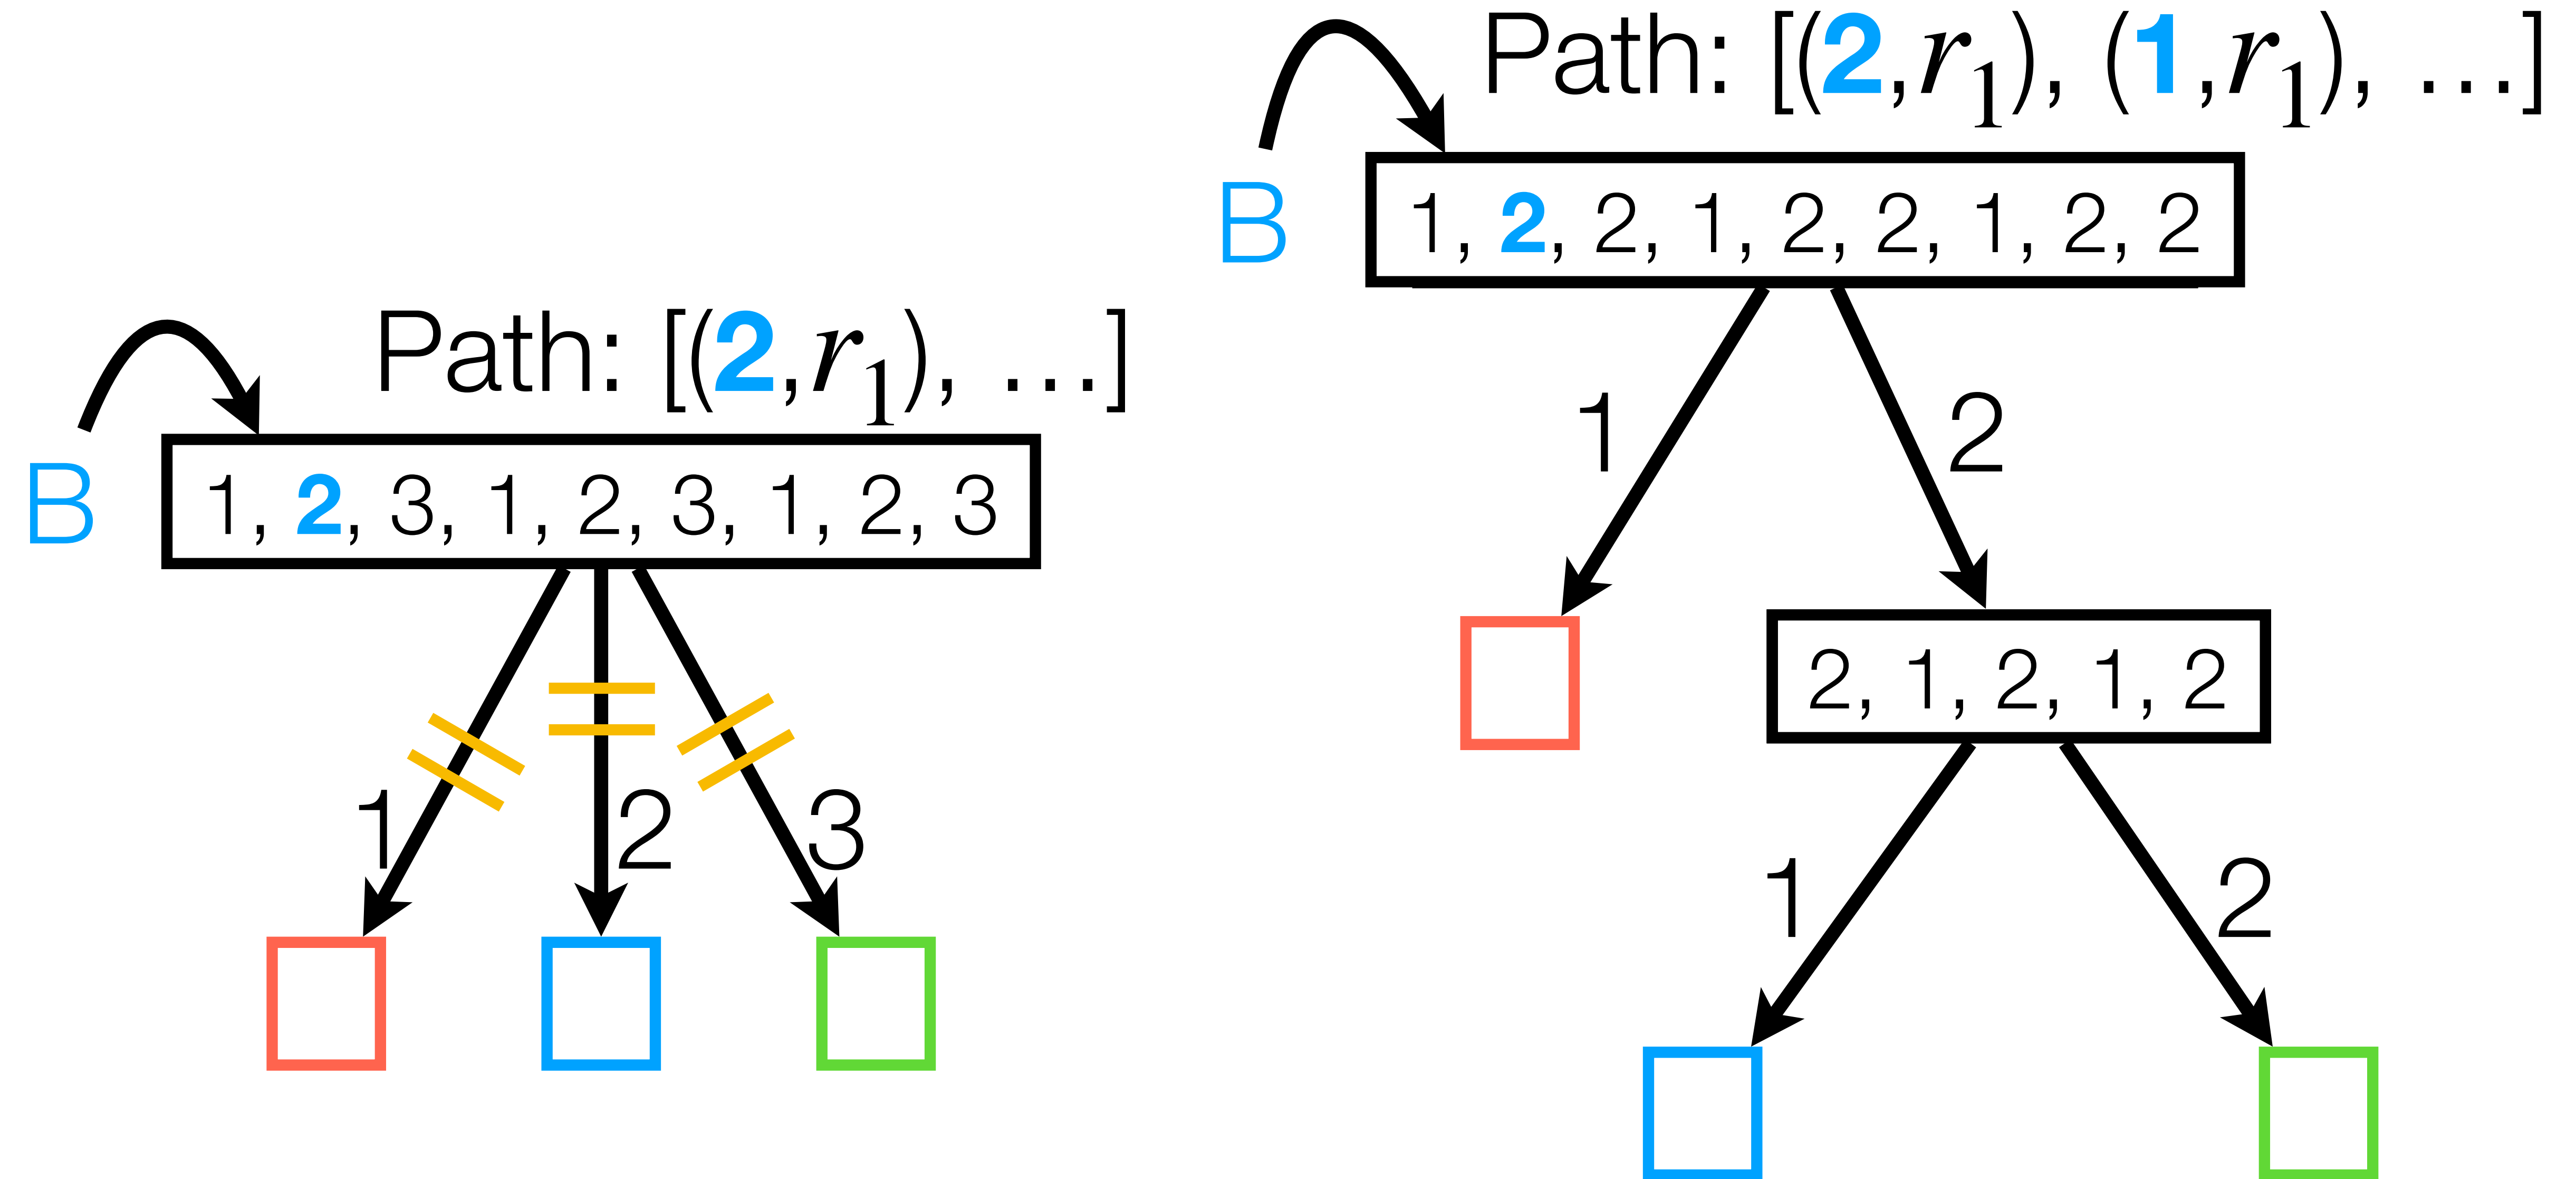
Compiling programs



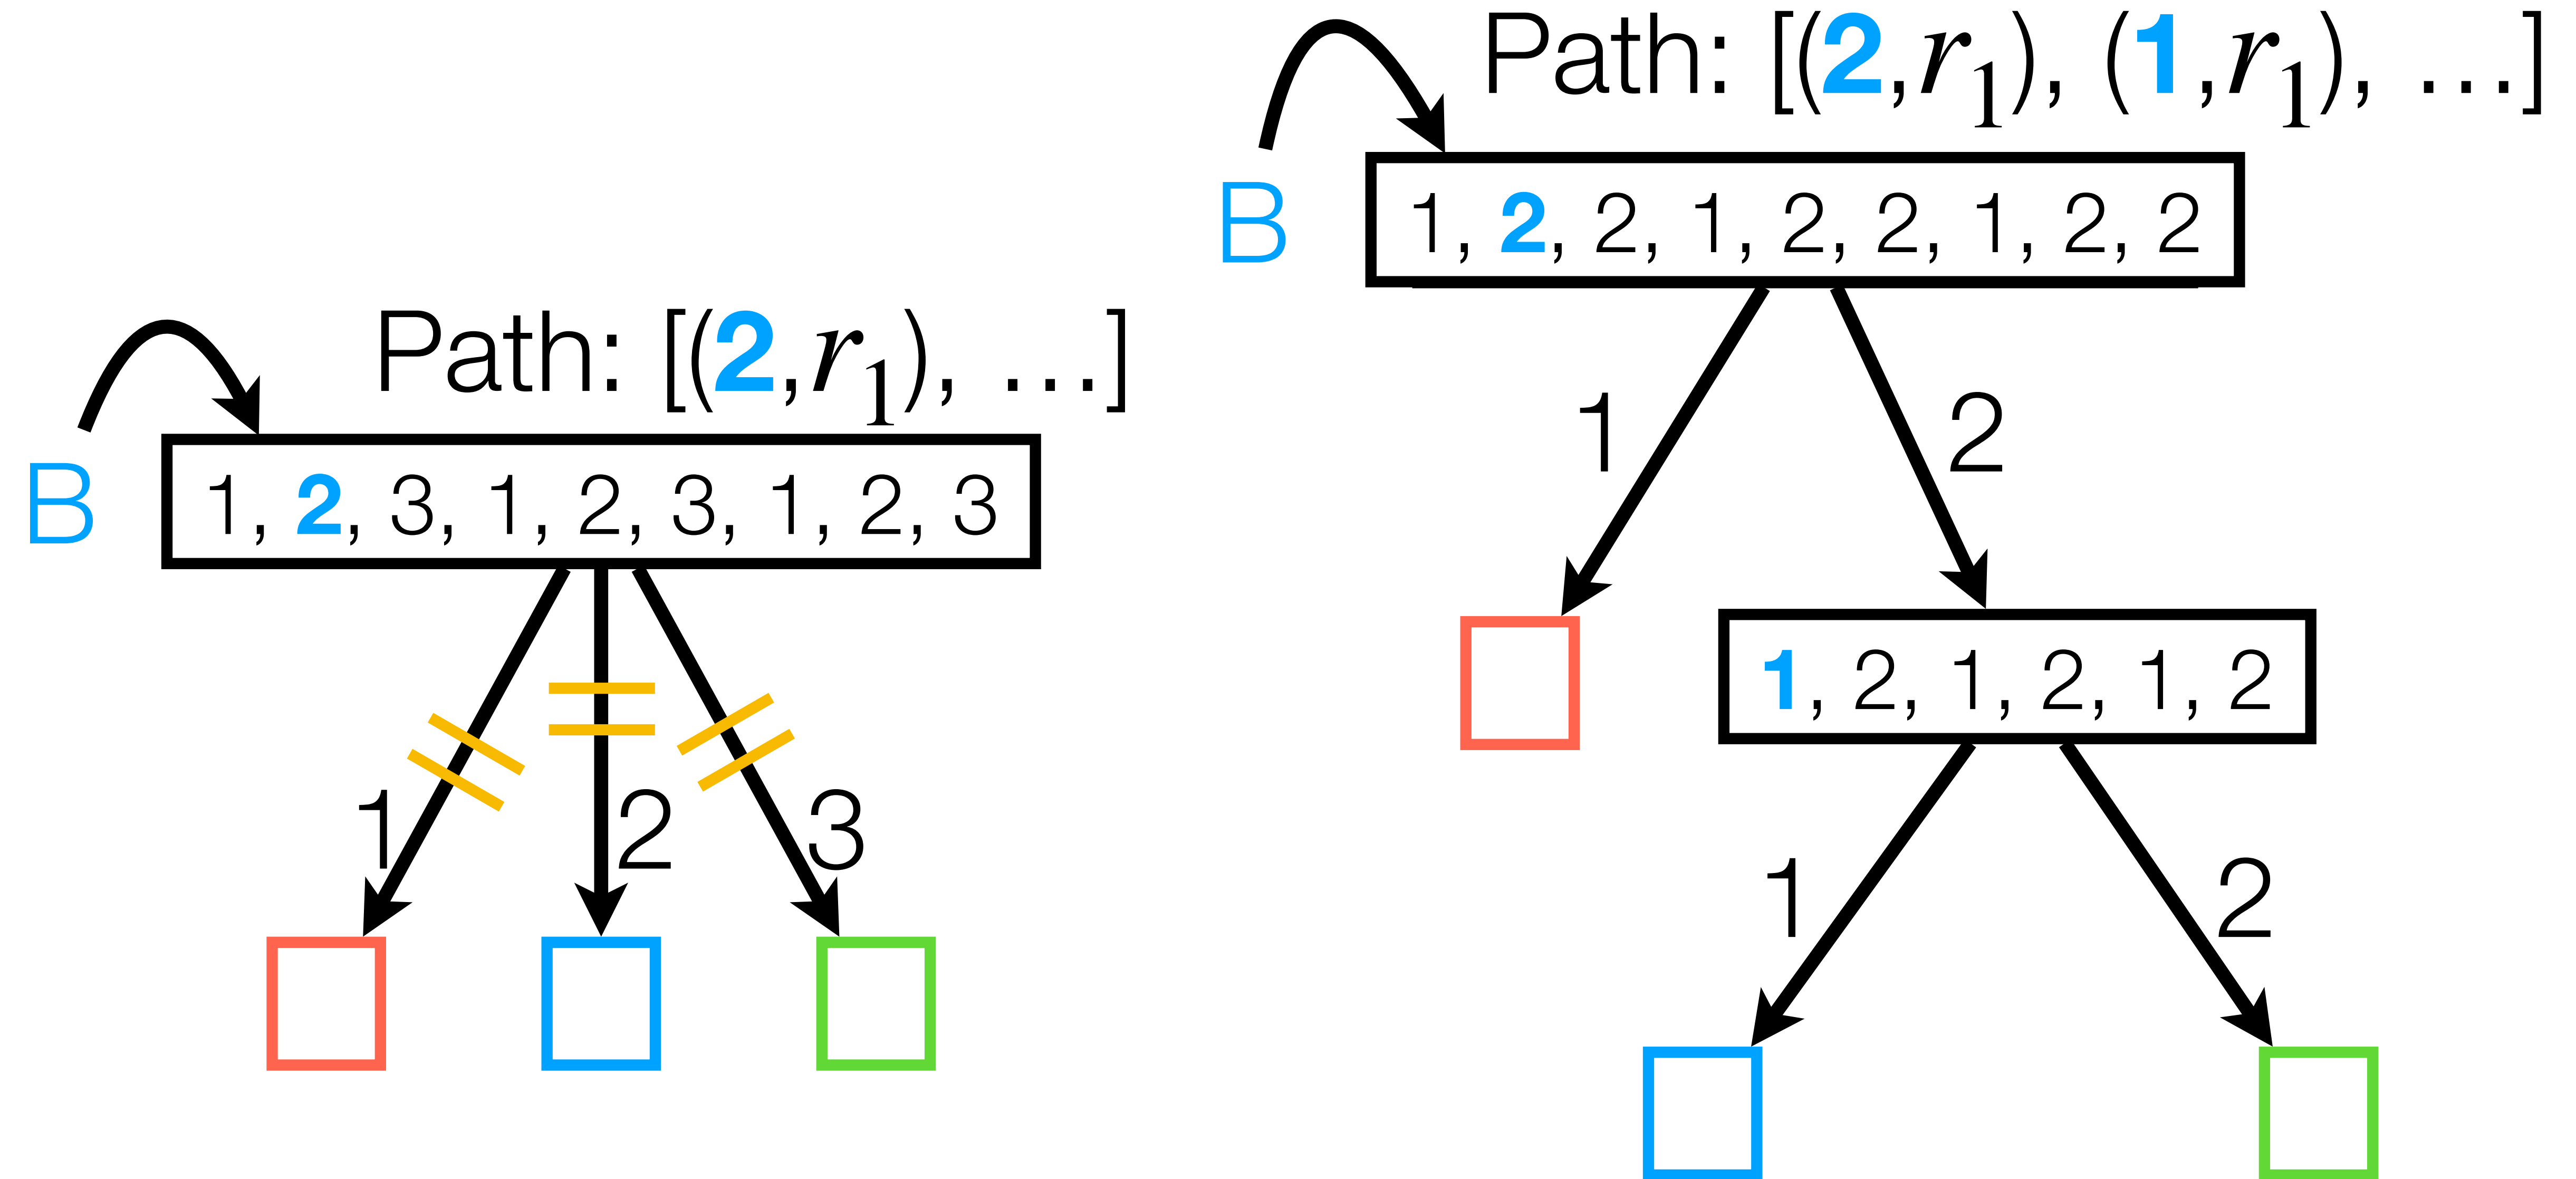
Compiling programs



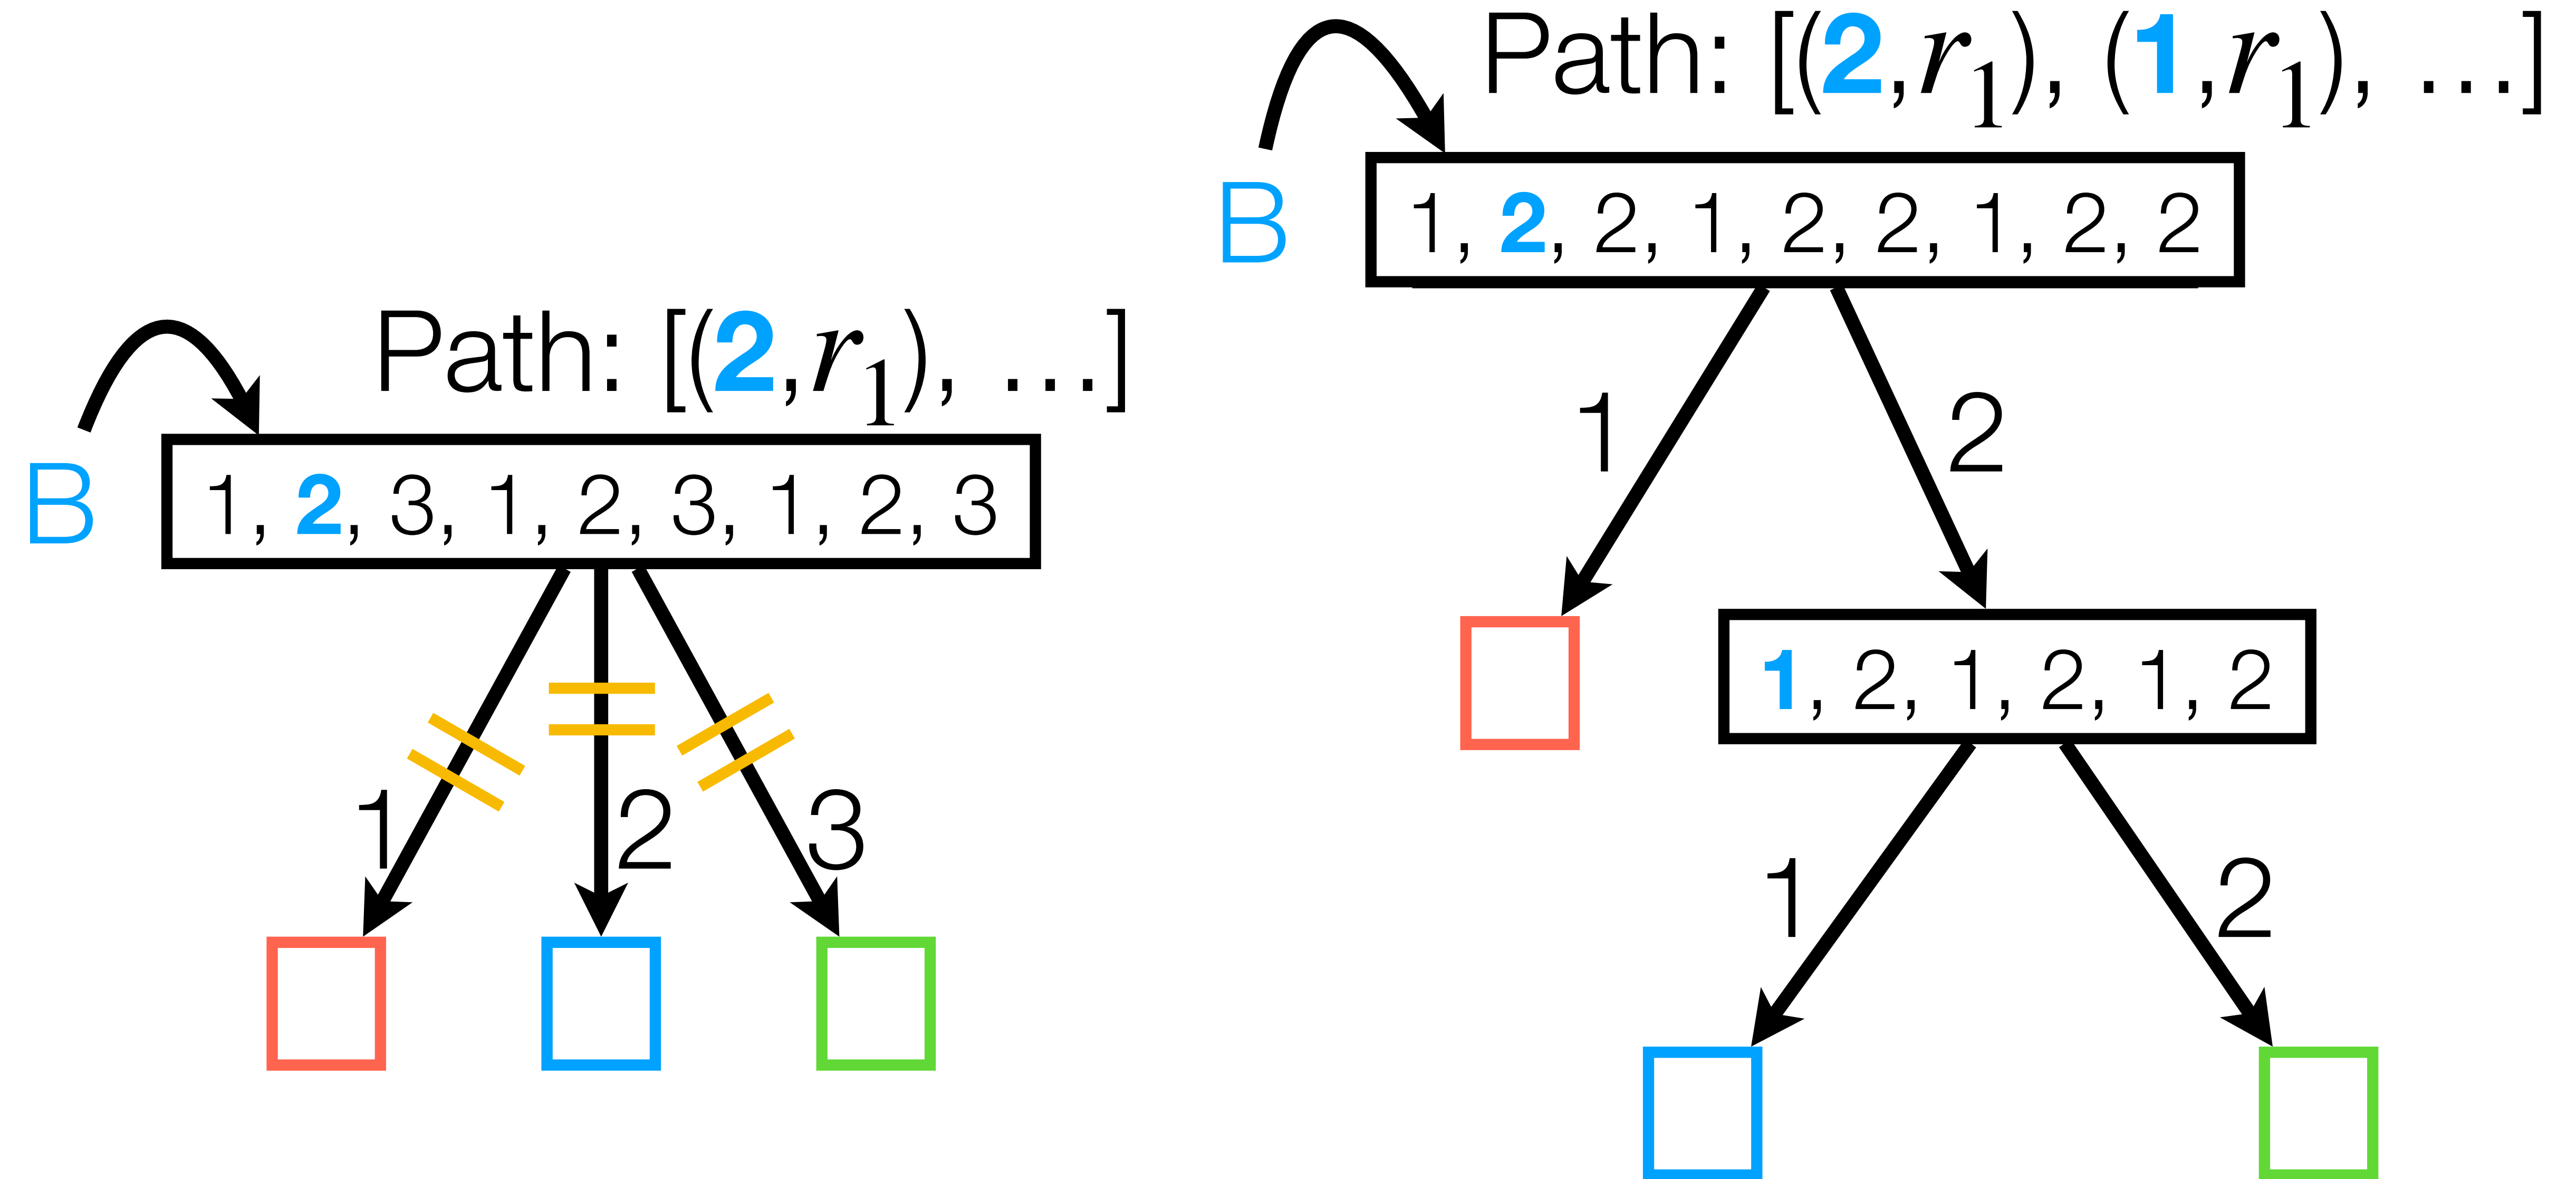
Compiling programs



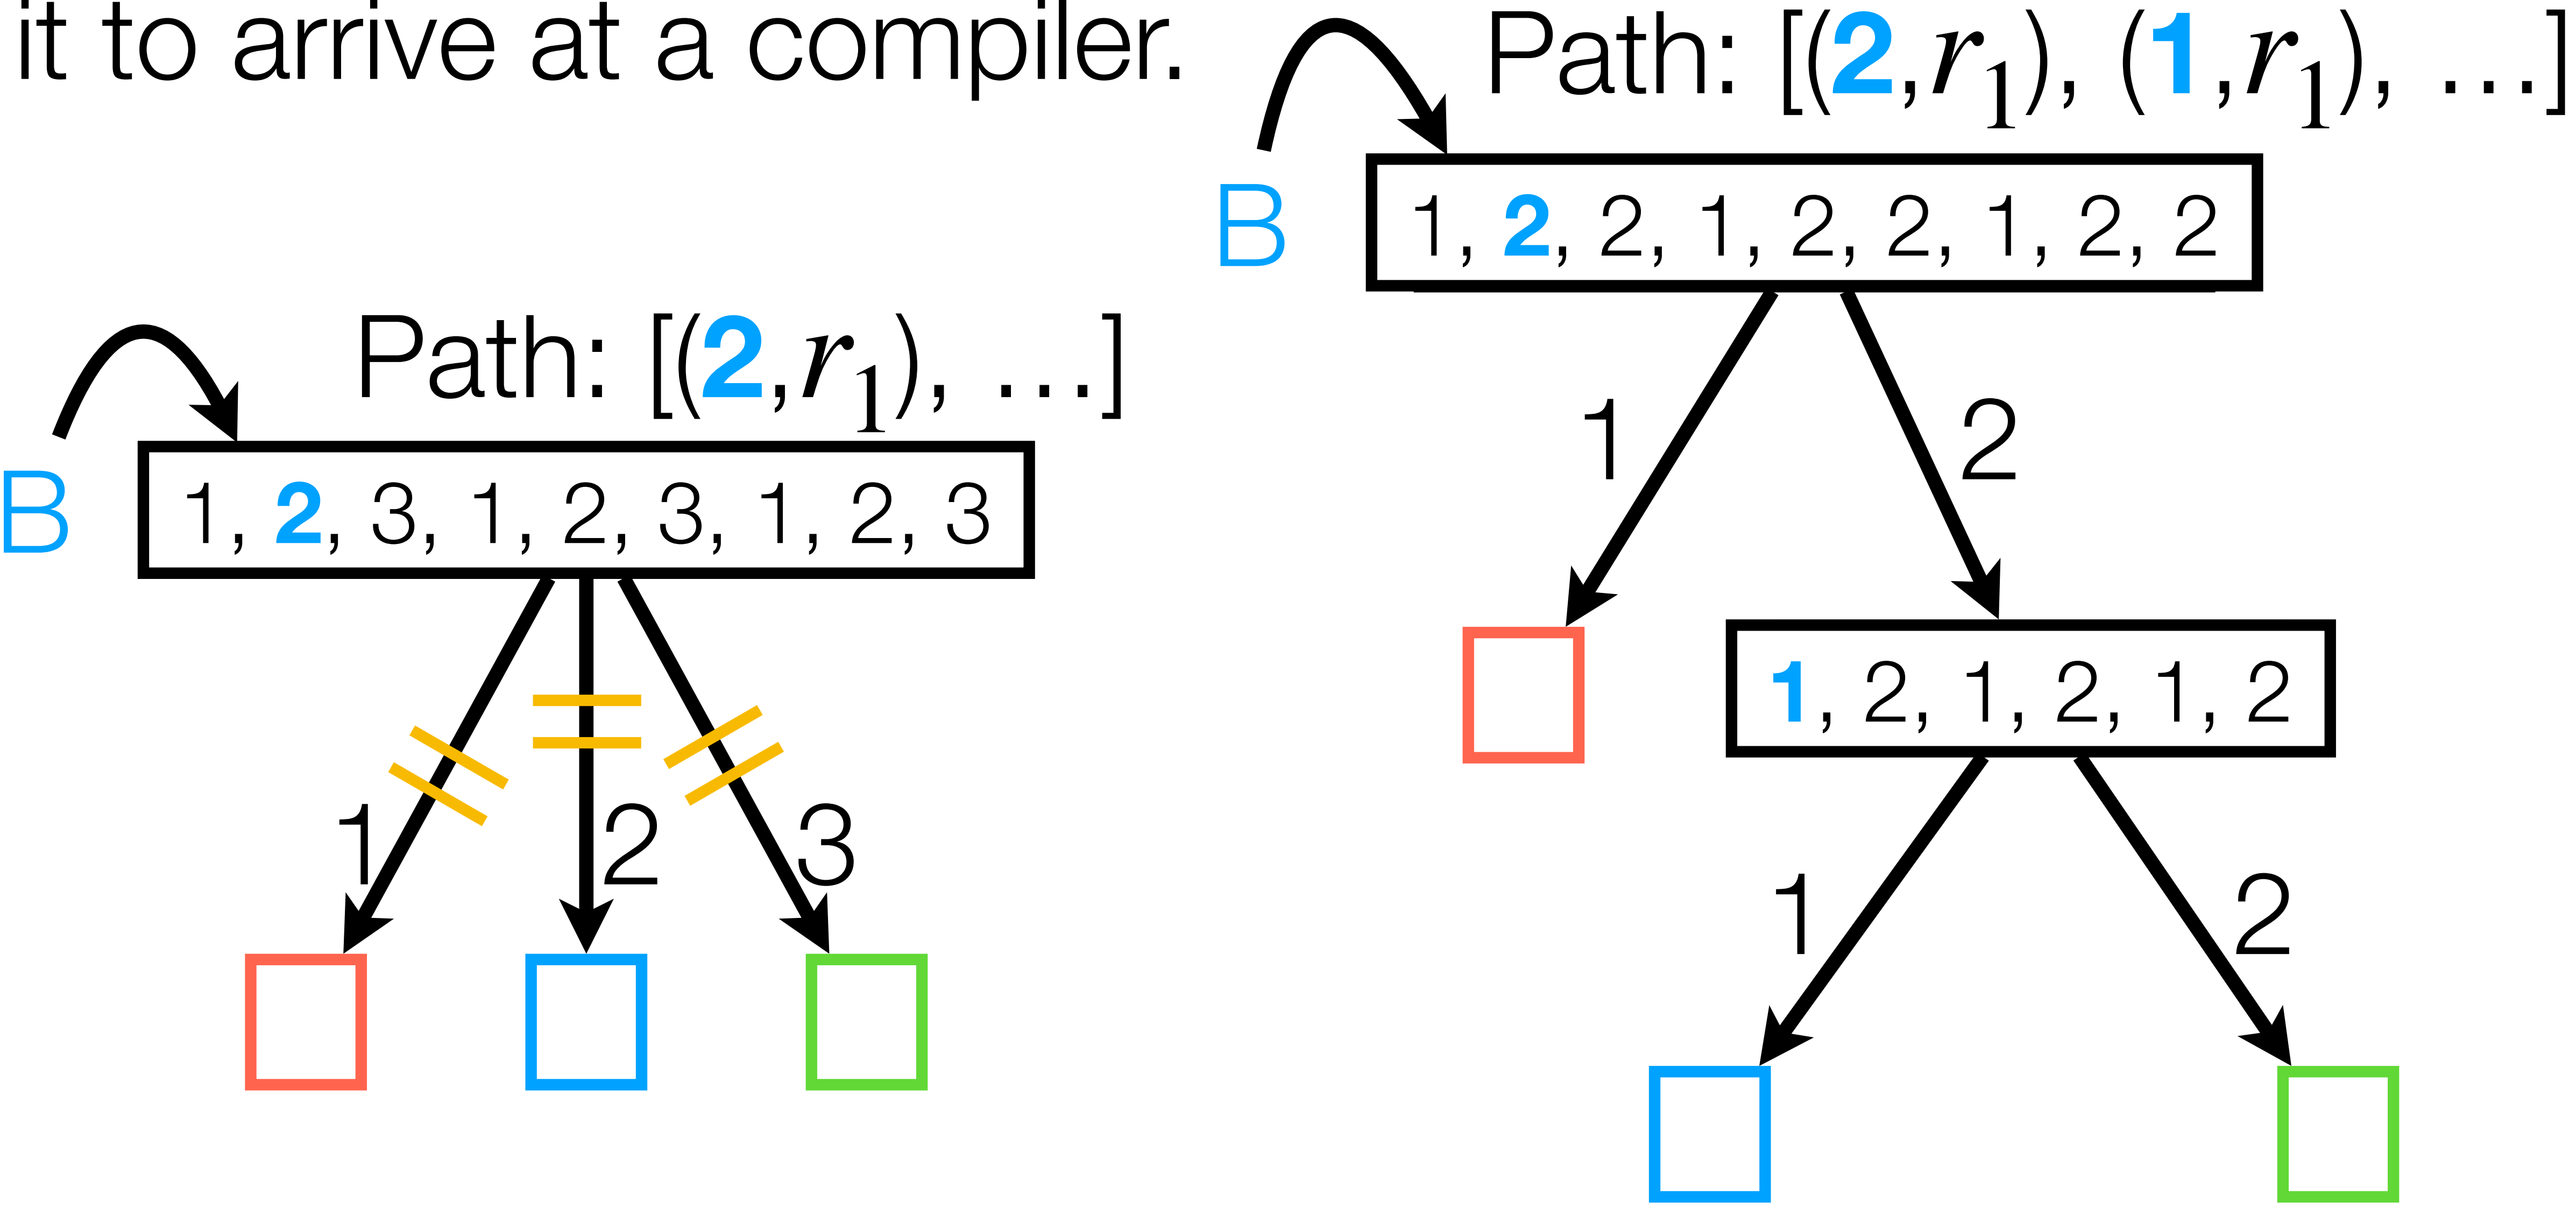
Compiling programs



Compiling programs



Given an embedding, we *lift* it to arrive at a compiler.

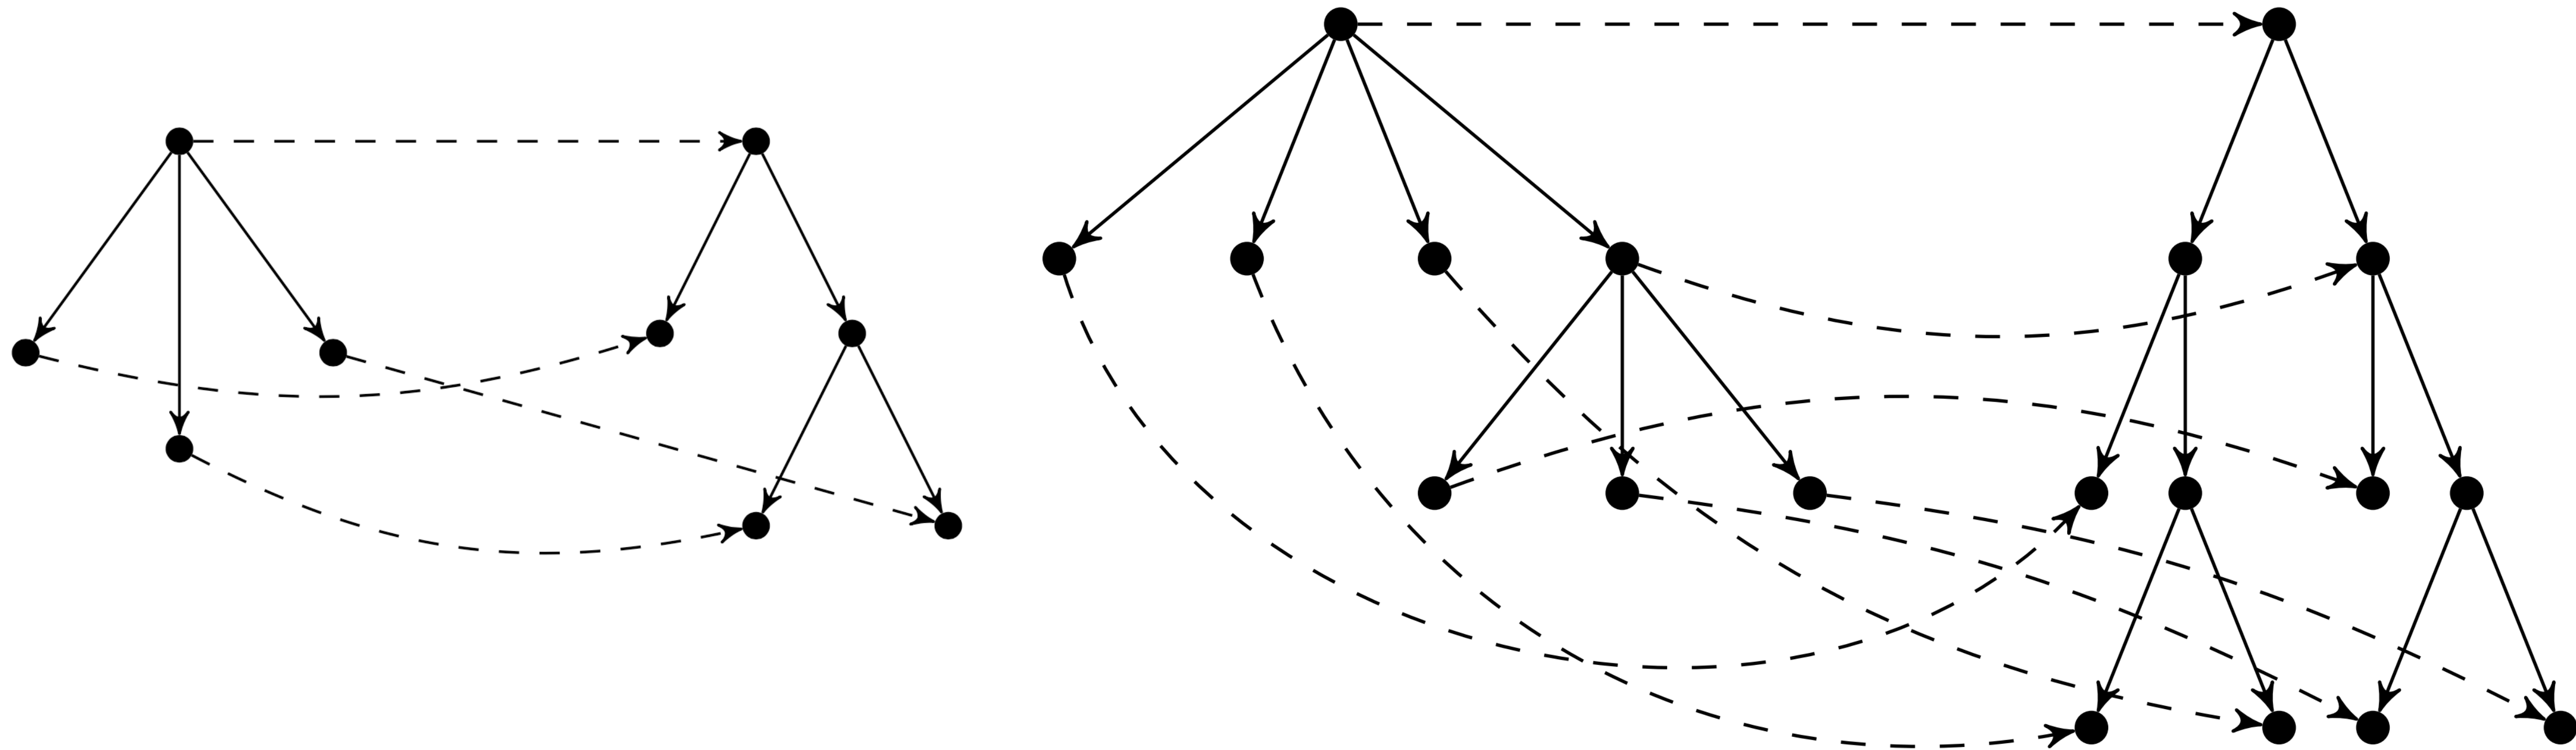


Generating embeddings automatically!

Generating embeddings automatically!

Homomorphic embedding.

Map root to root, leaves to leaves. Respect ancestry.



Generating embeddings automatically!

Homomorphic embedding.

Map root to root, leaves to leaves. Respect ancestry.

Two new algorithms,
both starting with heterogeneous source trees.

Generating embeddings automatically!

Homomorphic embedding.

Map root to root, leaves to leaves. Respect ancestry.

Two new algorithms,
both starting with heterogeneous source trees.

1. If target tree is regular d -ary for some d .

Generating embeddings automatically!

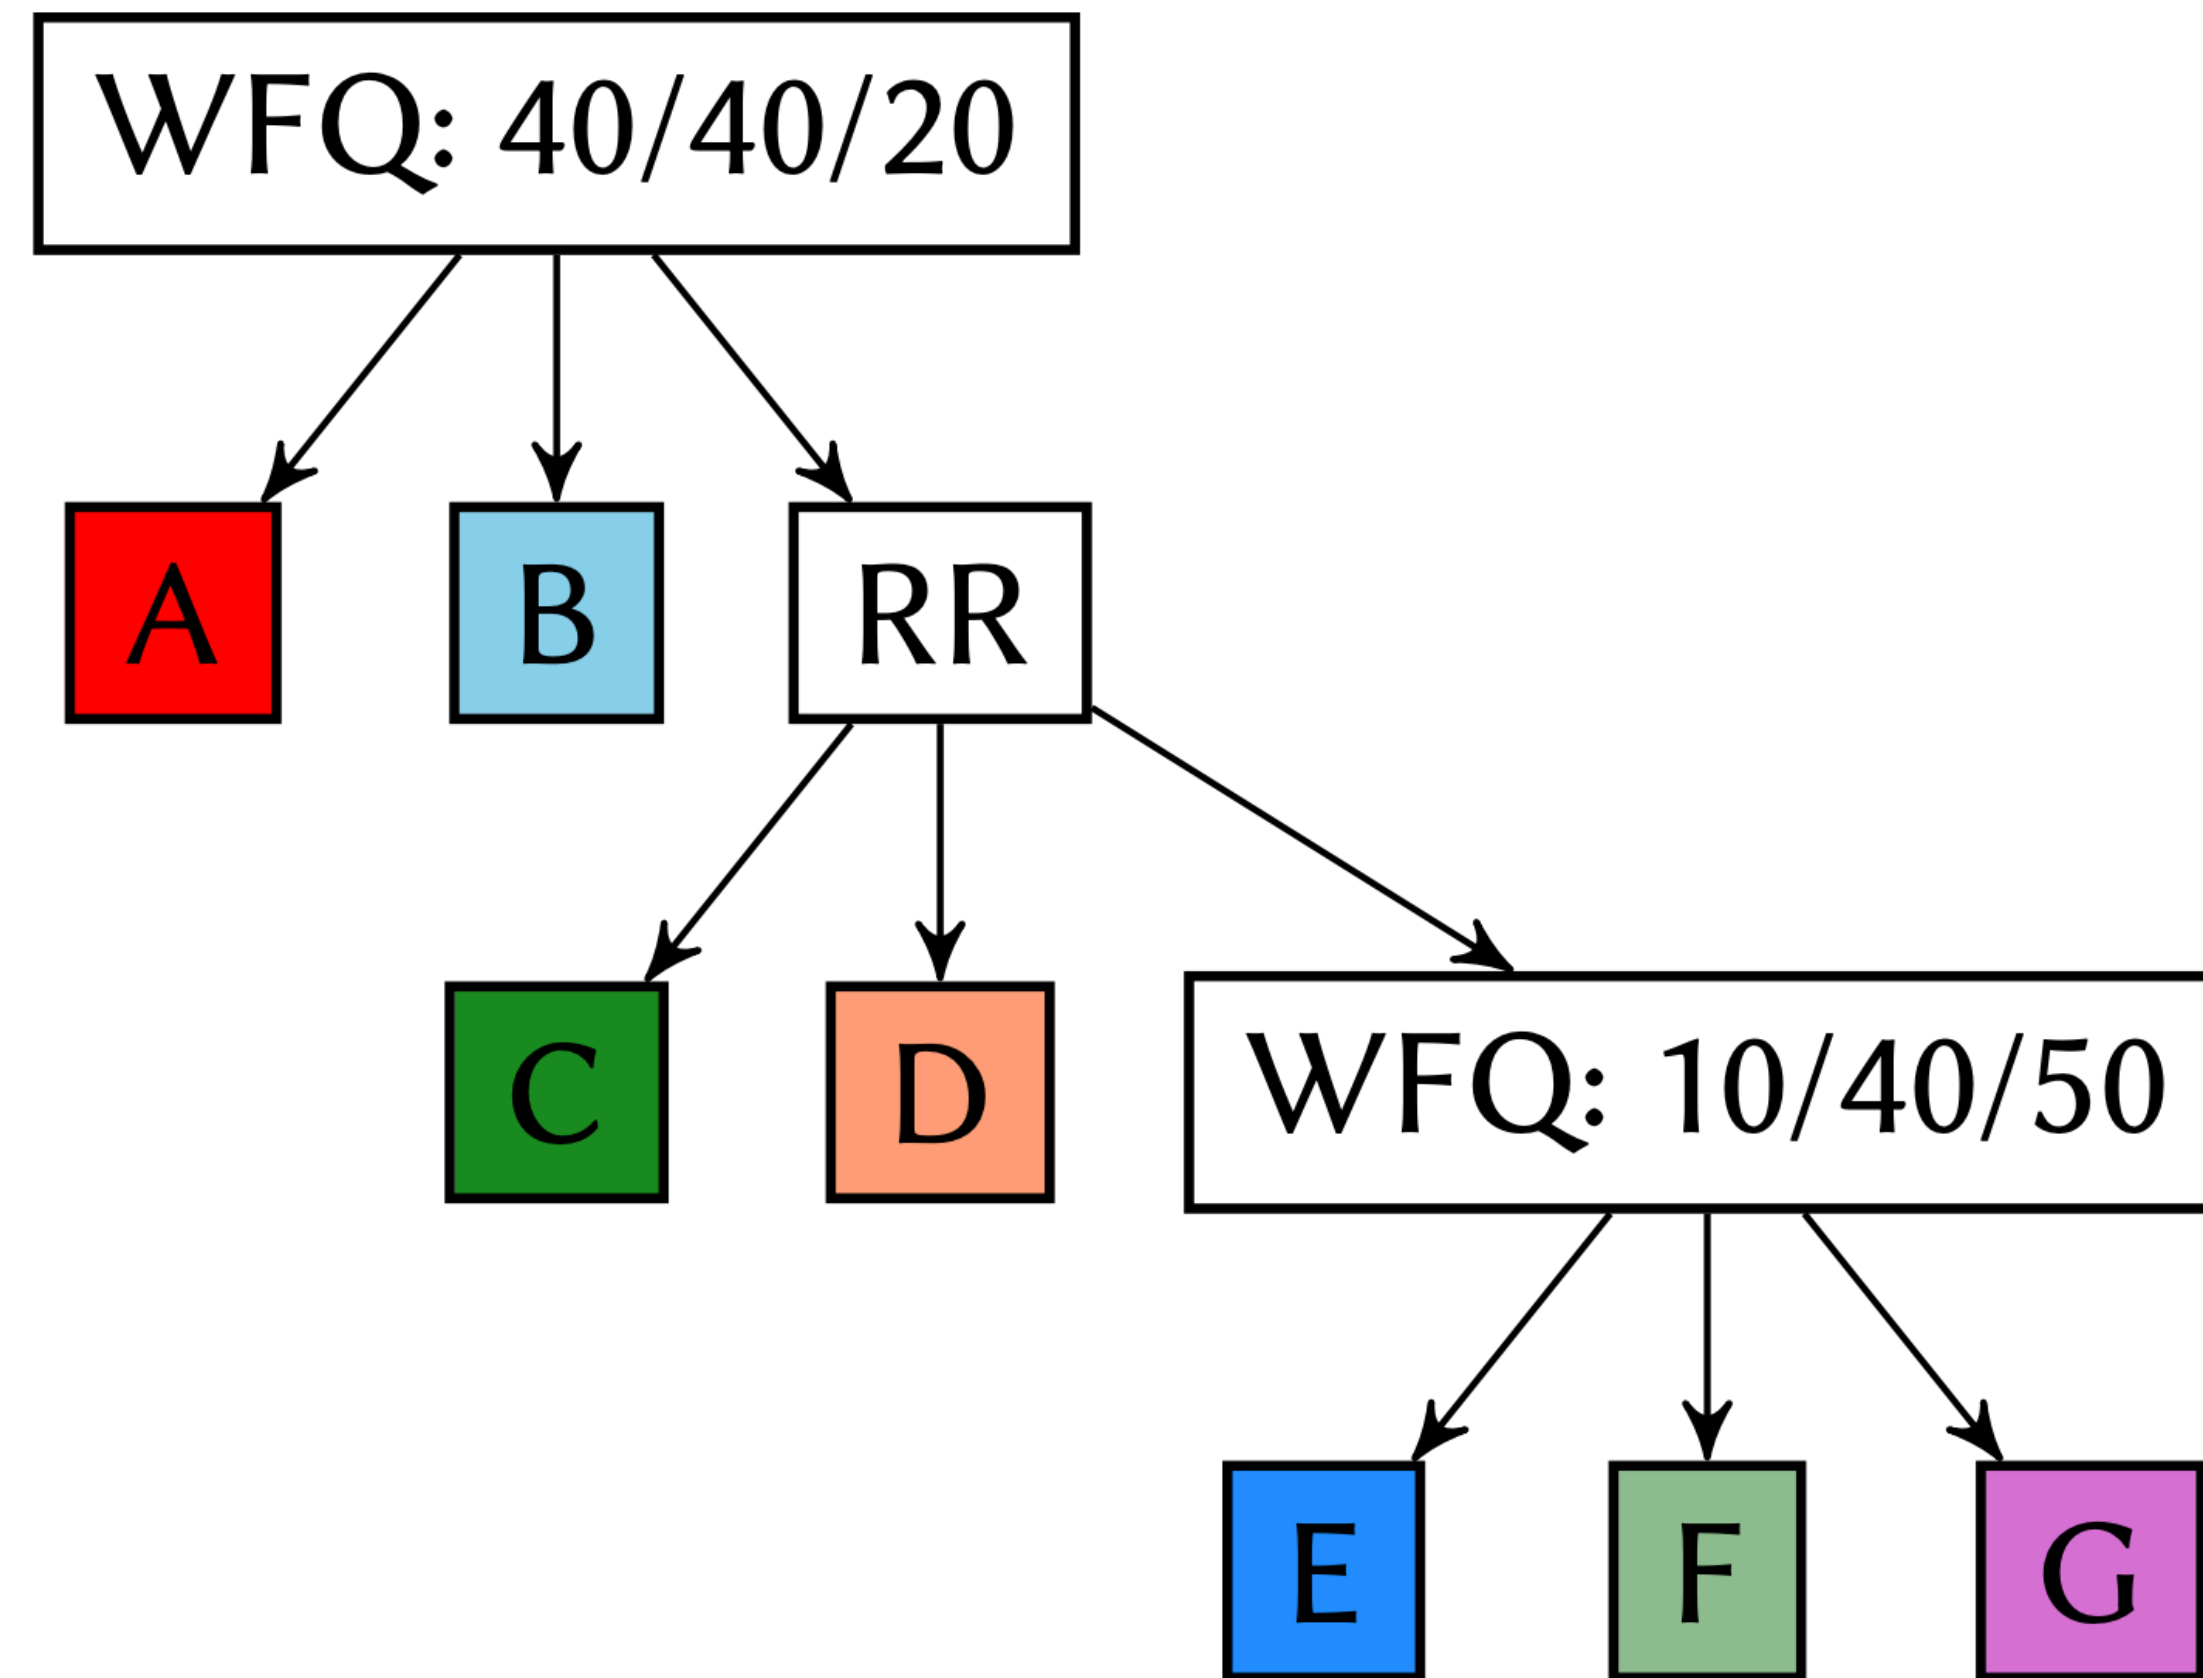
Homomorphic embedding.

Map root to root, leaves to leaves. Respect ancestry.

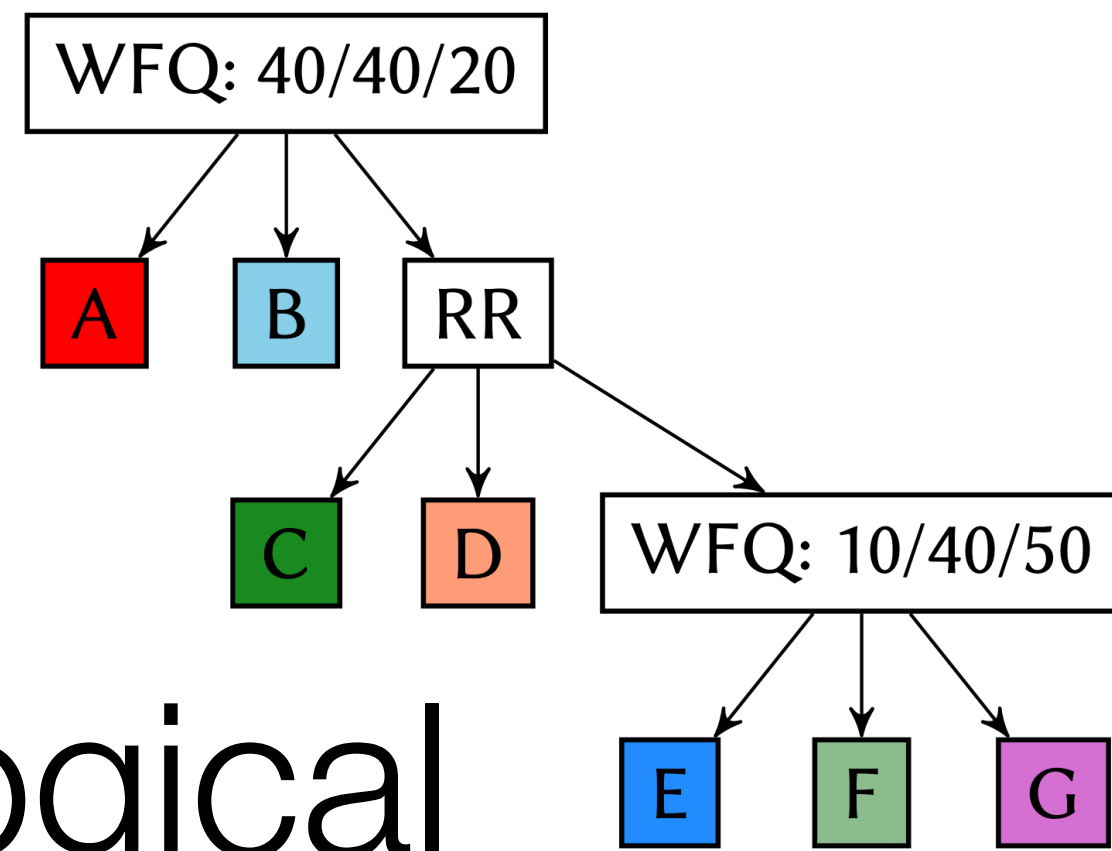
Two new algorithms,
both starting with heterogeneous source trees.

1. If target tree is regular d -ary for some d .
2. If target tree is itself heterogeneous.

Workflow

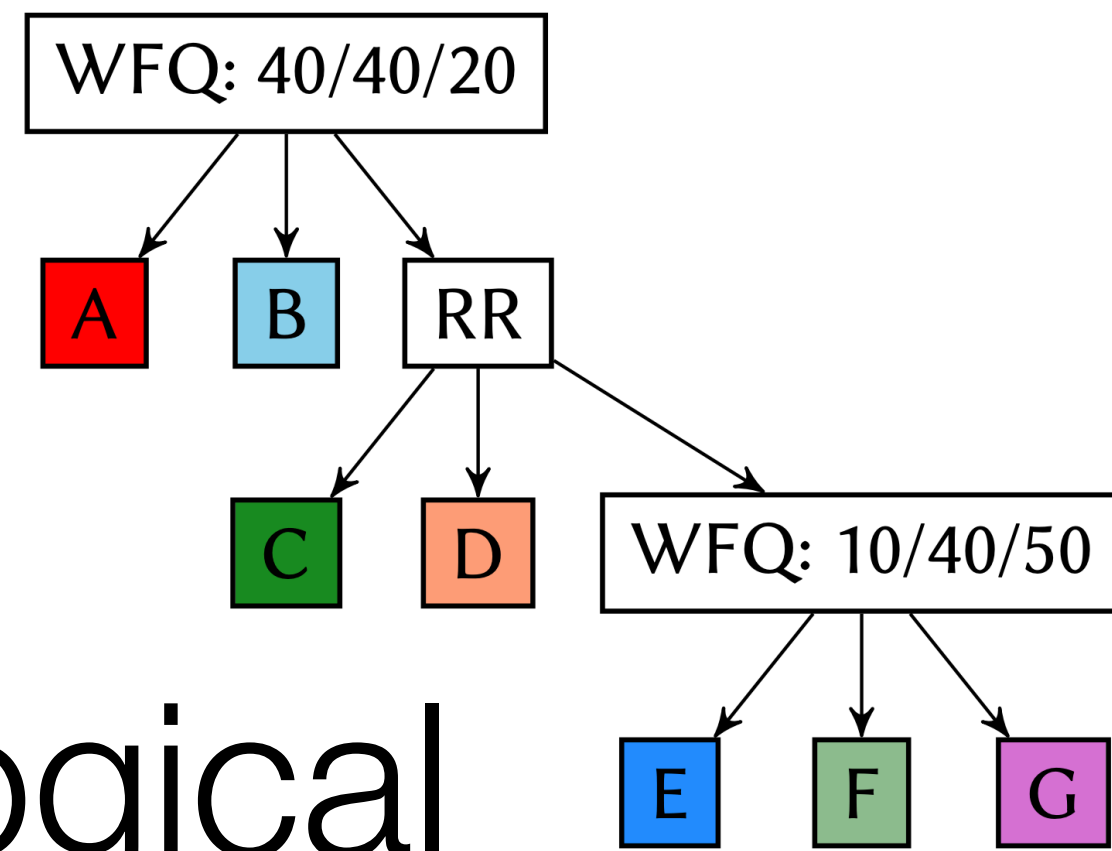


Workflow



logical

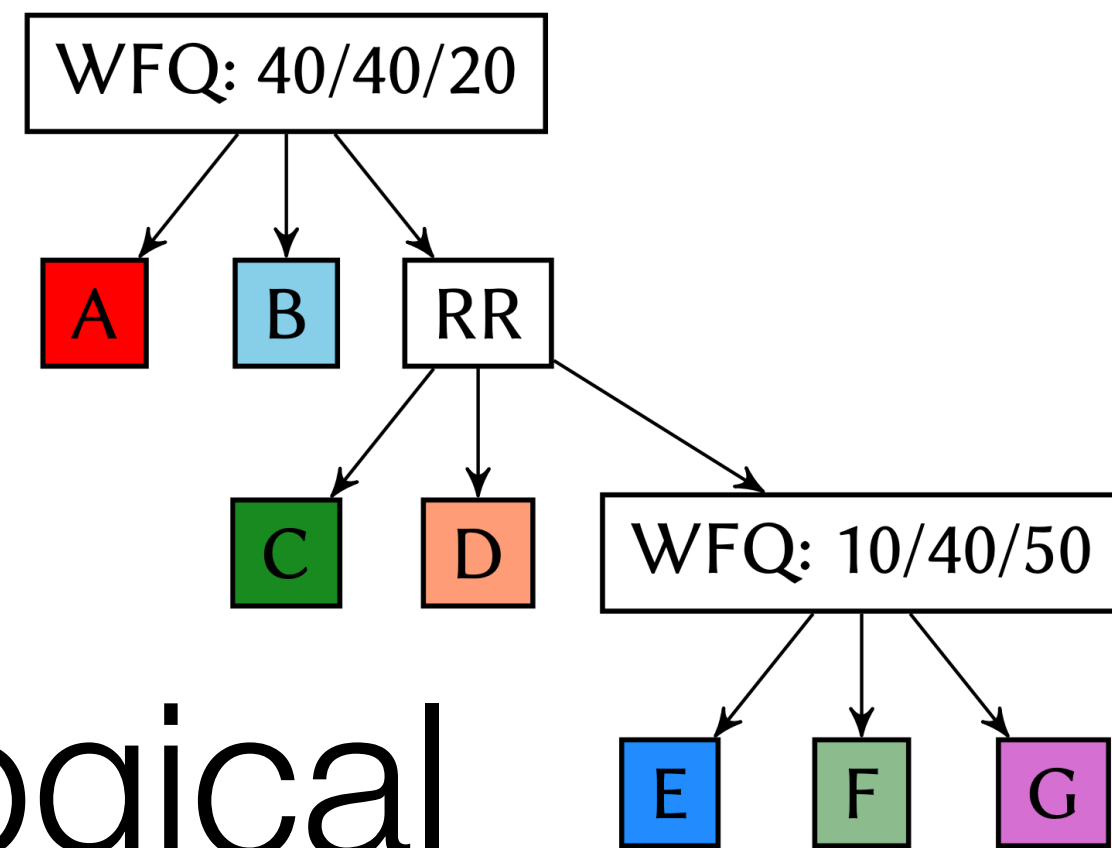
Workflow



logical

But the hardware supports
a regular-branching binary tree.

Workflow

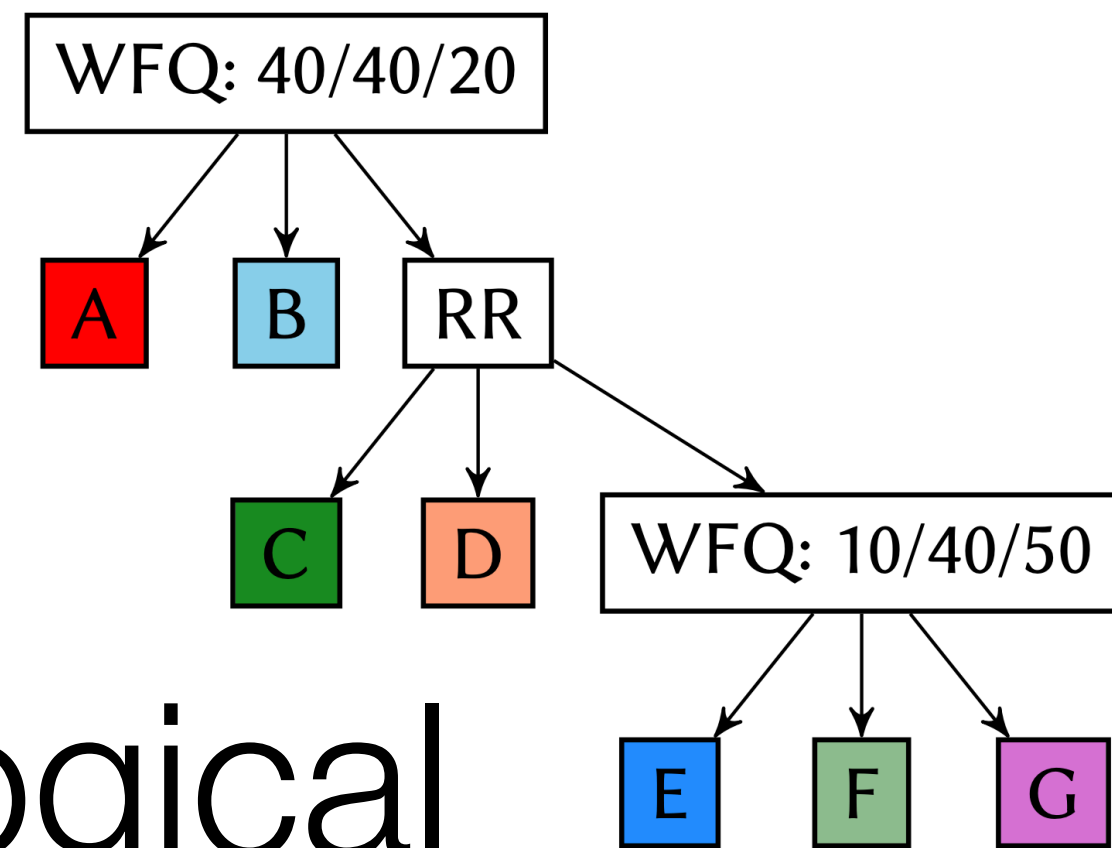


logical

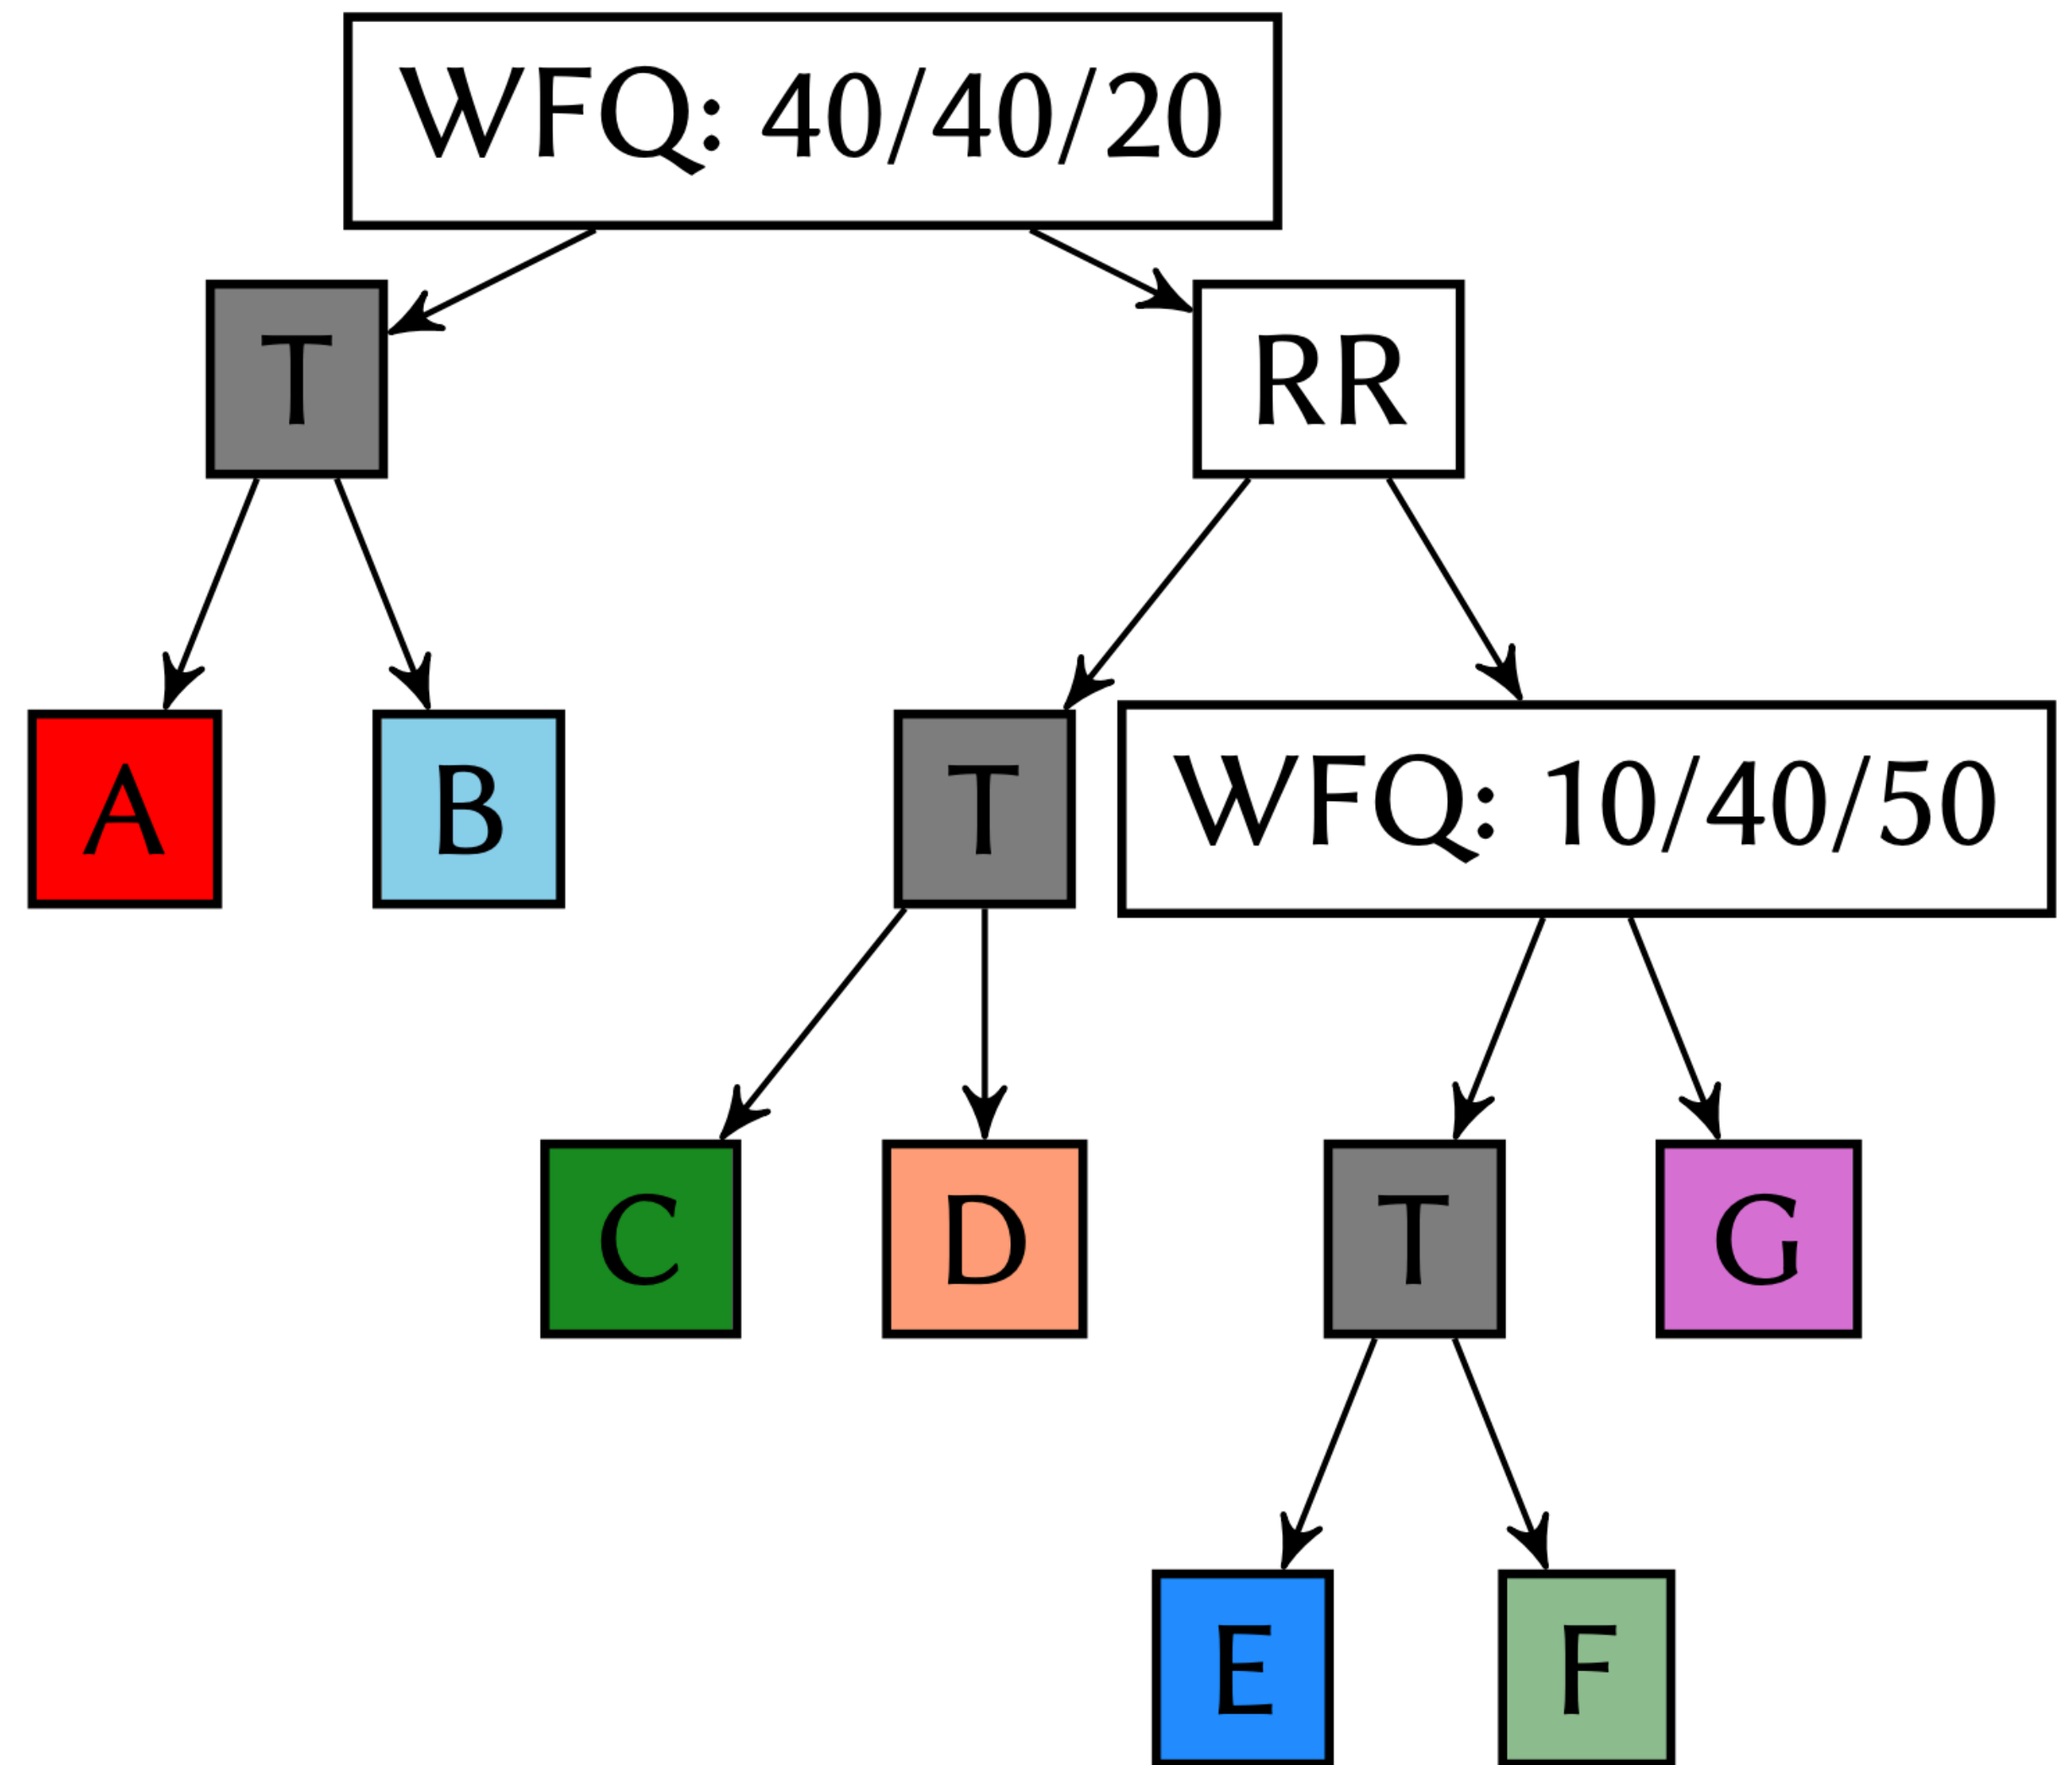
But the hardware supports
a regular-branching binary tree.

Here's how I'll use that tree.

Workflow

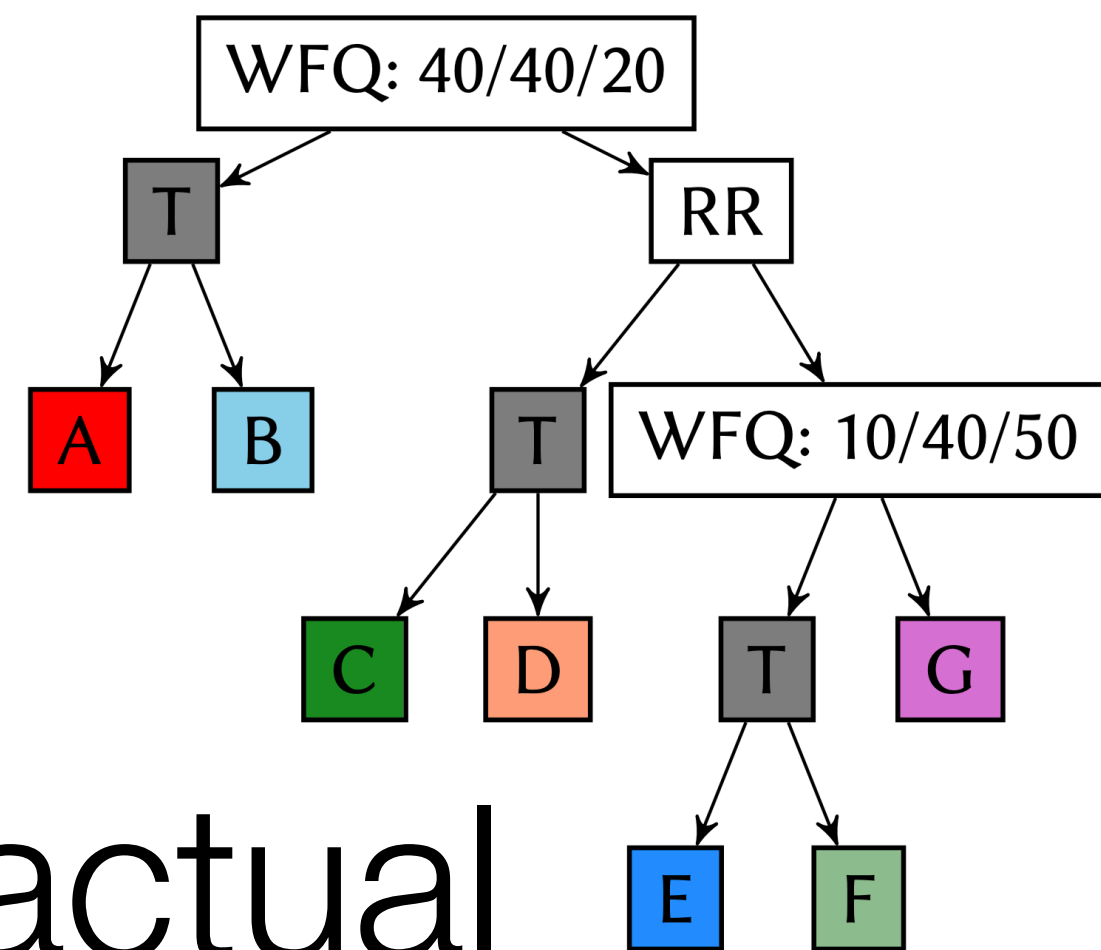
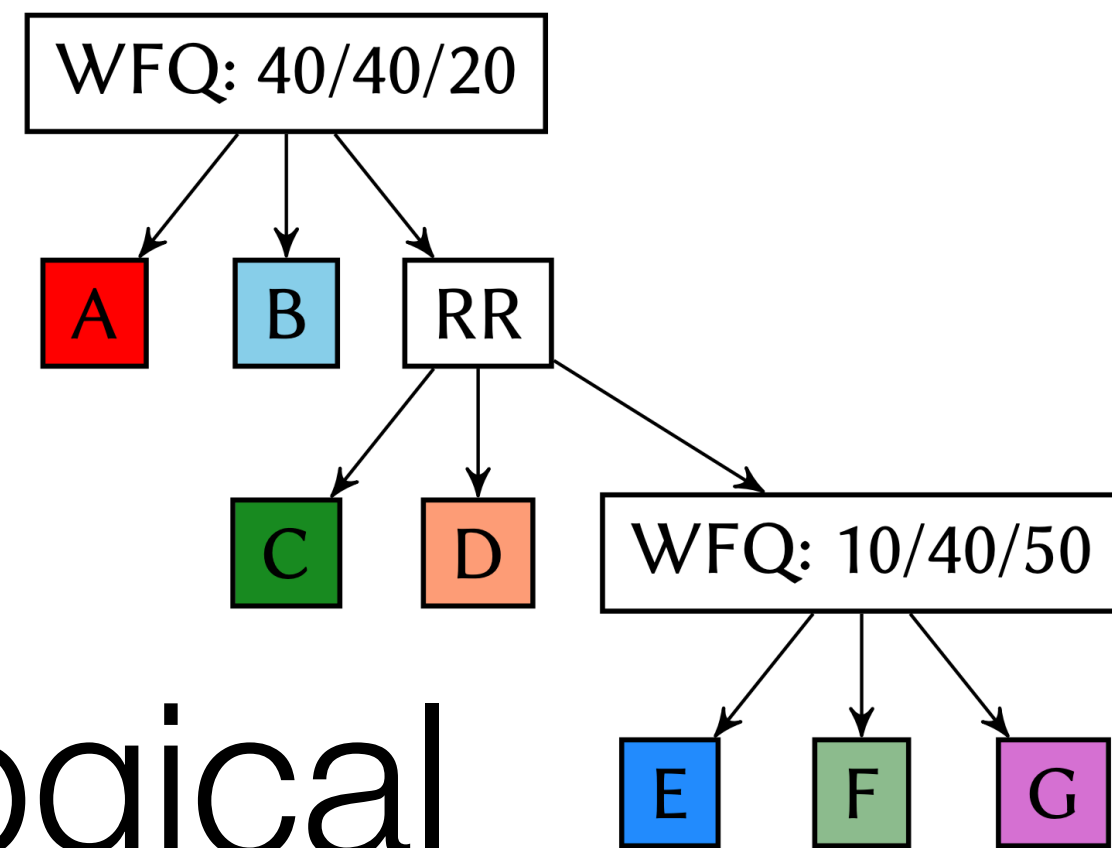


logical

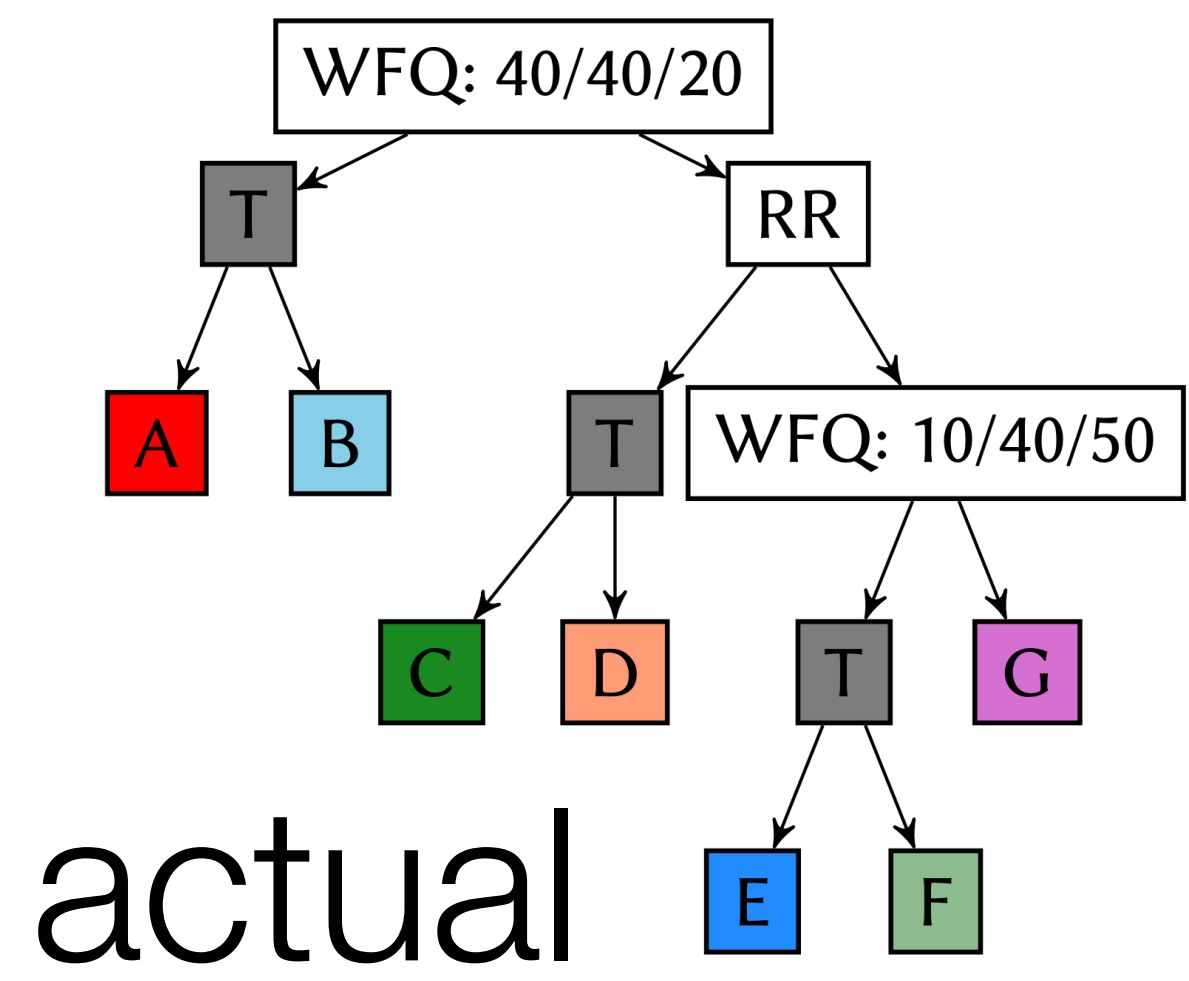
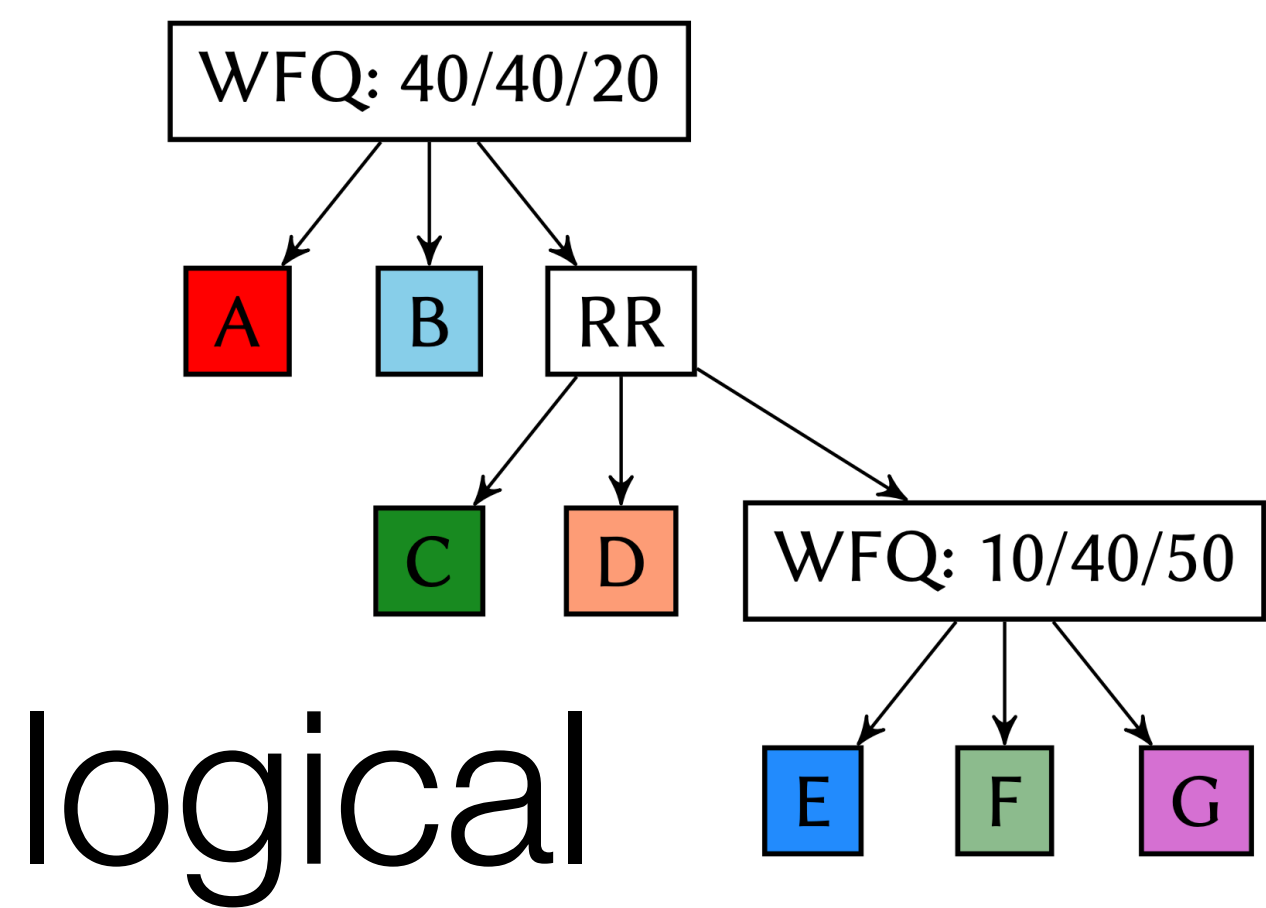


Here's how I'll use that tree.

Workflow



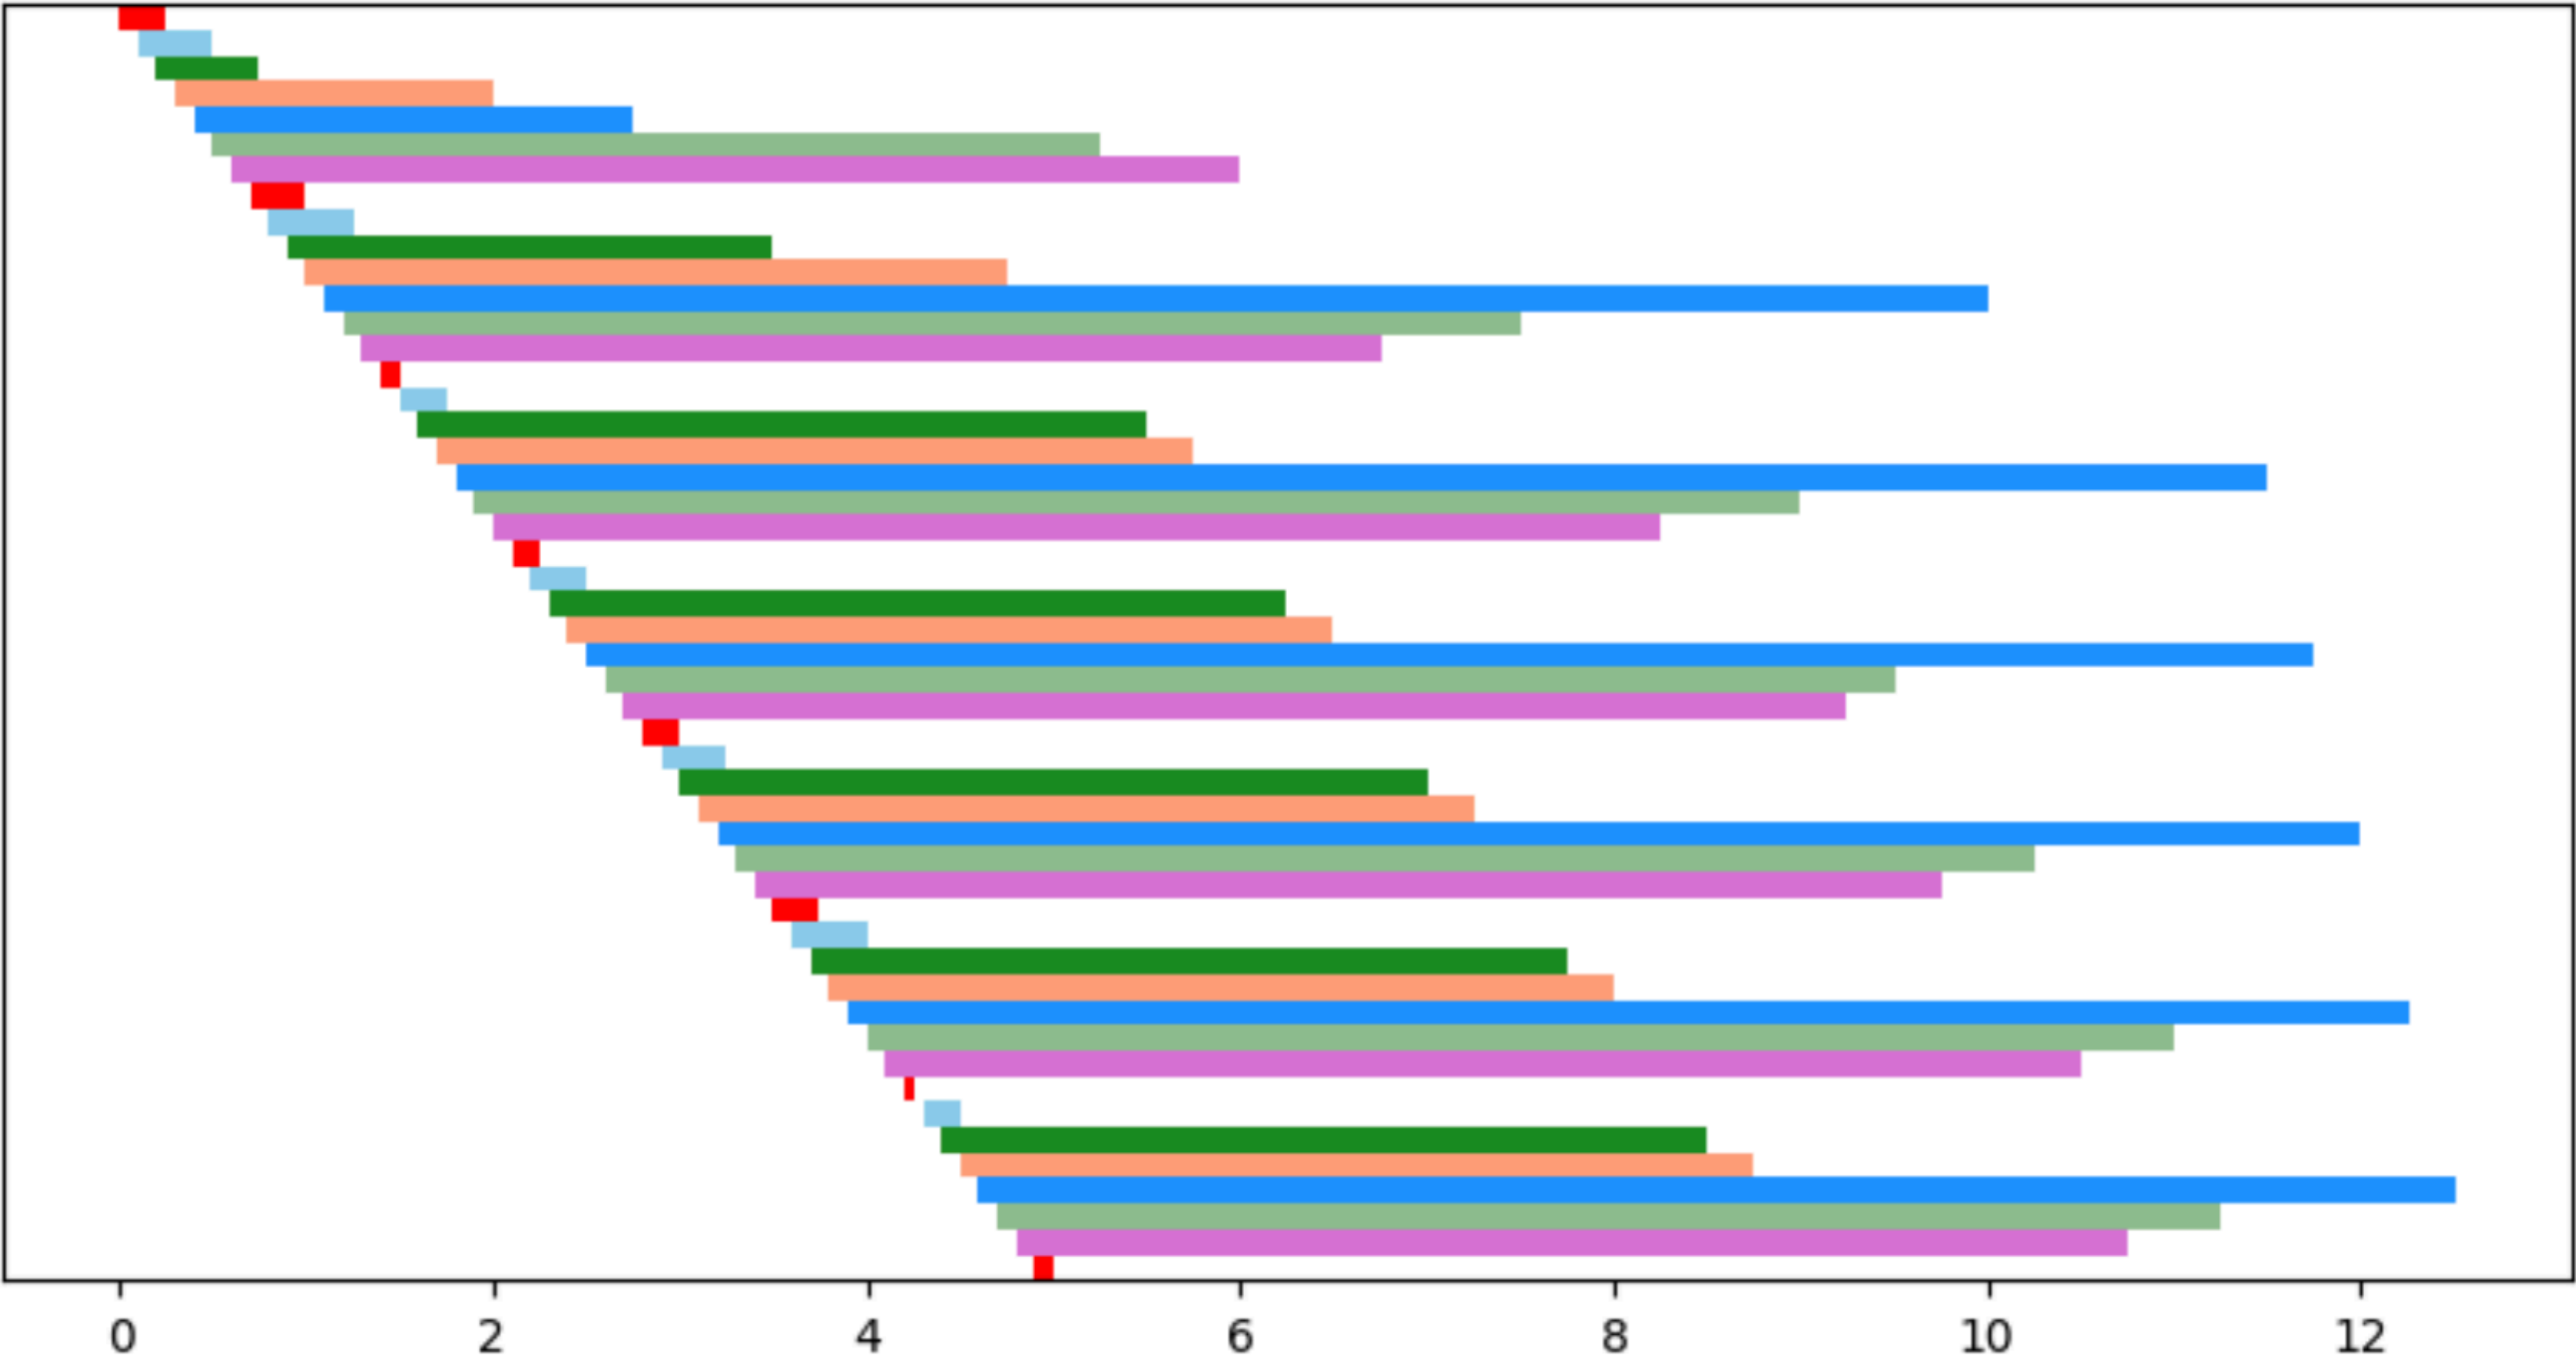
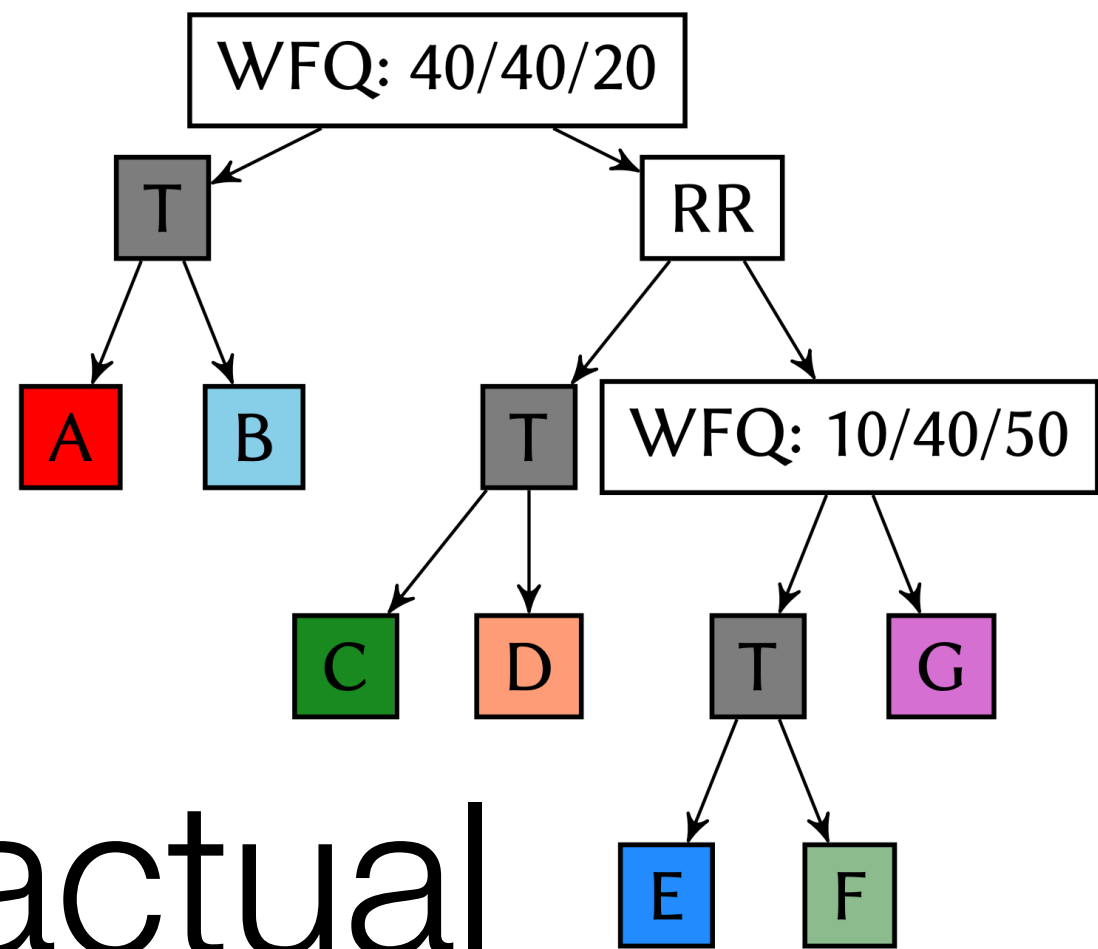
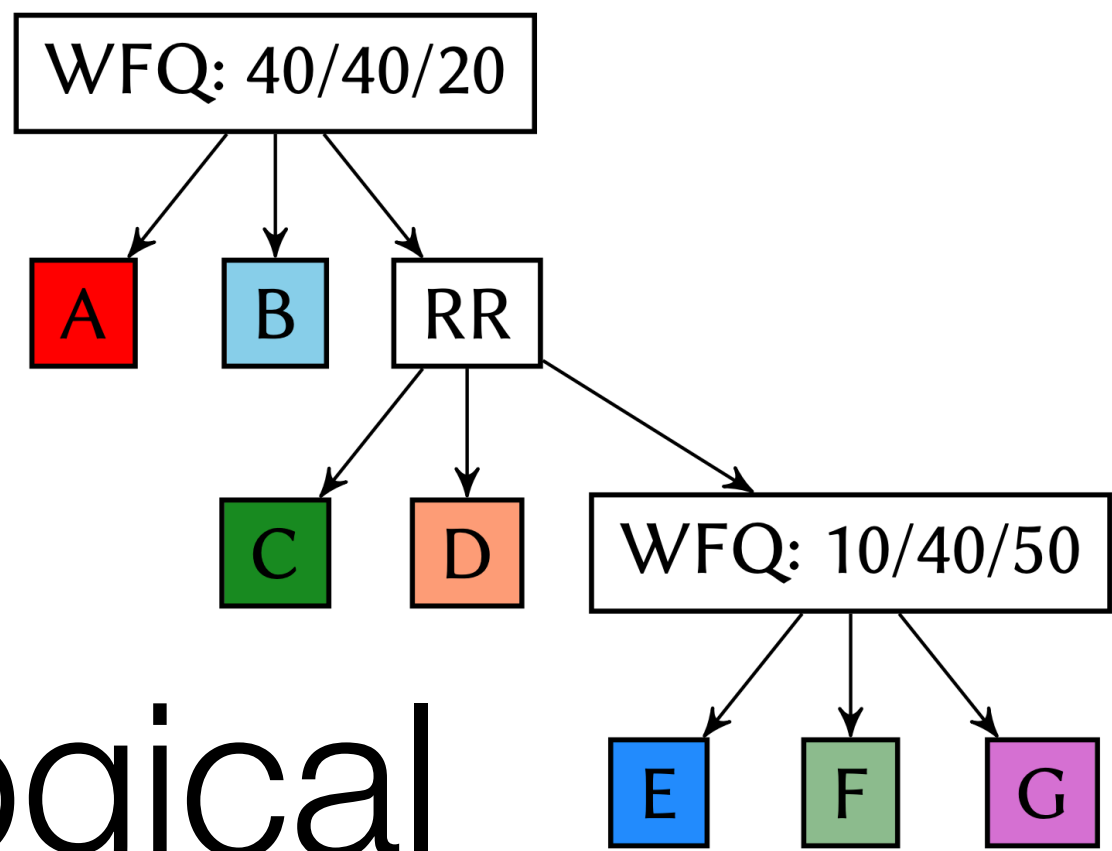
Simulation



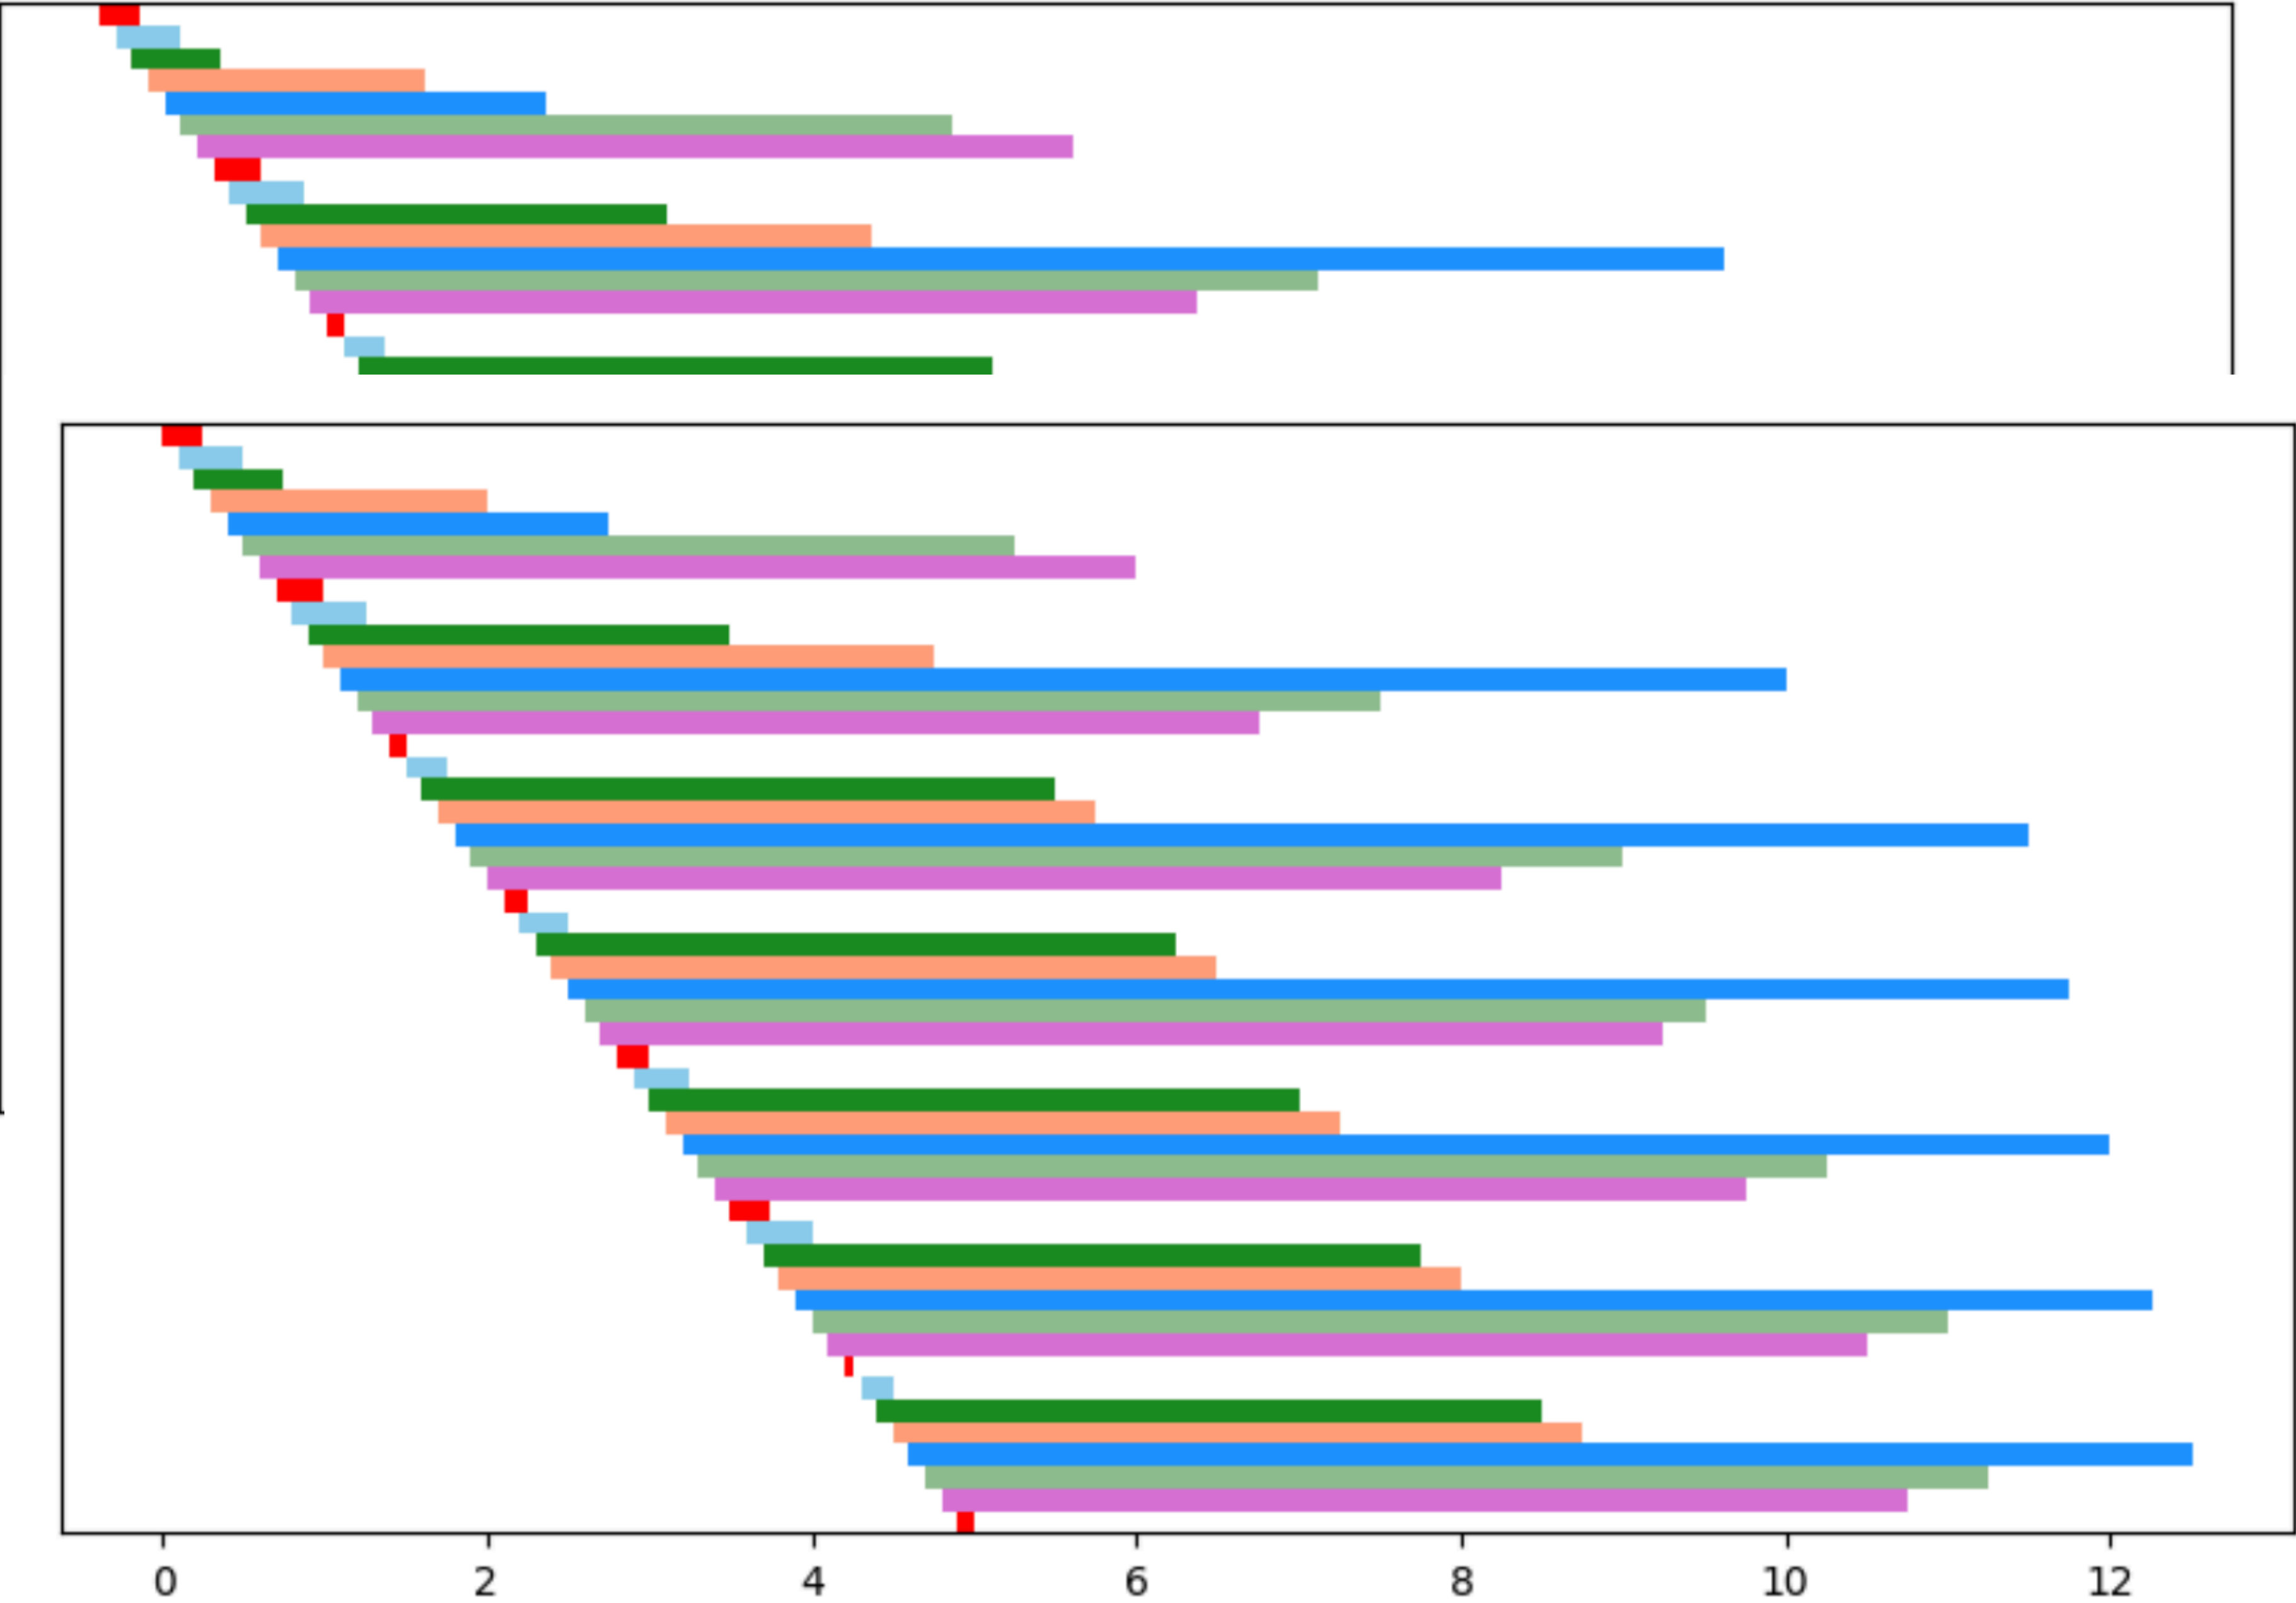
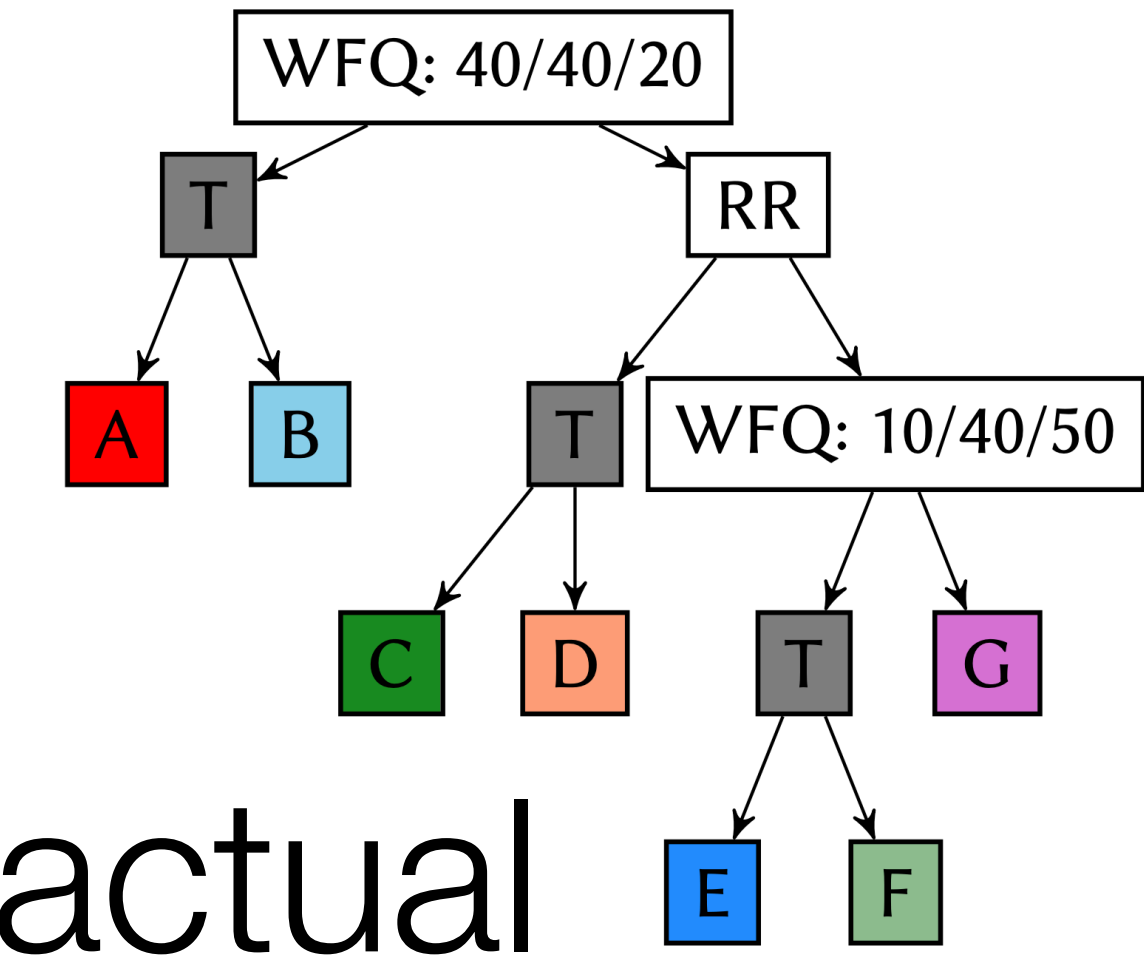
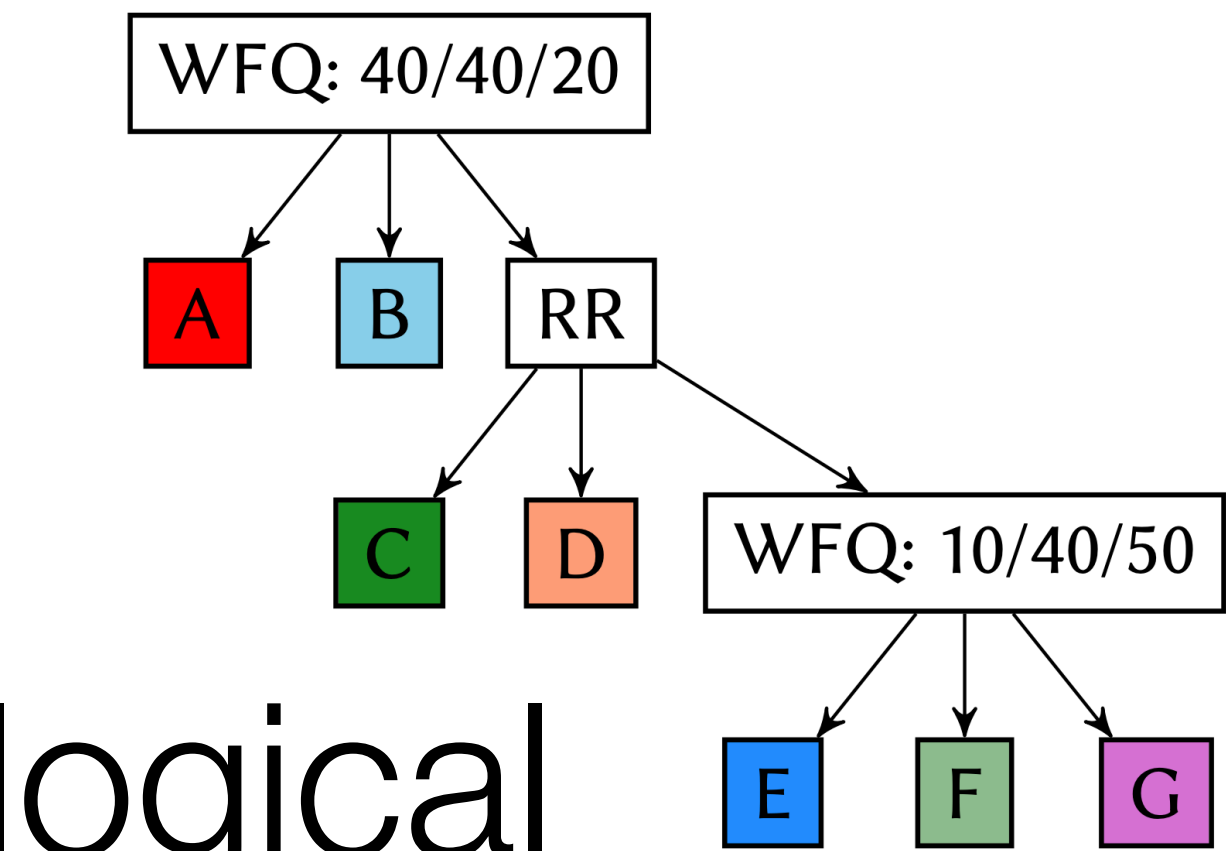
Simulation

logical

actual



Simulation



Underlying formalism

$$\frac{}{* \in \text{Topo}}$$

$$\frac{n \in \mathbb{N} \quad ts \in \text{Topo}^n}{\text{Node}(ts) \in \text{Topo}}$$

$$\frac{p \in \text{PIFO}(\text{Pkt})}{\text{Leaf}(p) \in \text{PIFOTree}(\underbrace{*}_{\text{Topo}})}$$

$$\frac{n \in \mathbb{N} \quad ts \in \text{Topo}^n \quad p \in \text{PIFO}(\{1, \dots, n\}) \quad \forall 1 \leq i \leq n. qs[i] \in \text{PIFOTree}(ts[i])}{\text{Internal}(qs, p) \in \text{PIFOTree}(\underbrace{\text{Node}(ts)}_{\text{Topo}})}$$

$$\frac{r \in \text{Rk}}{r \in \text{Path}(\underbrace{*}_{\text{Topo}})}$$

$$\frac{ts \in \text{Topo}^n \quad 1 \leq i \leq n \quad r \in \text{Rk} \quad pt \in \text{Path}(ts[i])}{(i, r) :: \text{Path}(\underbrace{\text{Node}(ts)}_{\text{Topo}})}$$

$$\frac{\text{PUSH}(p, pkt, r) = p'}{\text{push}(\text{Leaf}(p), pkt, r) = \underbrace{\text{Leaf}(p')}_{\text{PIFOTree}}}$$

$$\frac{\text{push}(qs[i], pkt, pt) = q' \quad \text{PUSH}(p, i, r) = p'}{\text{push}(\text{Internal}(qs, p), pkt, \underbrace{(i, r) :: pt}_{\text{Path}}) = \underbrace{\text{Internal}(qs[i/q'], p')}_{\text{PIFOTree}}}$$

A general way to deploy PIFO trees.

A general way to deploy PIFO trees.



Let the hardware support some tree.

A general way to deploy PIFO trees.



Let the human
program against some tree.



Let the hardware
support some tree.

A general way to deploy PIFO trees.



Let the human
program against some tree.



Let the hardware
support some tree.

A general way to deploy PIFO trees.



A general way to deploy PIFO trees.



A general way to deploy PIFO trees.



A general way to deploy PIFO trees.



Let the human
program against some tree.



Let the hardware
support some tree.

A general way to deploy PIFO trees.



A general way to deploy PIFO trees.



A general way to deploy PIFO trees.



A general way to deploy PIFO trees.



Let the human
program against some tree.

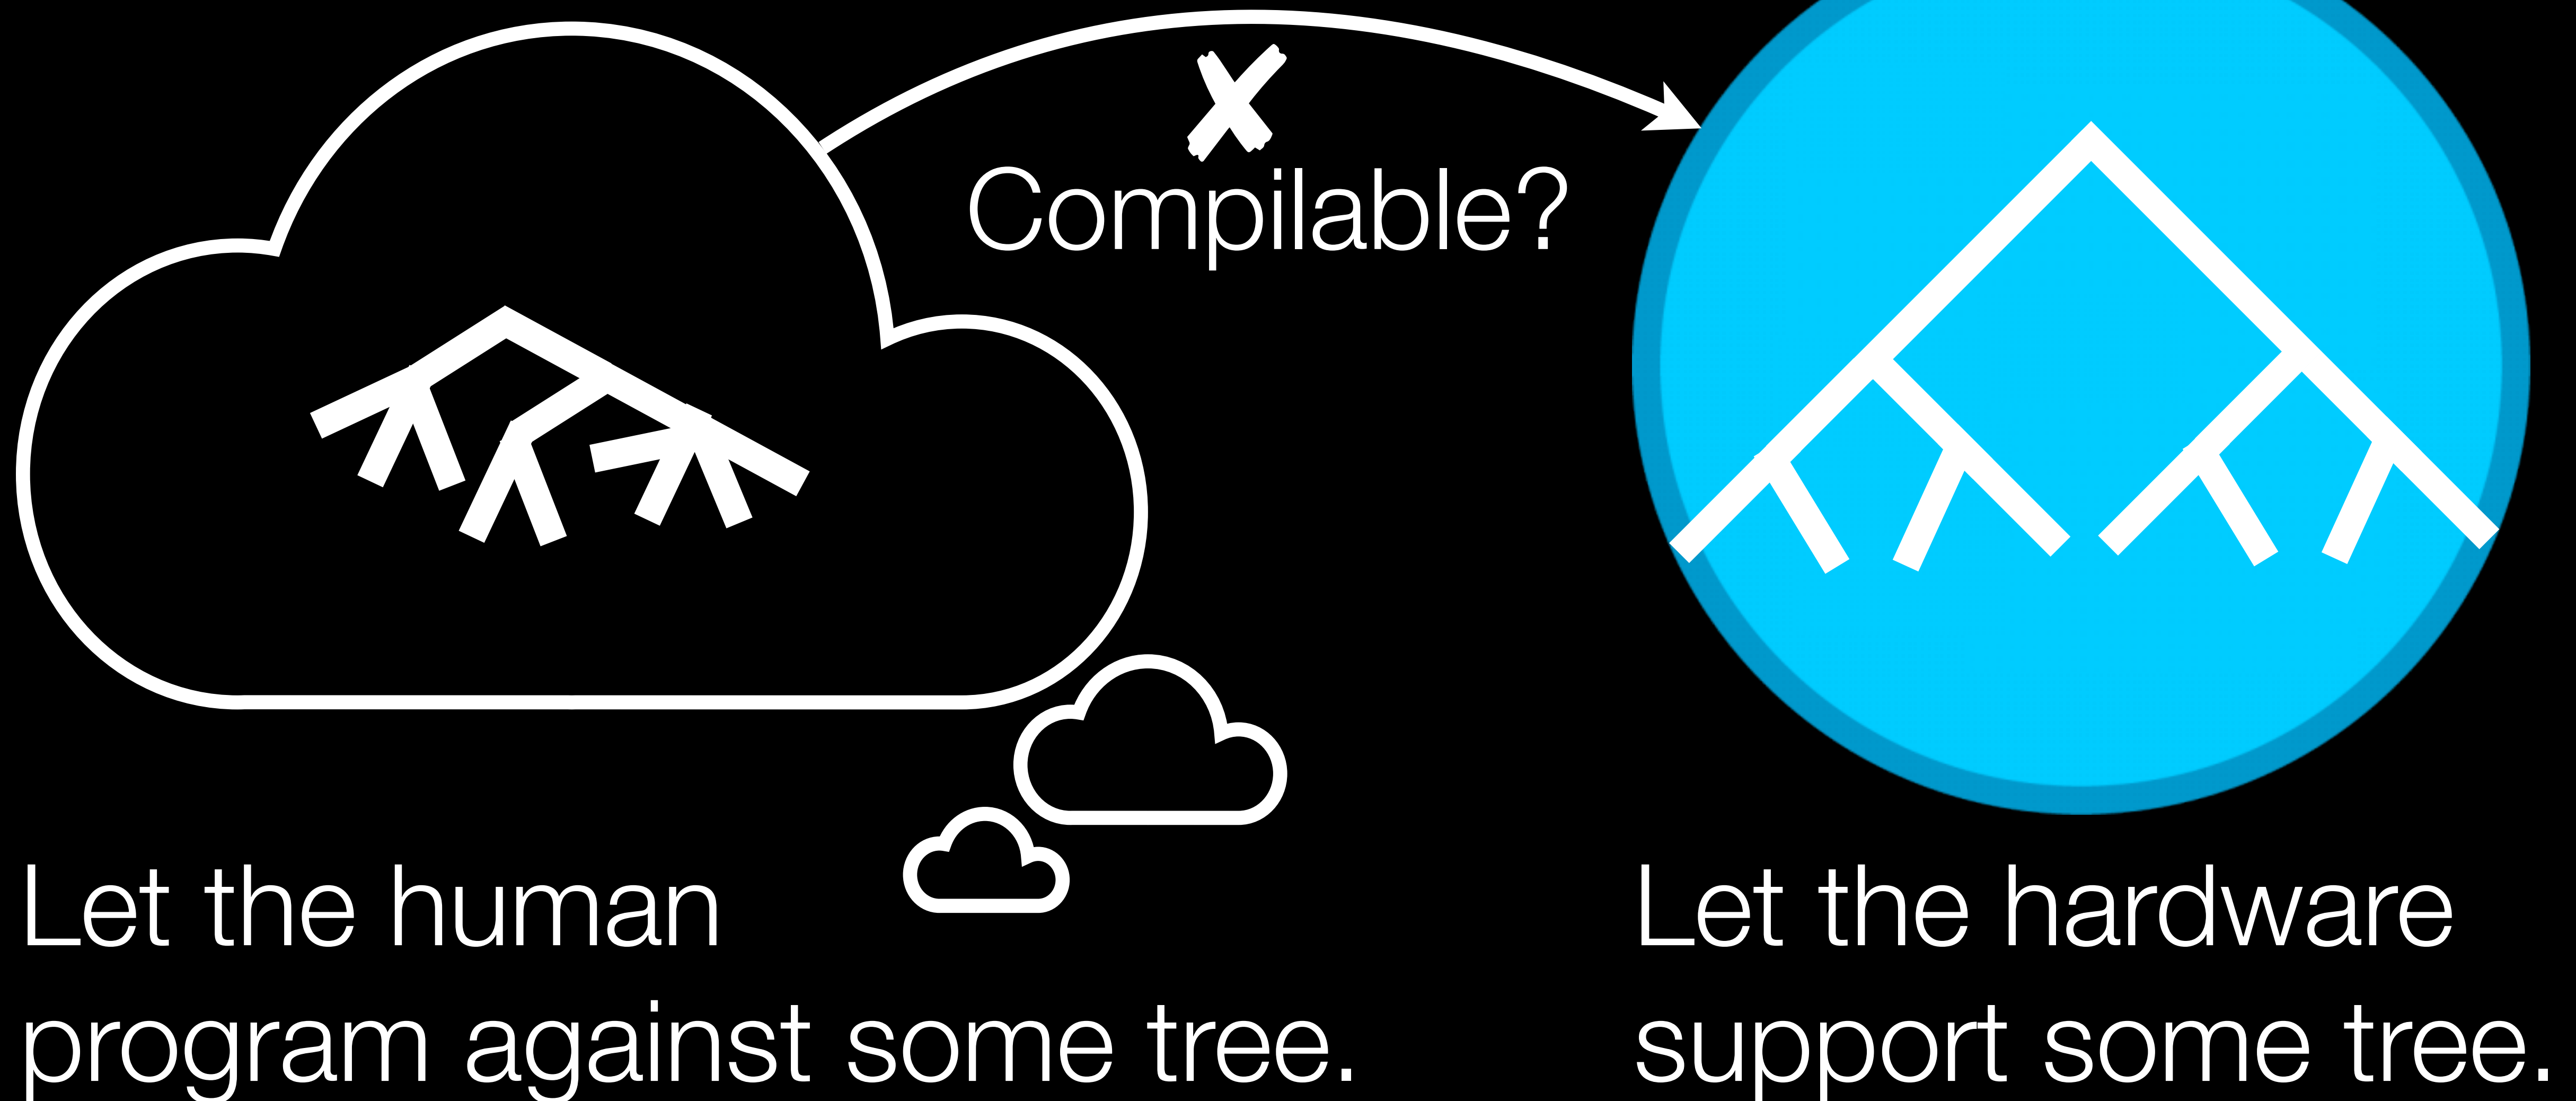


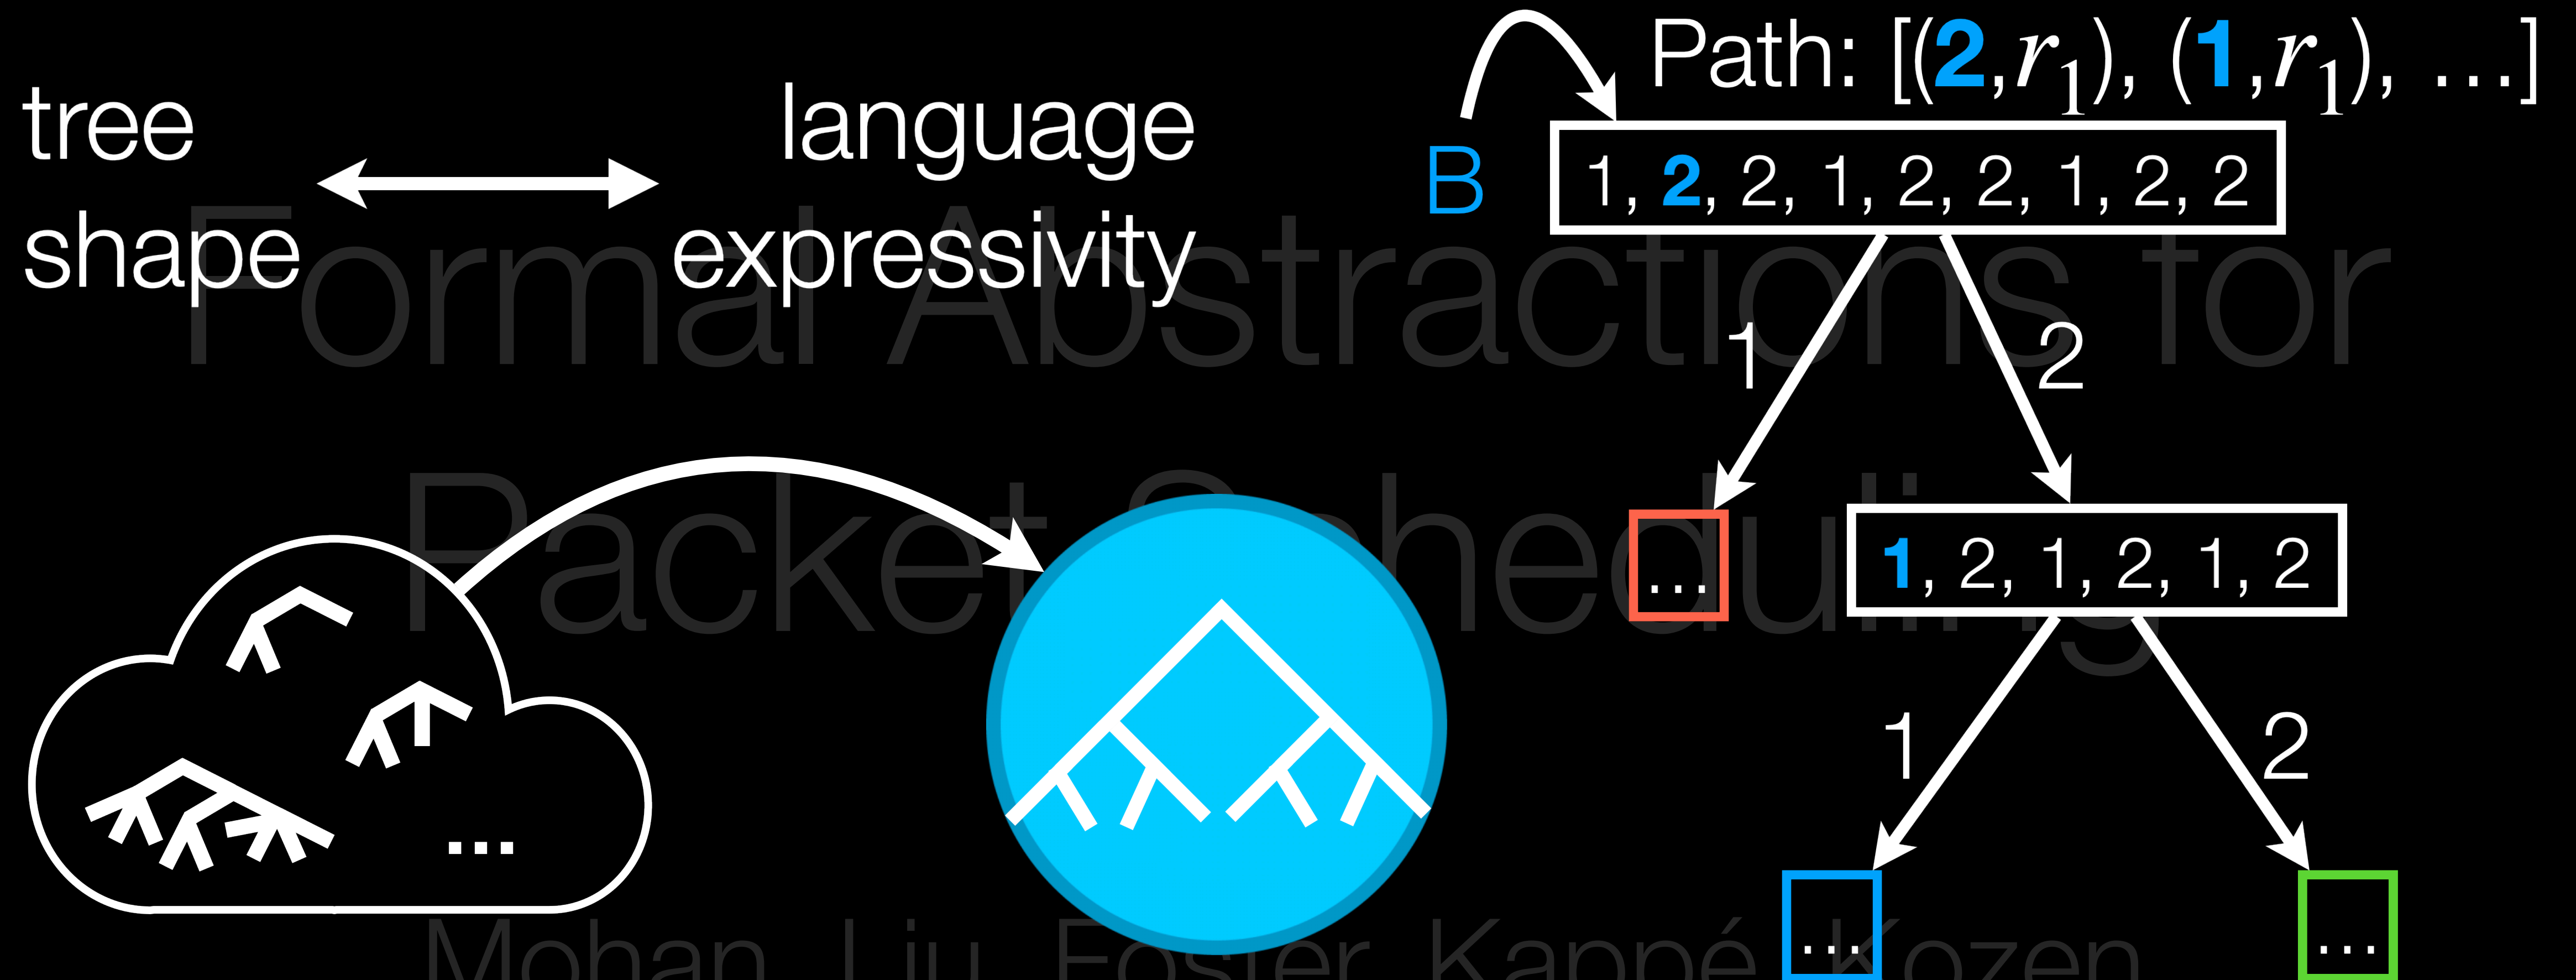
Let the hardware
support some tree.

A general way to deploy PIFO trees.



A general way to deploy PIFO trees.





cs.cornell.edu/~amohan