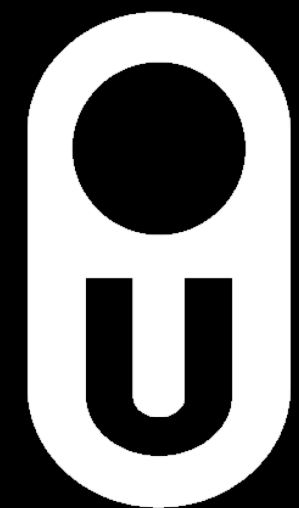


Formal Abstractions for Packet Scheduling

Mohan, Liu, Foster, Kappé, Kozen



SDN made networks programmable.



SDN made networks programmable.

Early goal: routing.



SDN made networks programmable.

Early goal: routing.

But now we need control over *scheduling*.



SDN made networks programmable.

Early goal: routing.

But now we need control over *scheduling*.



SDN made networks programmable.

Early goal: routing.

But now we need control over *scheduling*.

Basic tools work fine...

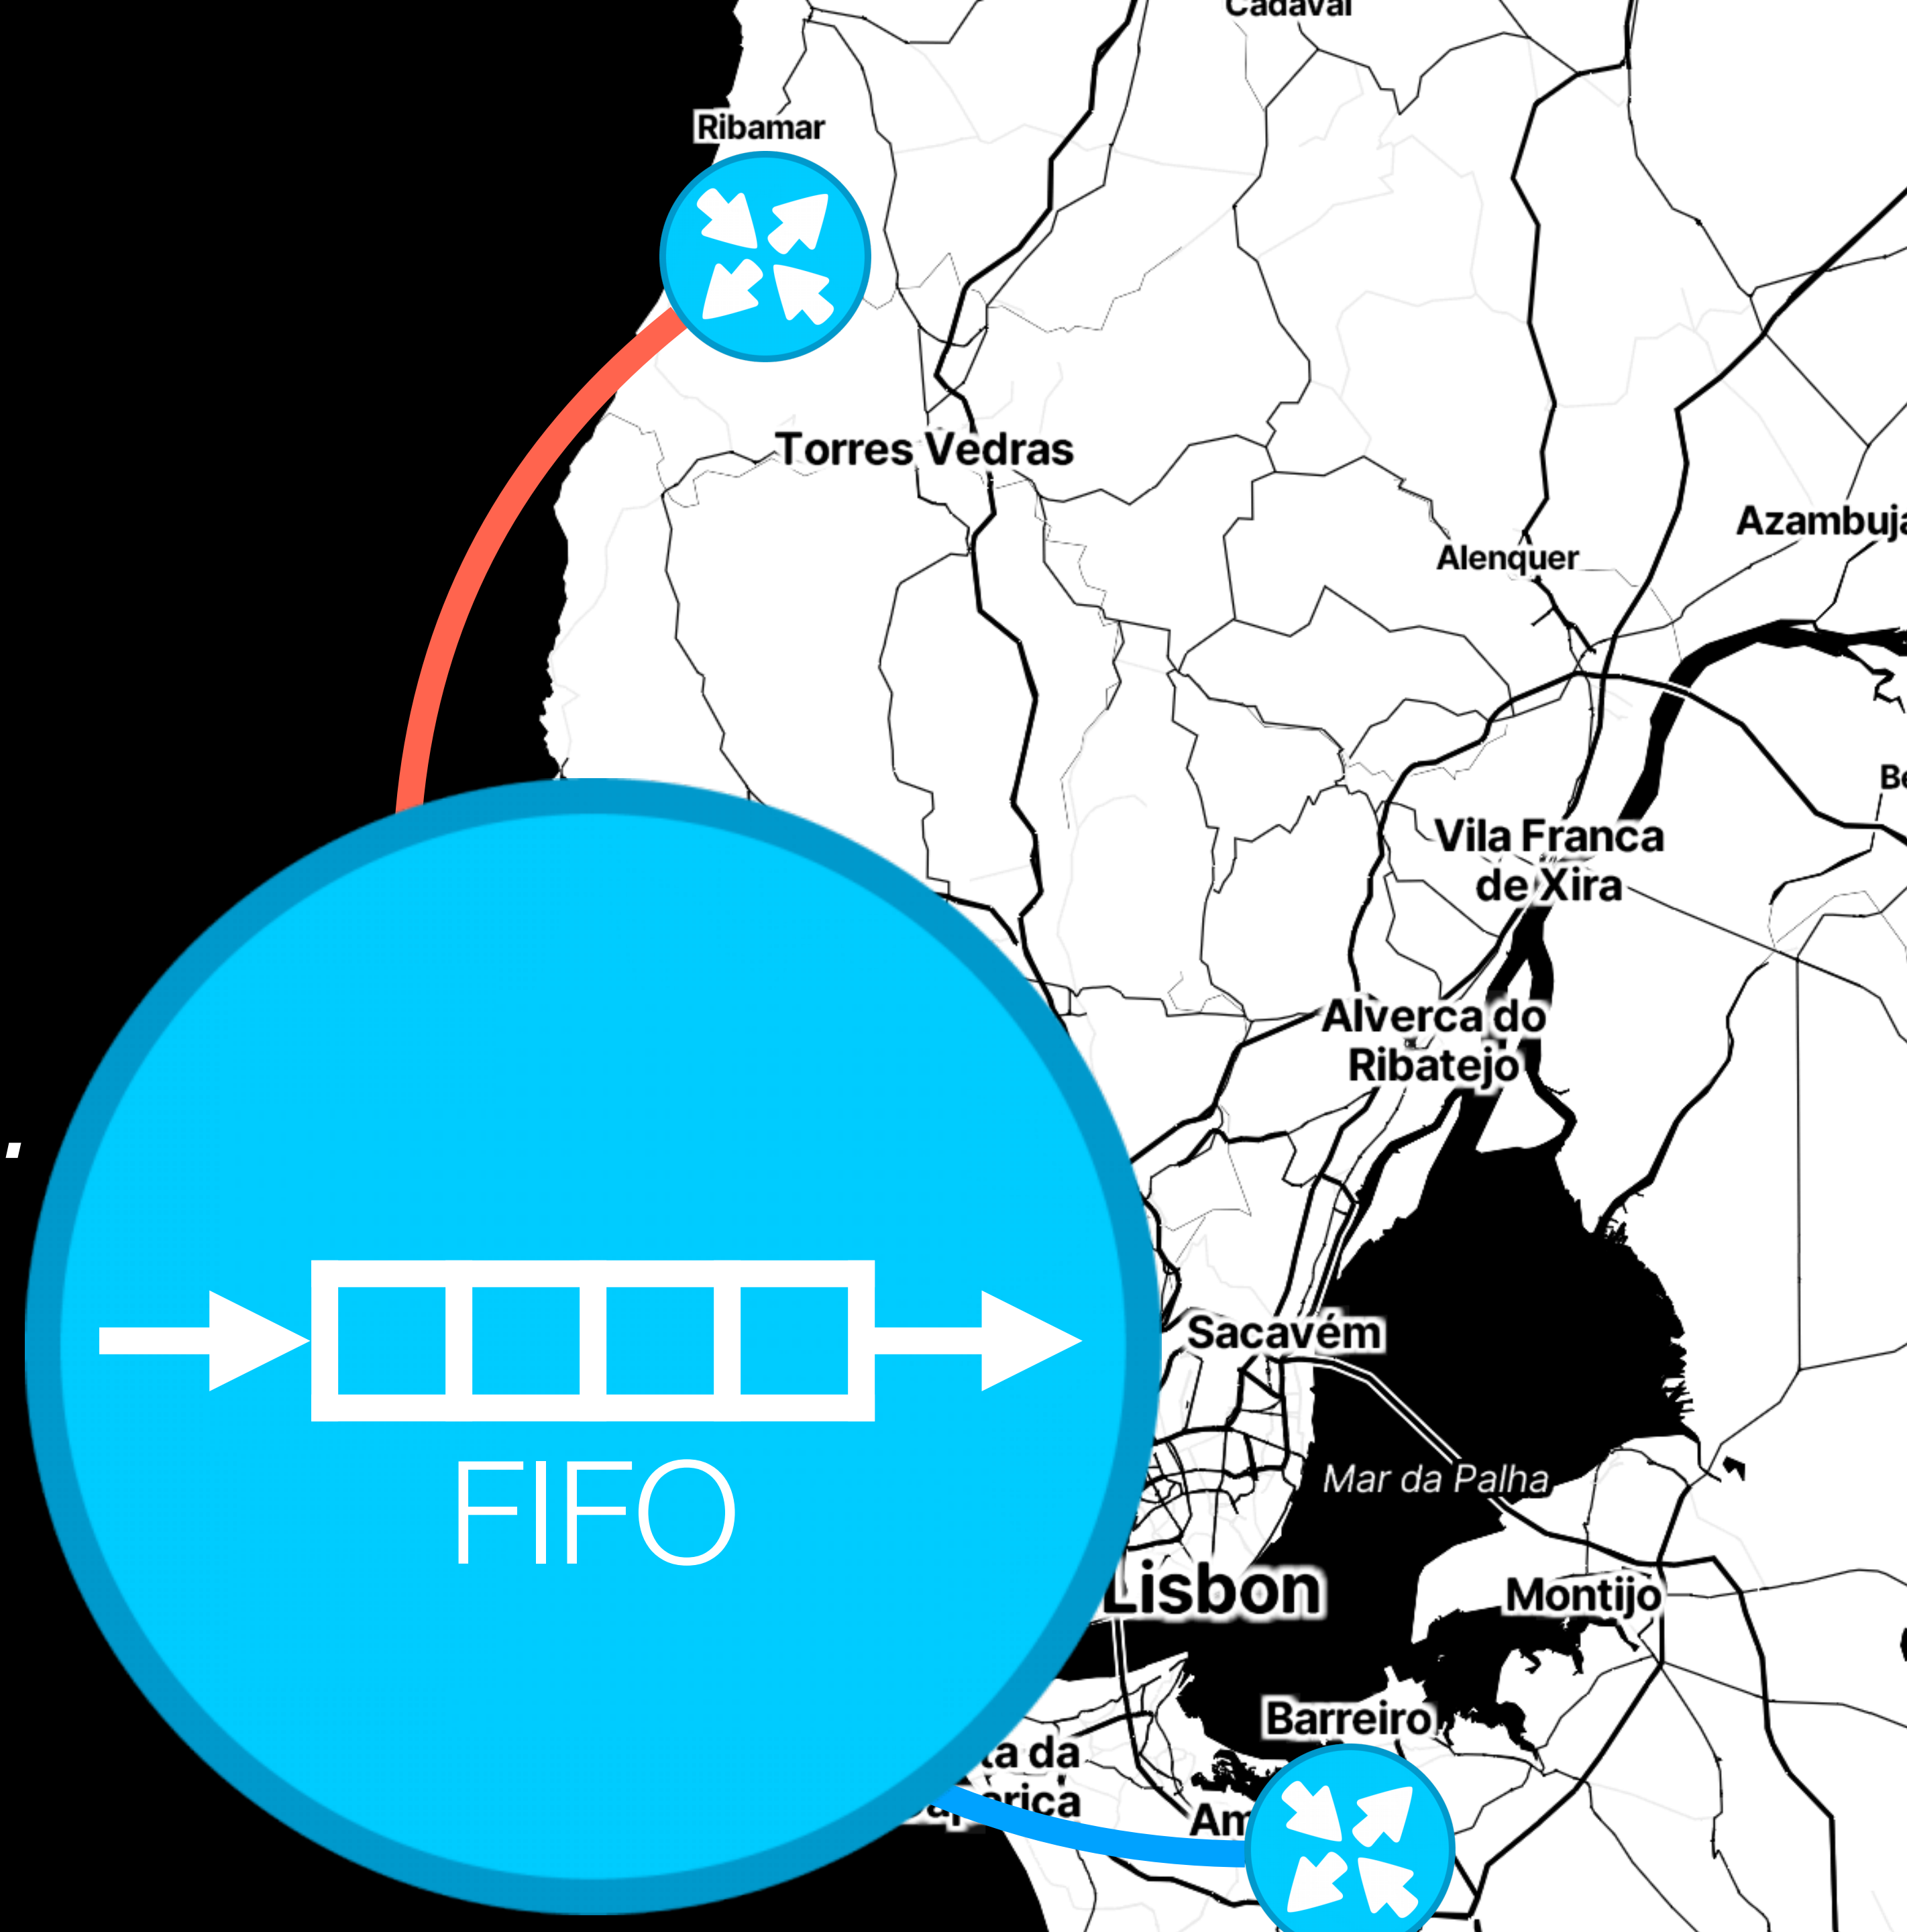


SDN made networks programmable.

Early goal: routing.

But now we need control over *scheduling*.

Basic tools work fine...

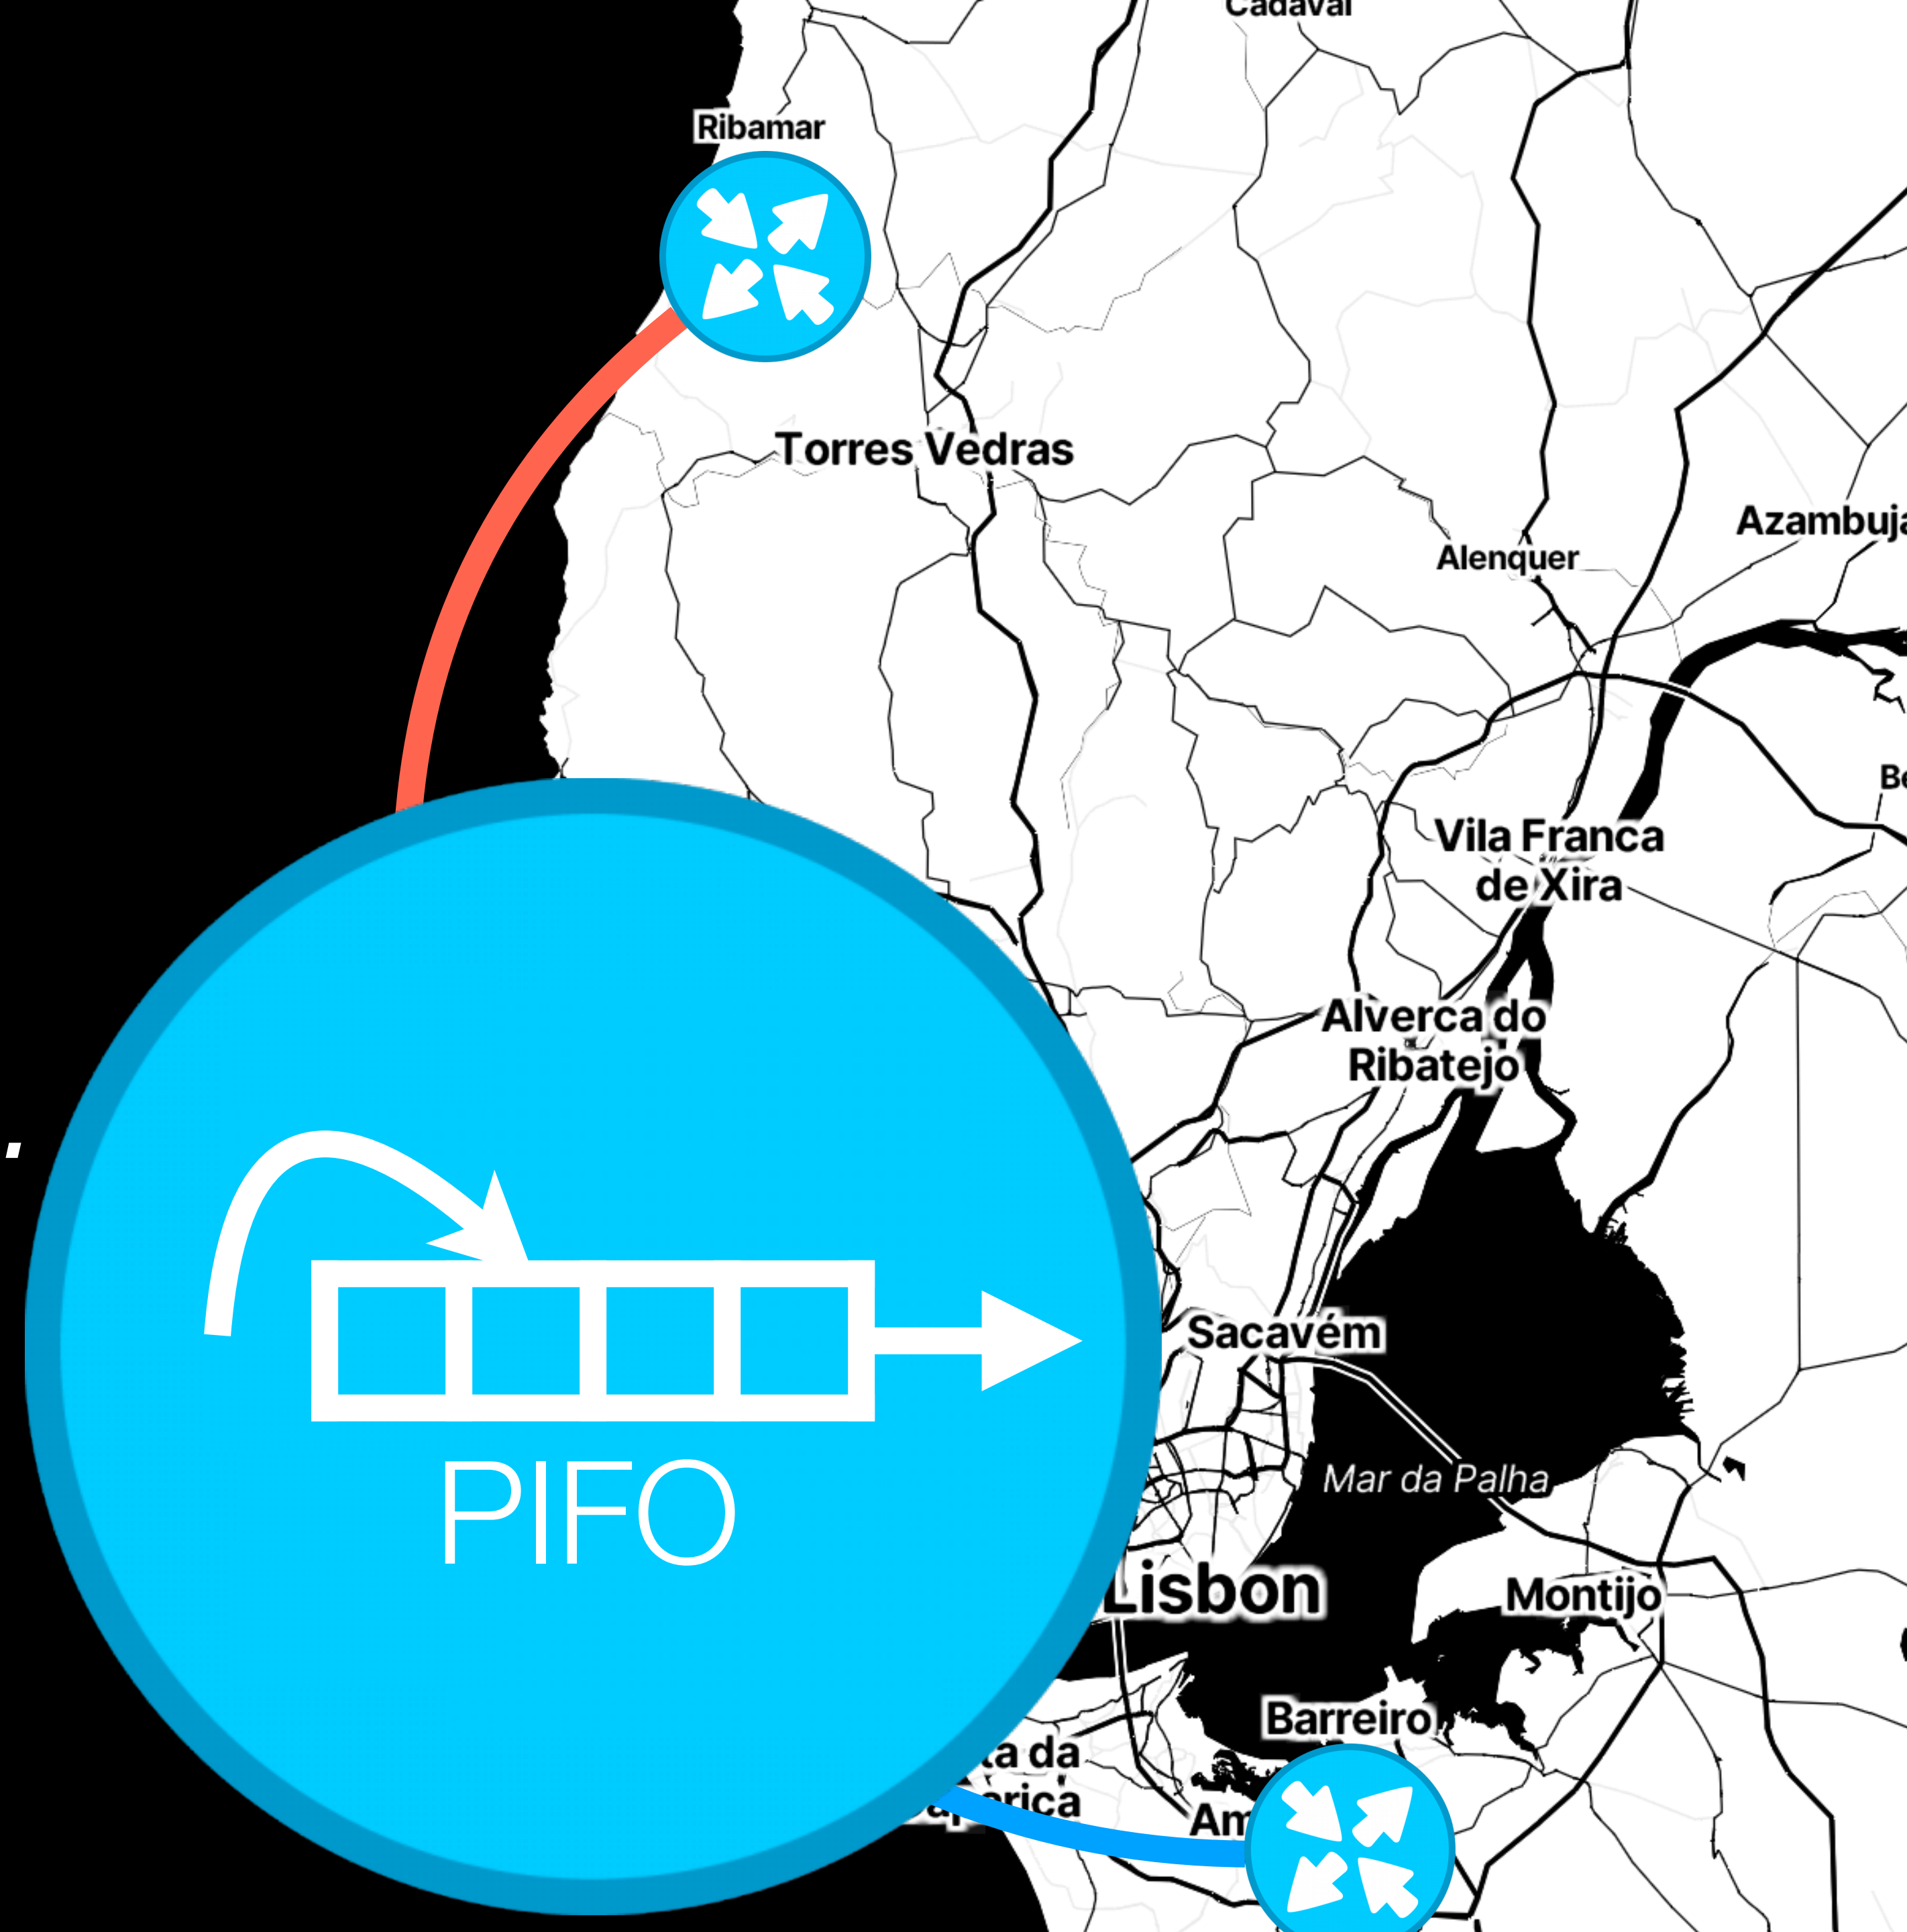


SDN made networks programmable.

Early goal: routing.

But now we need control over *scheduling*.

Basic tools work fine...




But modern scheduling requires more.



But modern scheduling requires more.

R traffic goes to either Porto or Torres Vedras.




But modern scheduling requires more.

R traffic goes to either **P** or **T**.

Goal:



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



But modern scheduling requires more.

R traffic goes to either Porto or Torres Vedras.

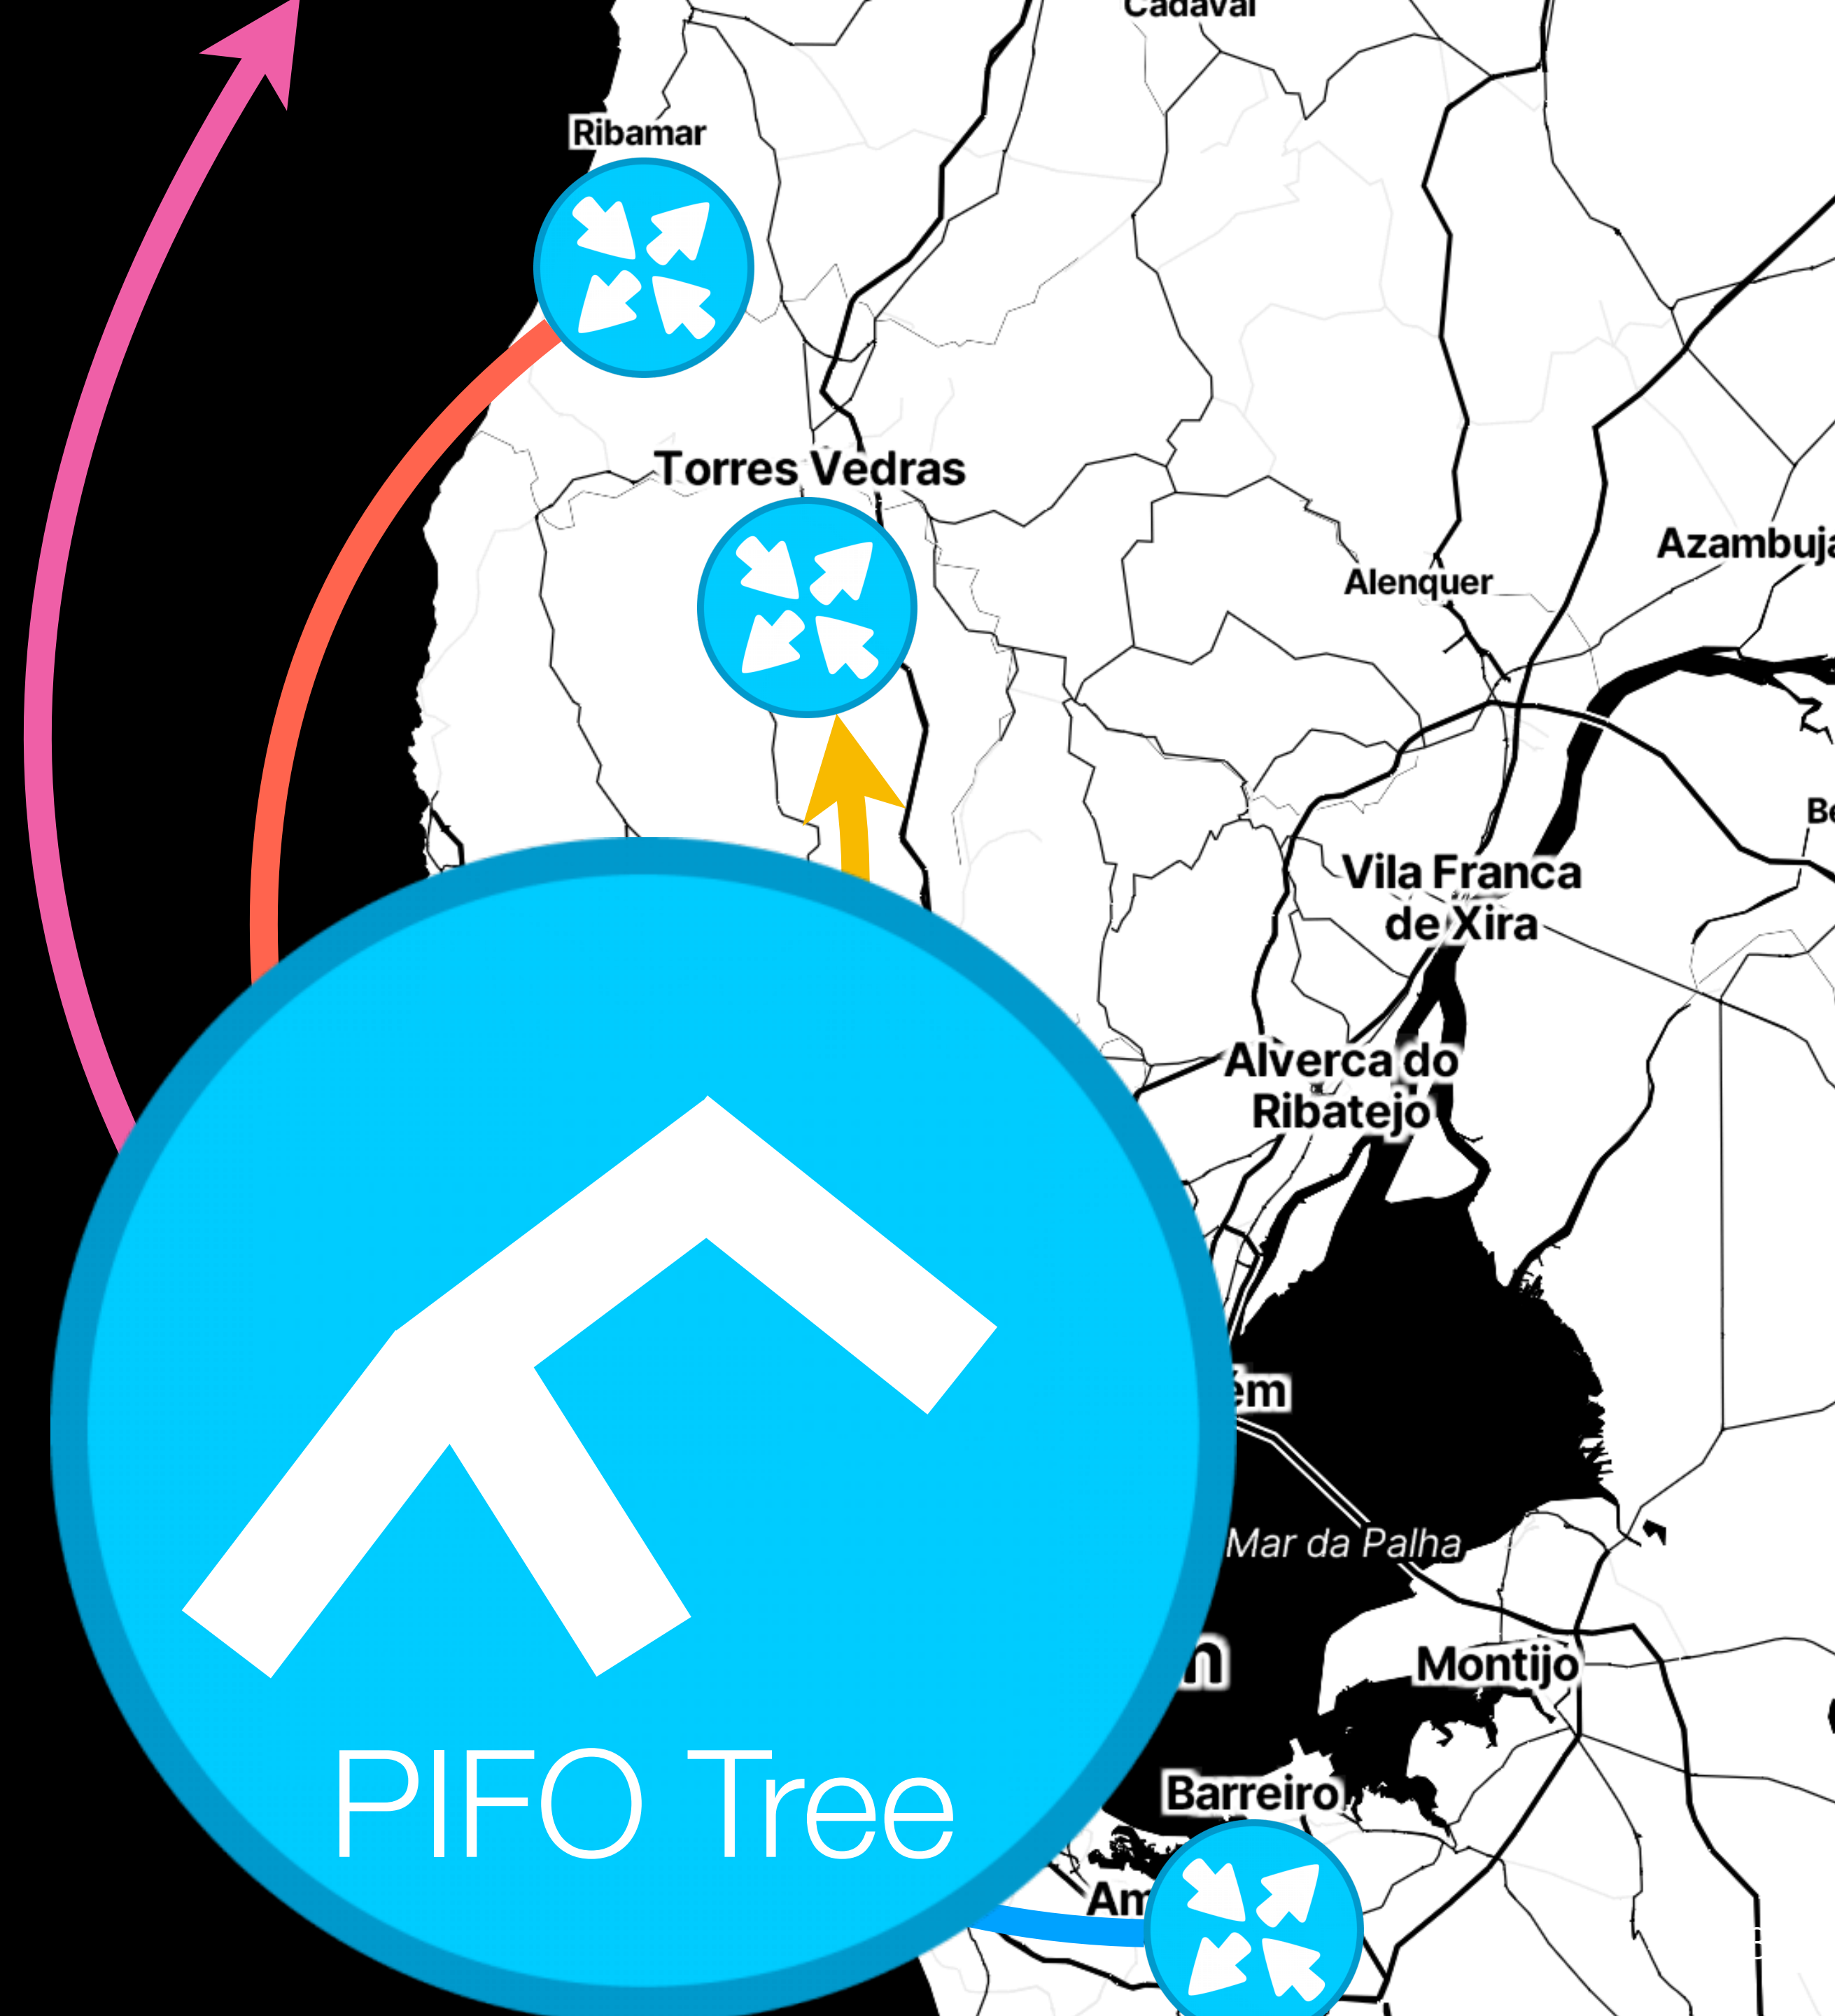

Goal:



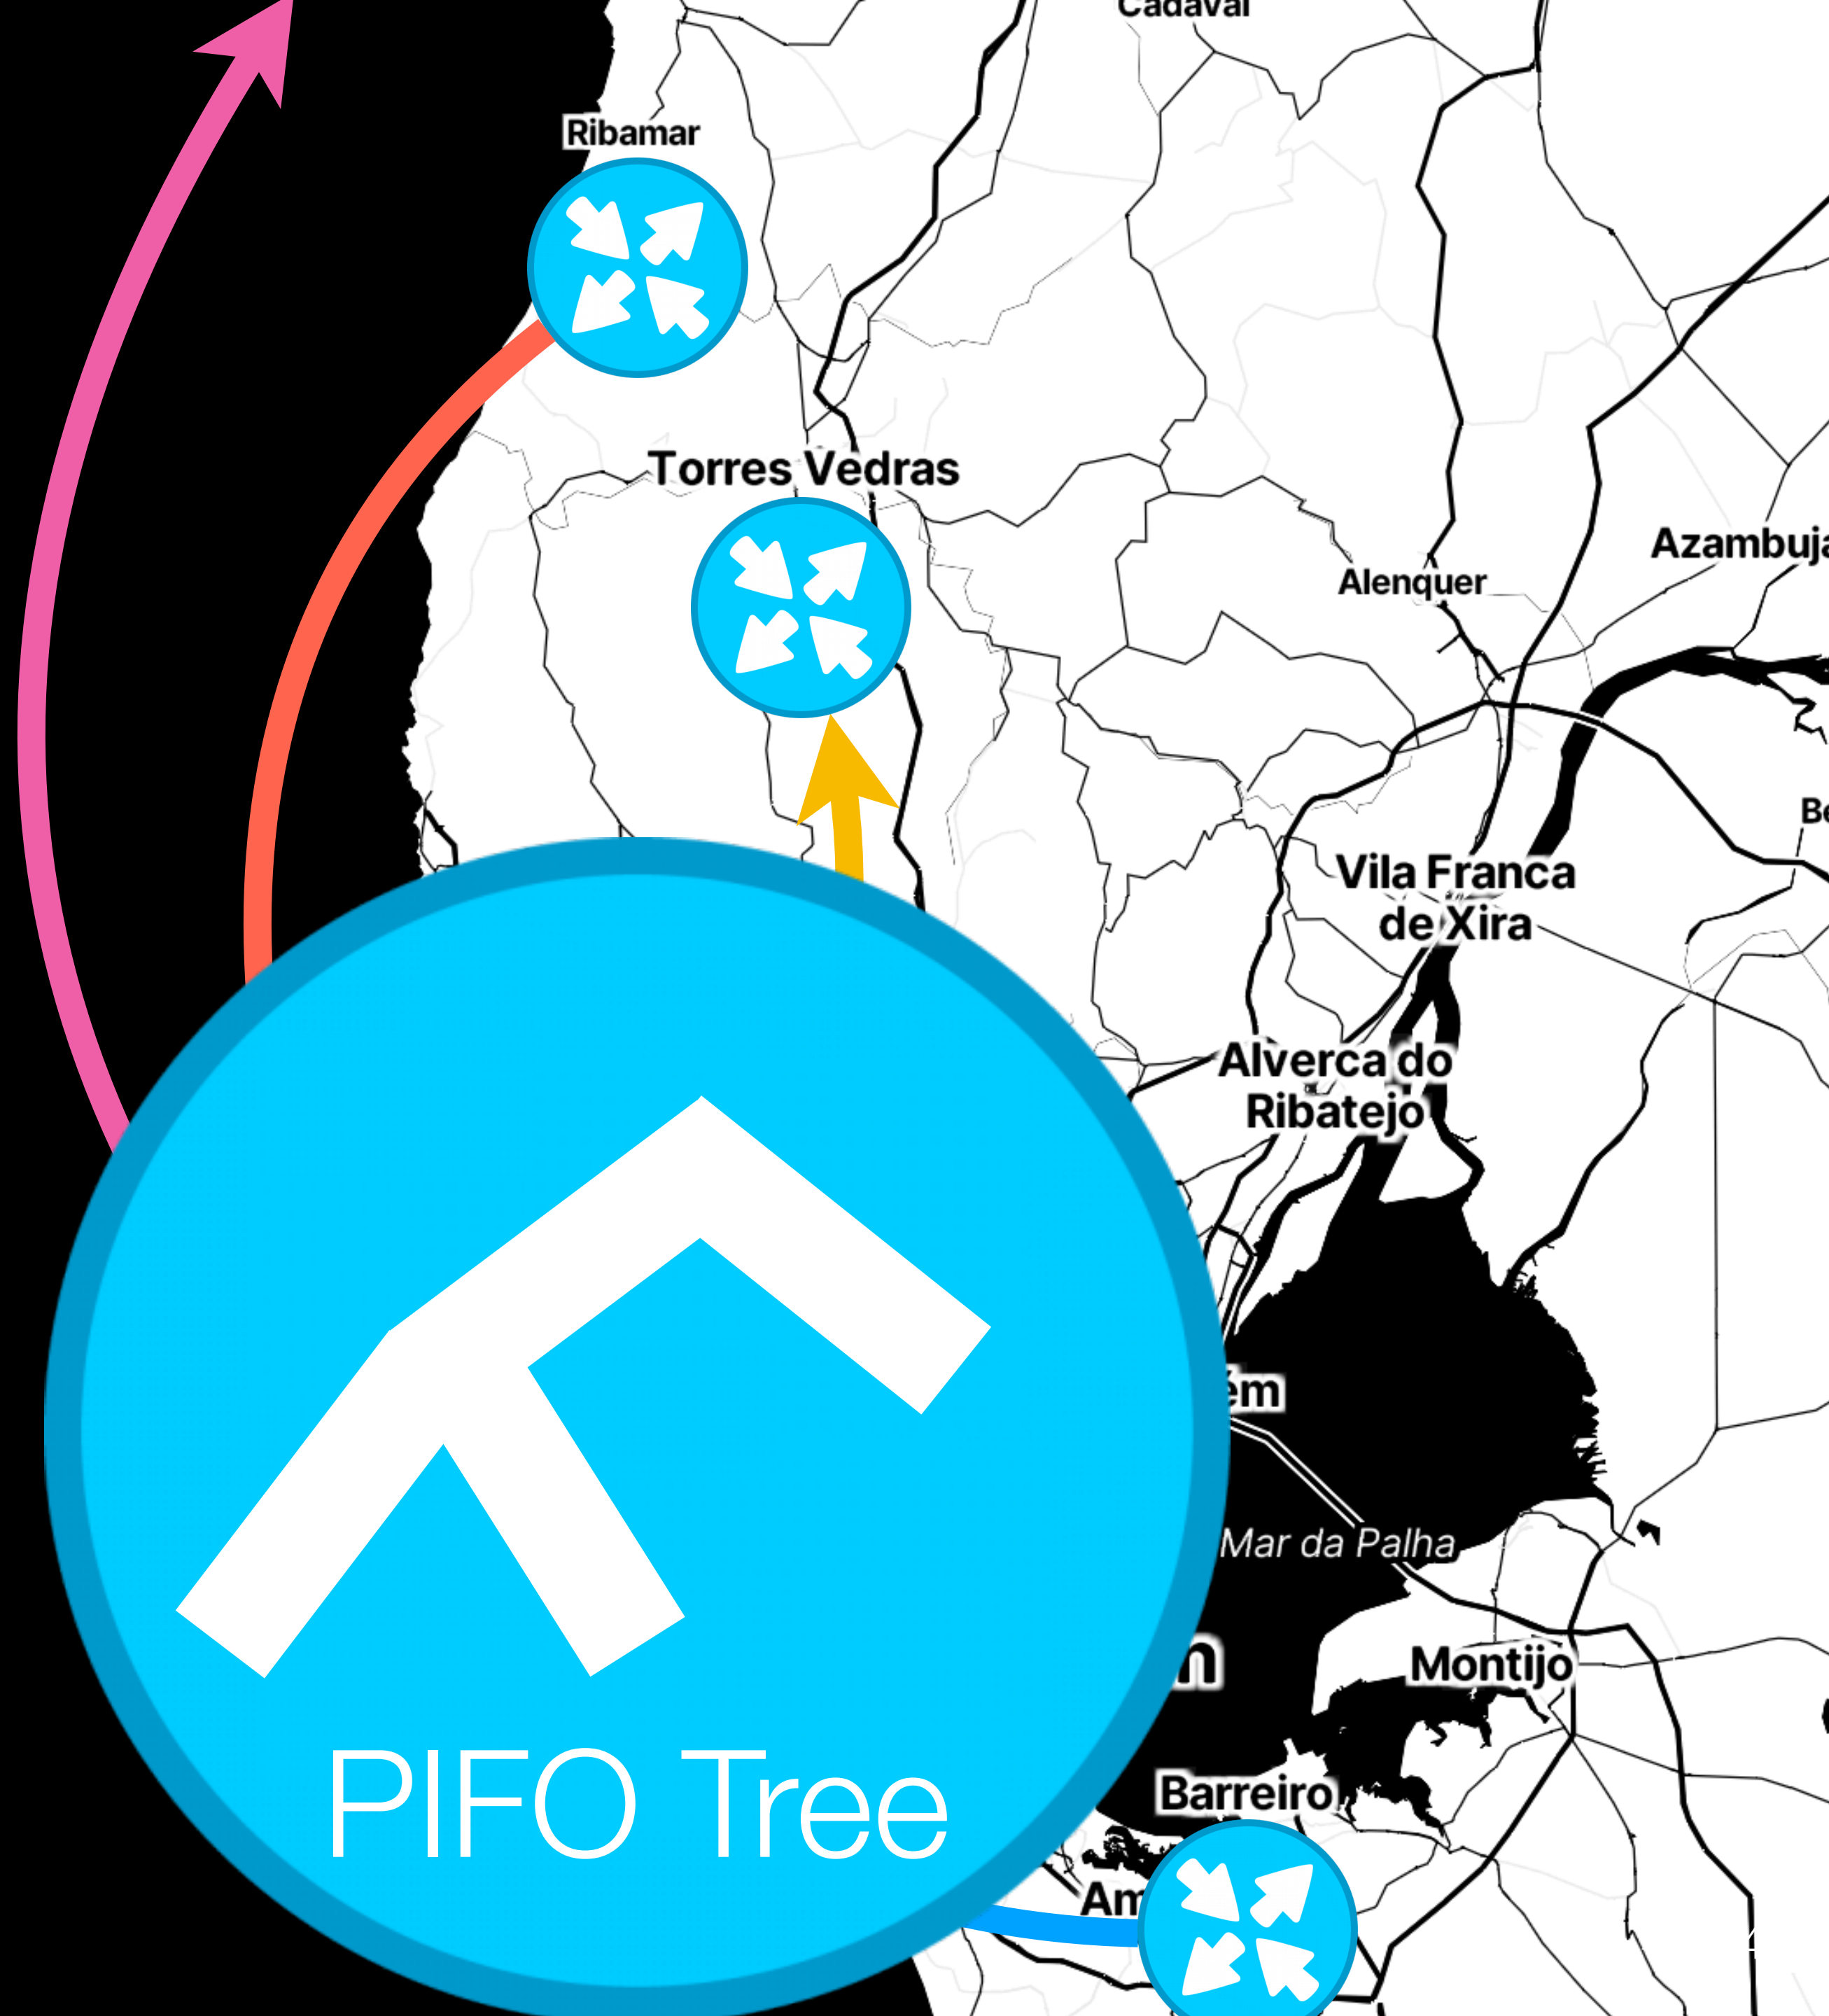
But modern scheduling requires more.

R traffic goes to either Porto or Torres Vedras.

Goal:




New plan!




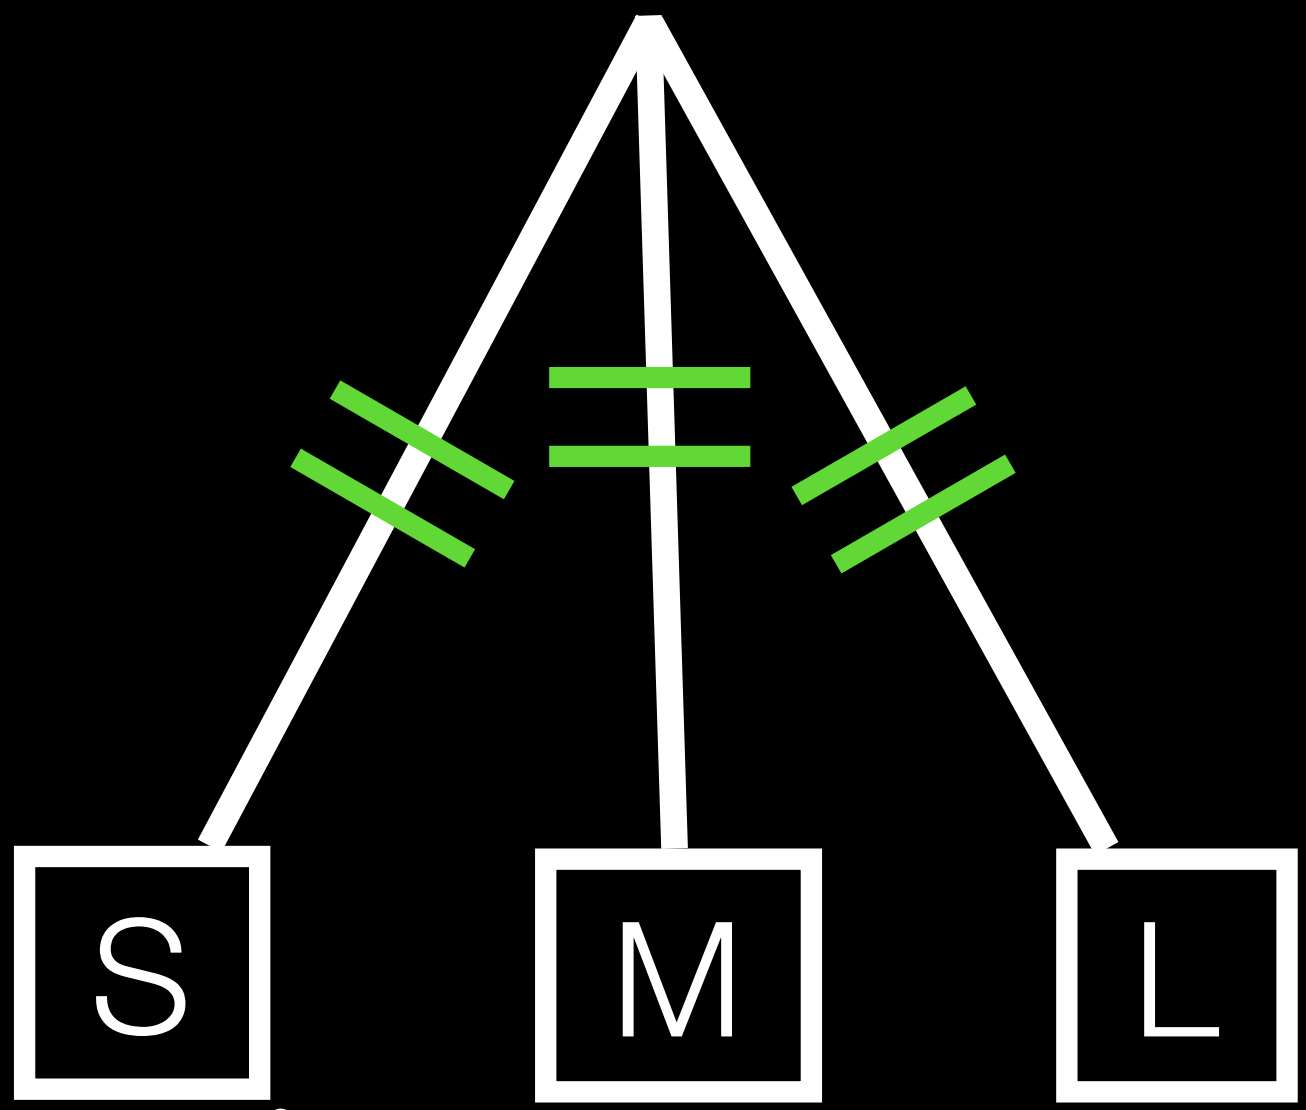
New plan!
Interleave
small, medium, and large
packets.



New plan!
Interleave
small, medium, and large
packets.



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.

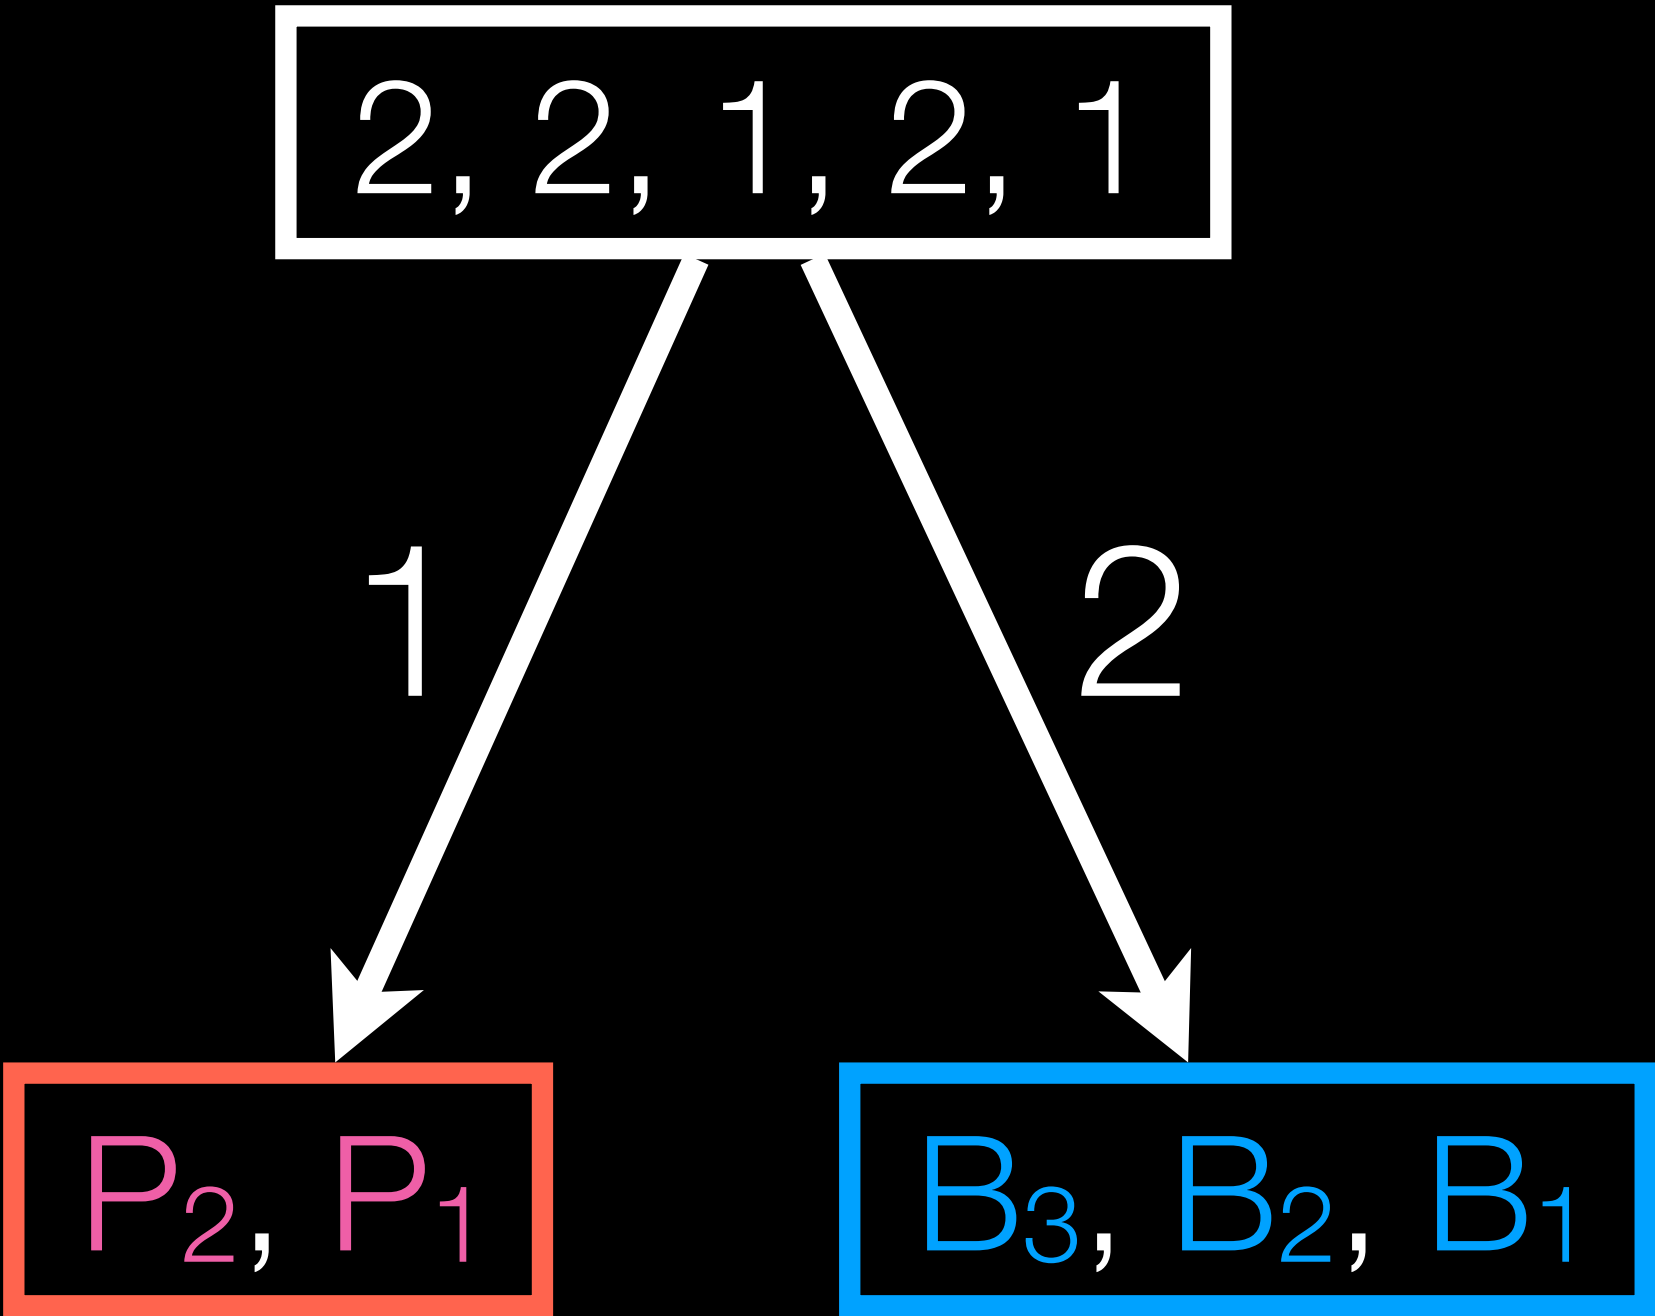


A human needs a *range* of trees.

The hardware wants to support *one* tree.


Aside: PIFO trees

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



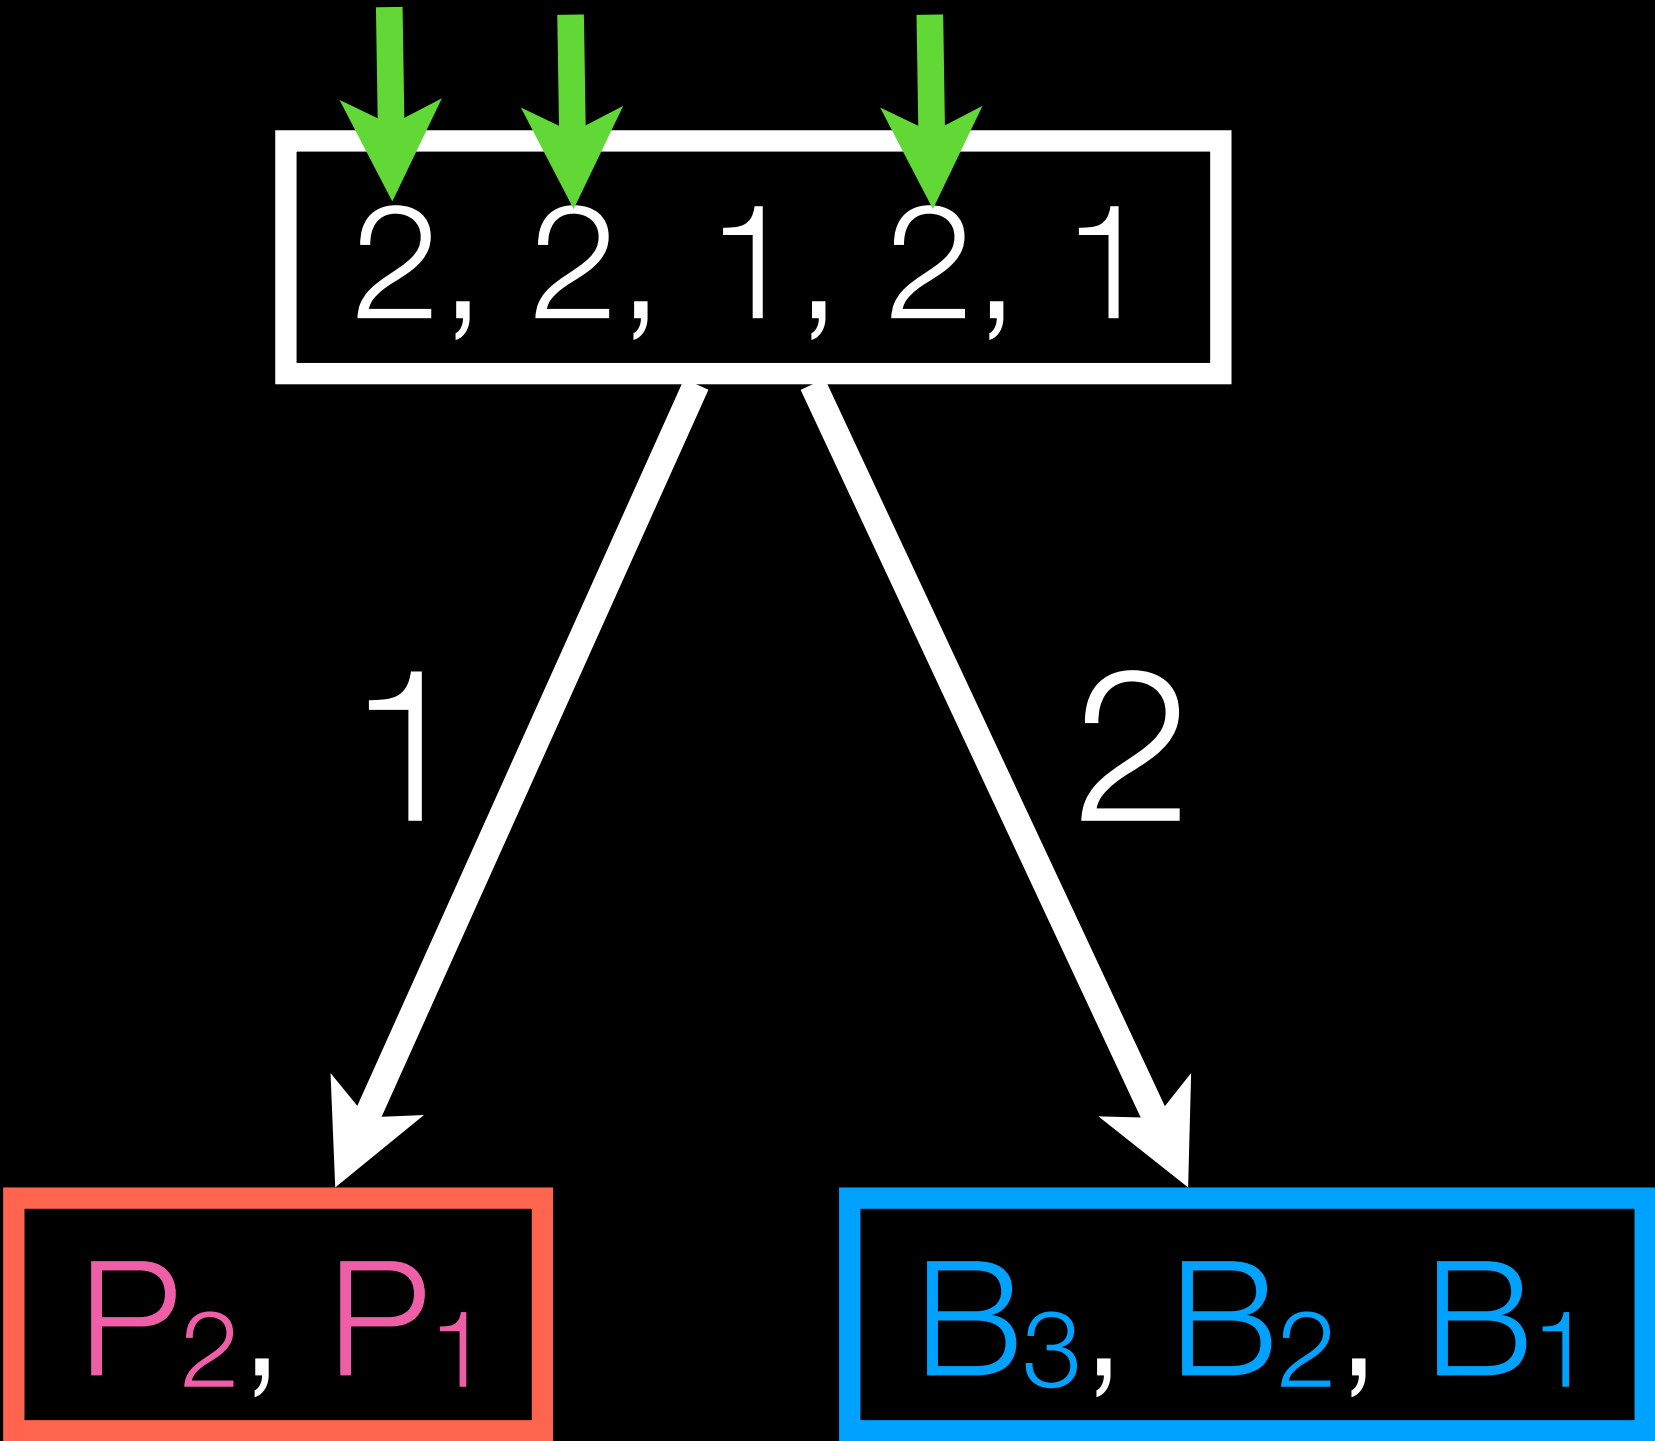
Aside: PIFO trees

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




Aside: PIFO trees

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



Aside: PIFO trees


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



This behaves like a queue!

Aside: PIFO trees


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



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

Aside: PIFO trees


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



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

Aside: PIFO trees


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



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

Aside: PIFO trees


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



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

Aside: PIFO trees

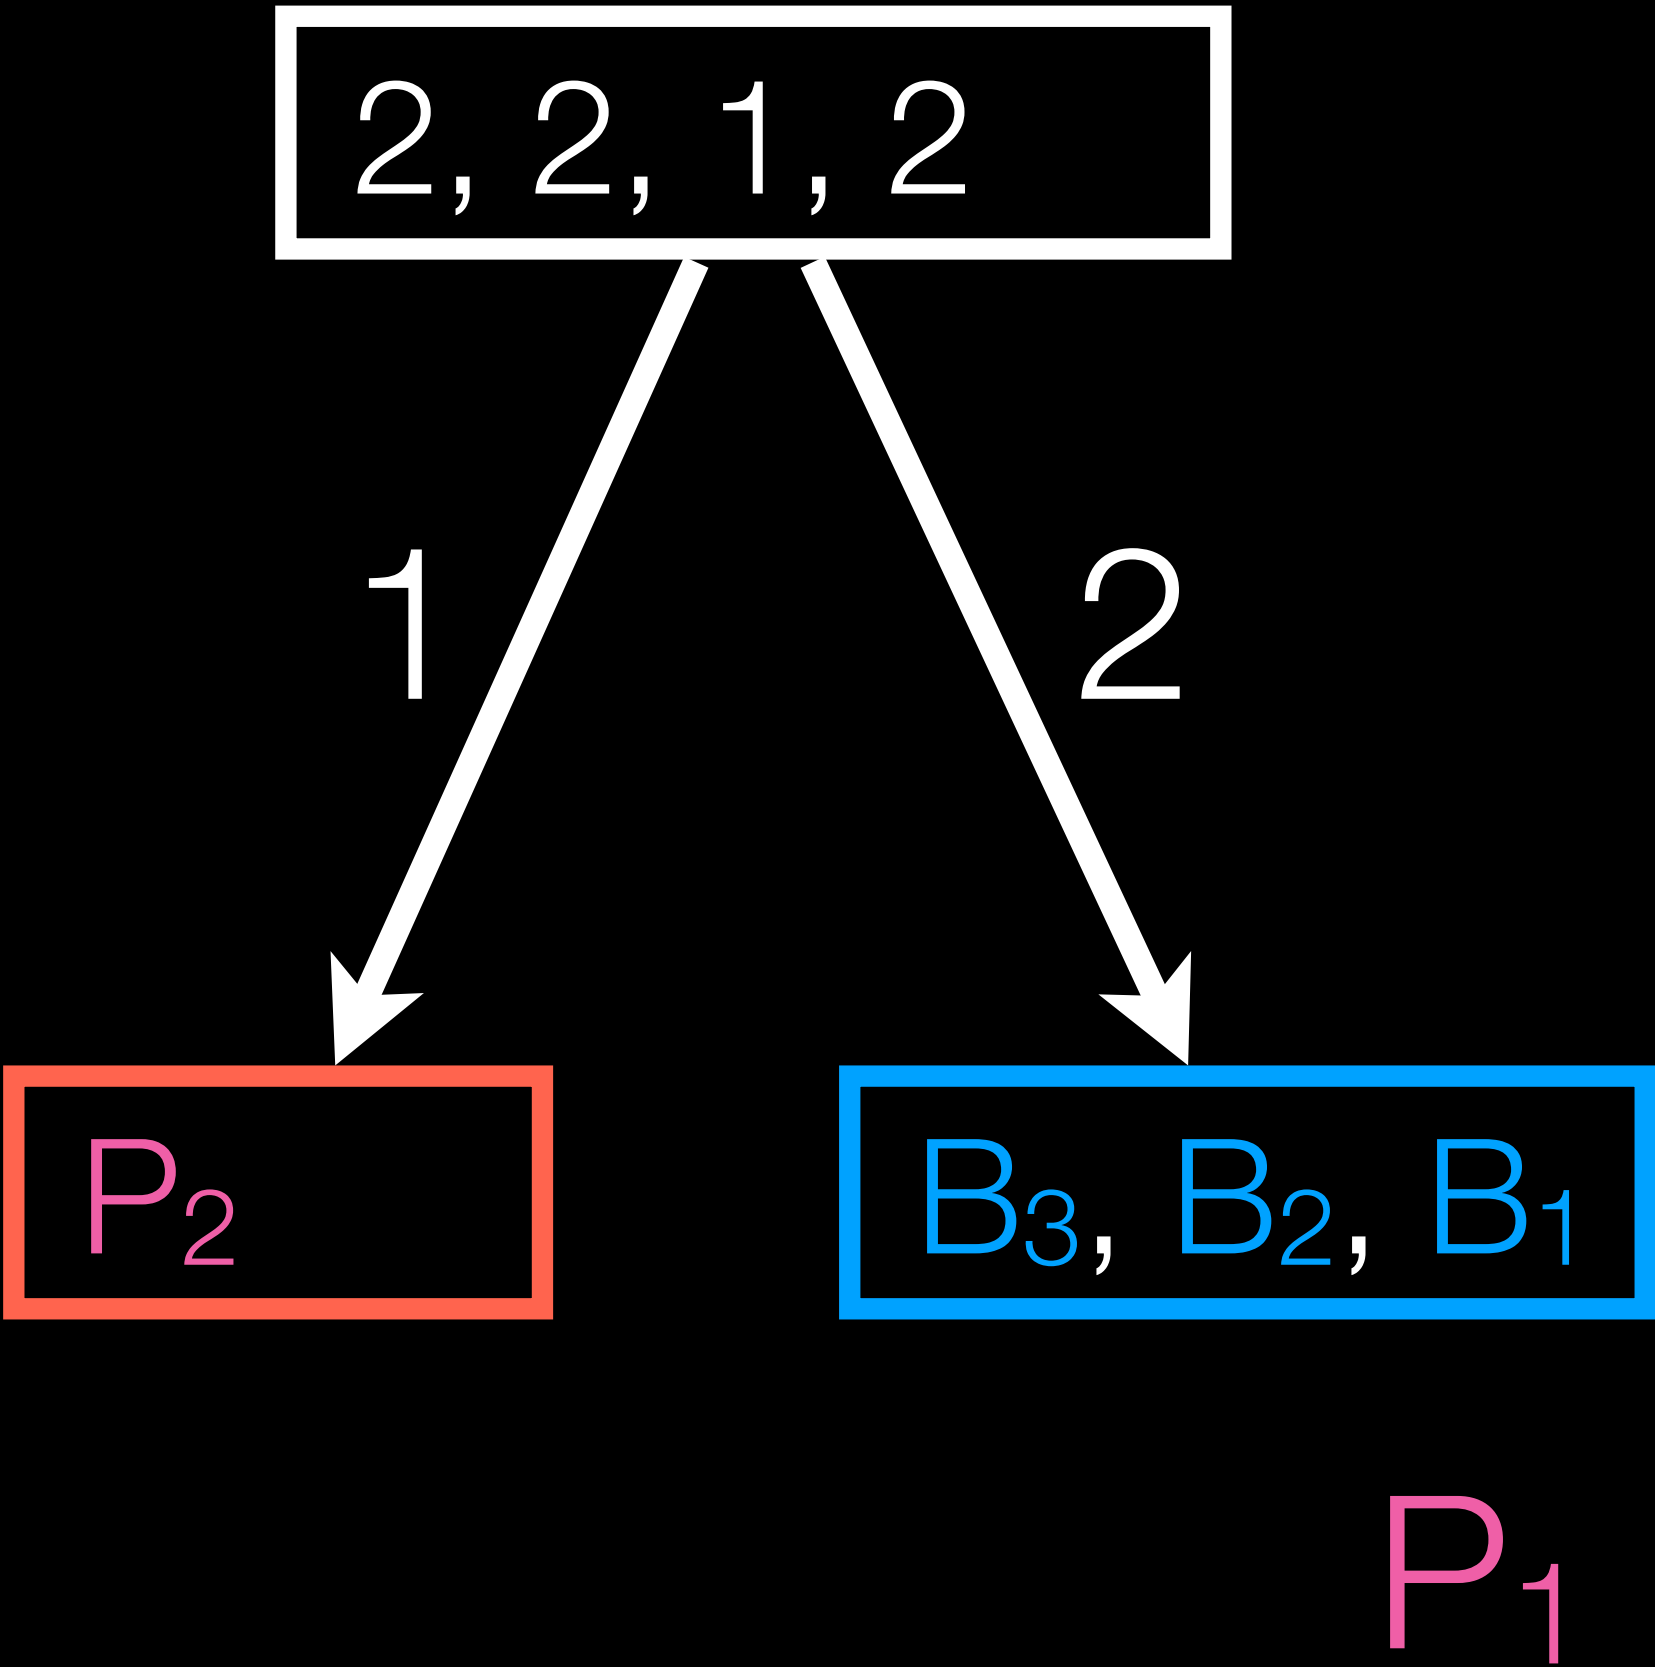
interleave **R** and **B**;
interleave **P** and **T**.



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

Aside: PIFO trees


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



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

Aside: PIFO trees

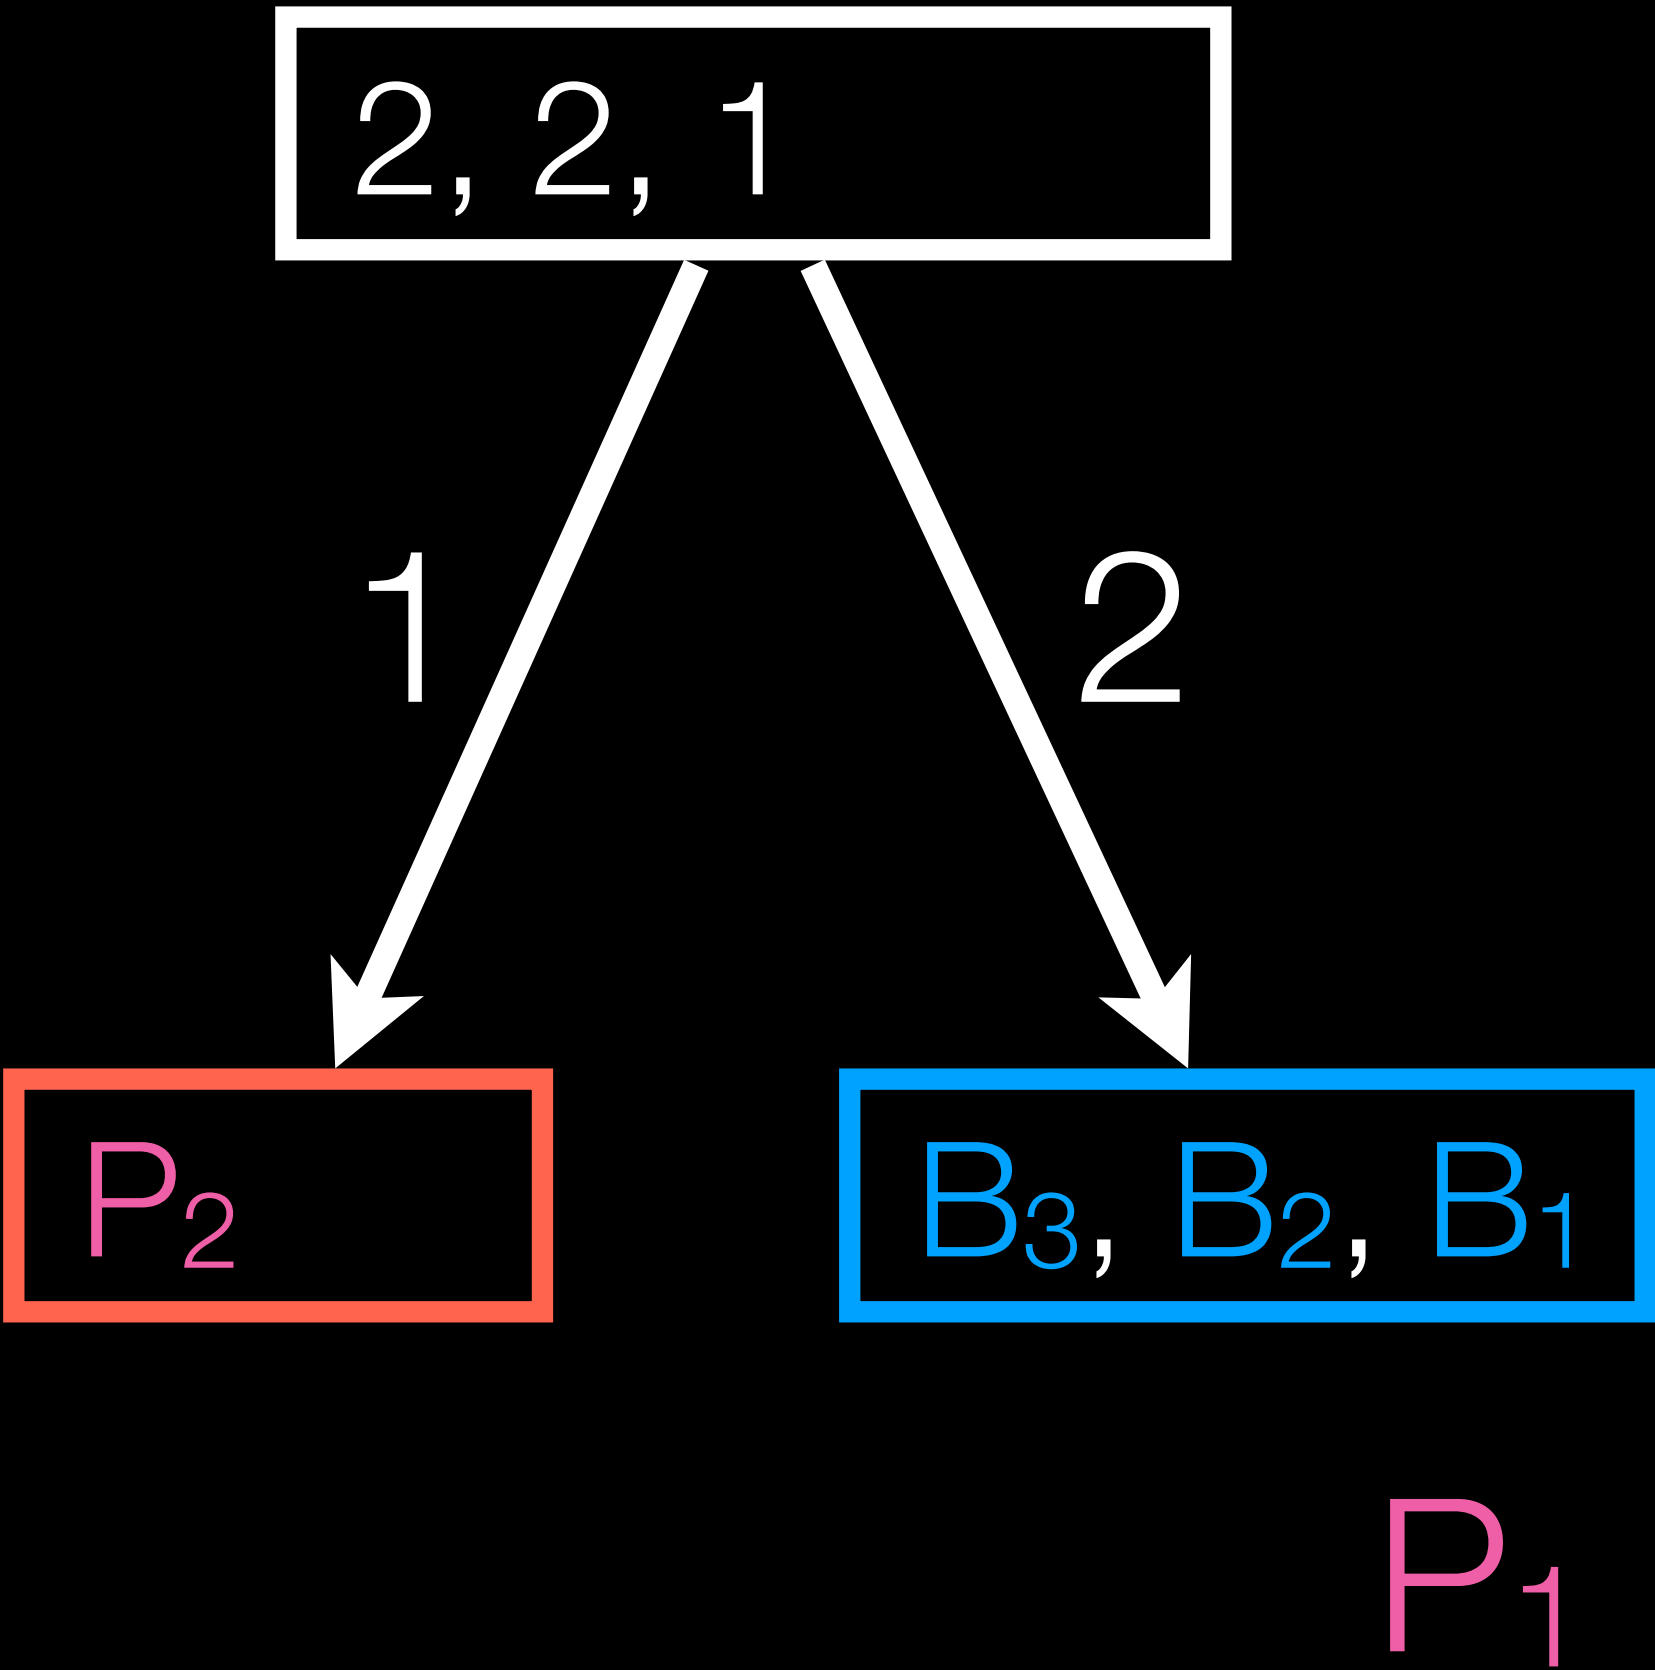
interleave **R** and **B**;
interleave **P** and **T**.



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

Aside: PIFO trees


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



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

Aside: PIFO trees

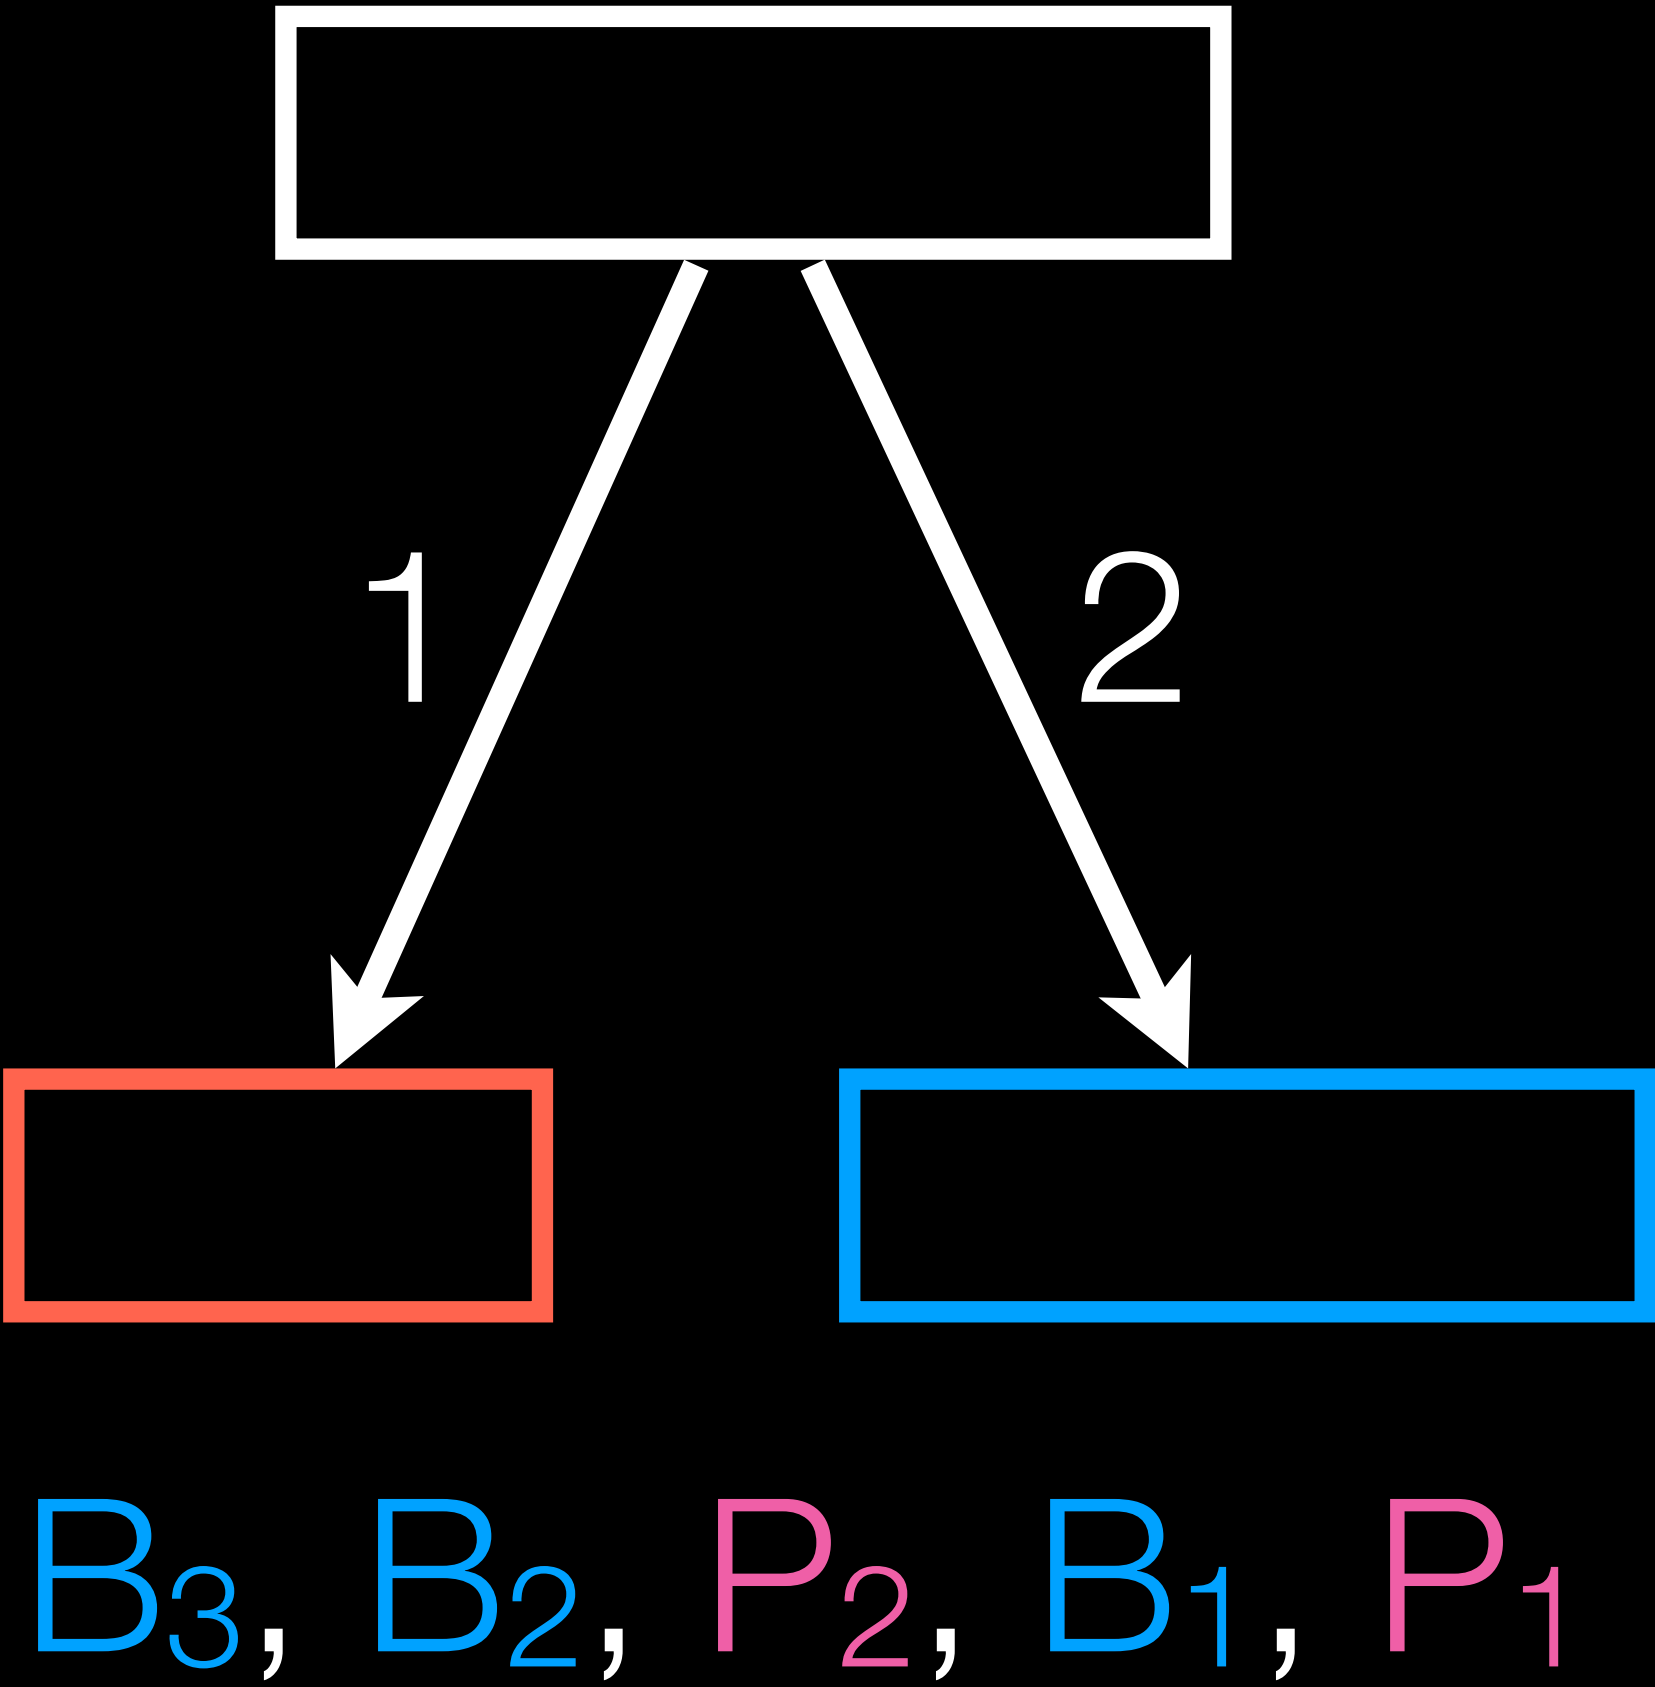
interleave **R** and **B**;
interleave **P** and **T**.



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

Aside: PIFO trees


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



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

Aside: PIFO trees


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



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

Aside: PIFO trees

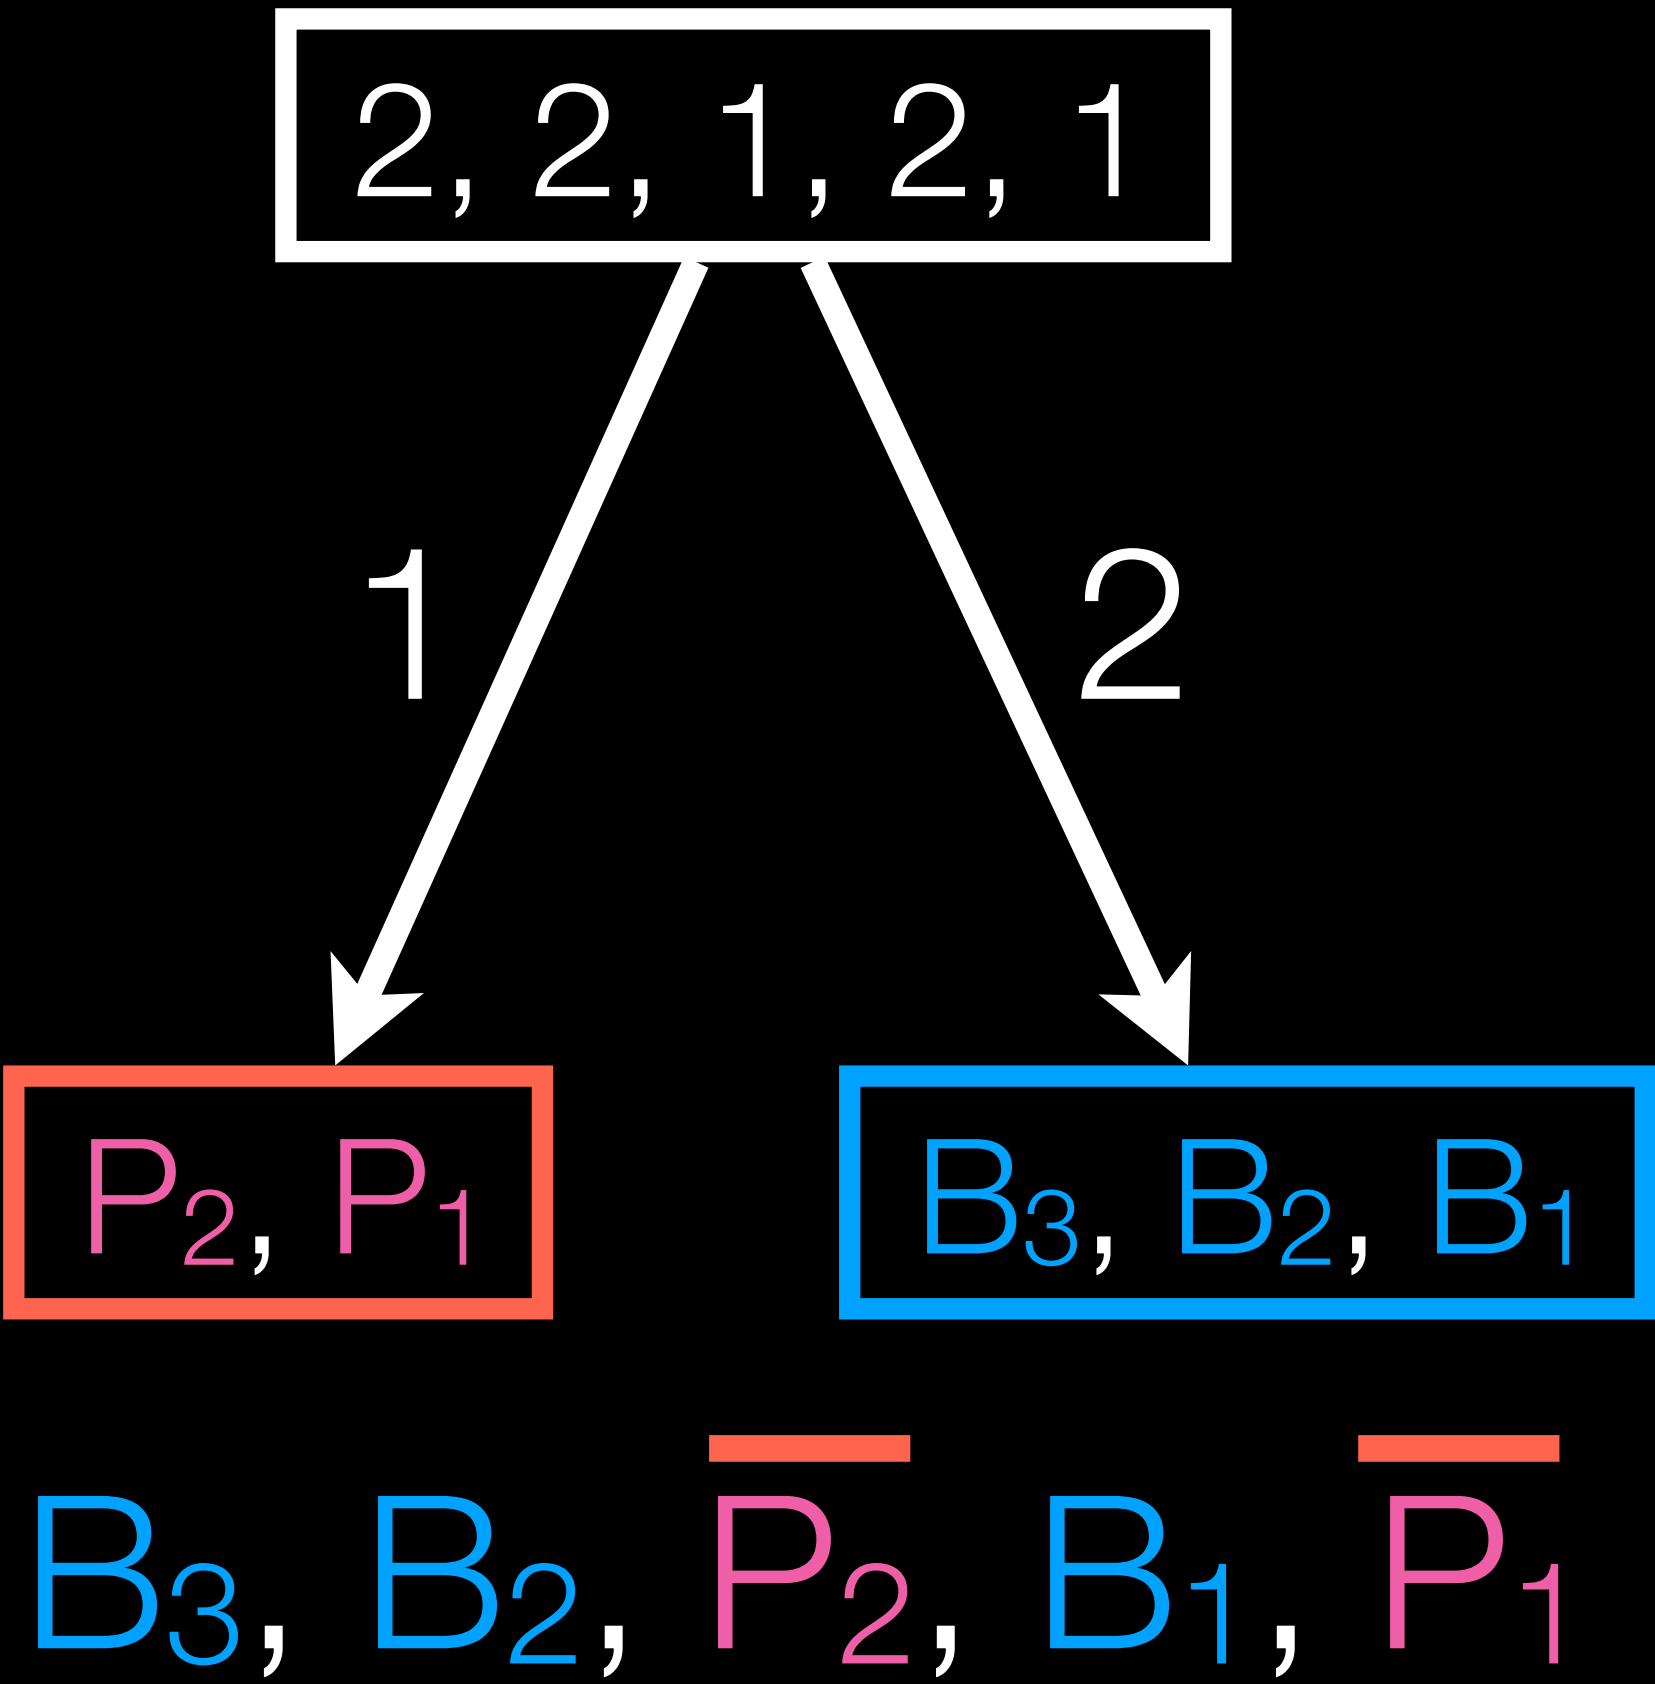
interleave **R** and **B**;
interleave **P** and **T**.



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

Aside: PIFO trees

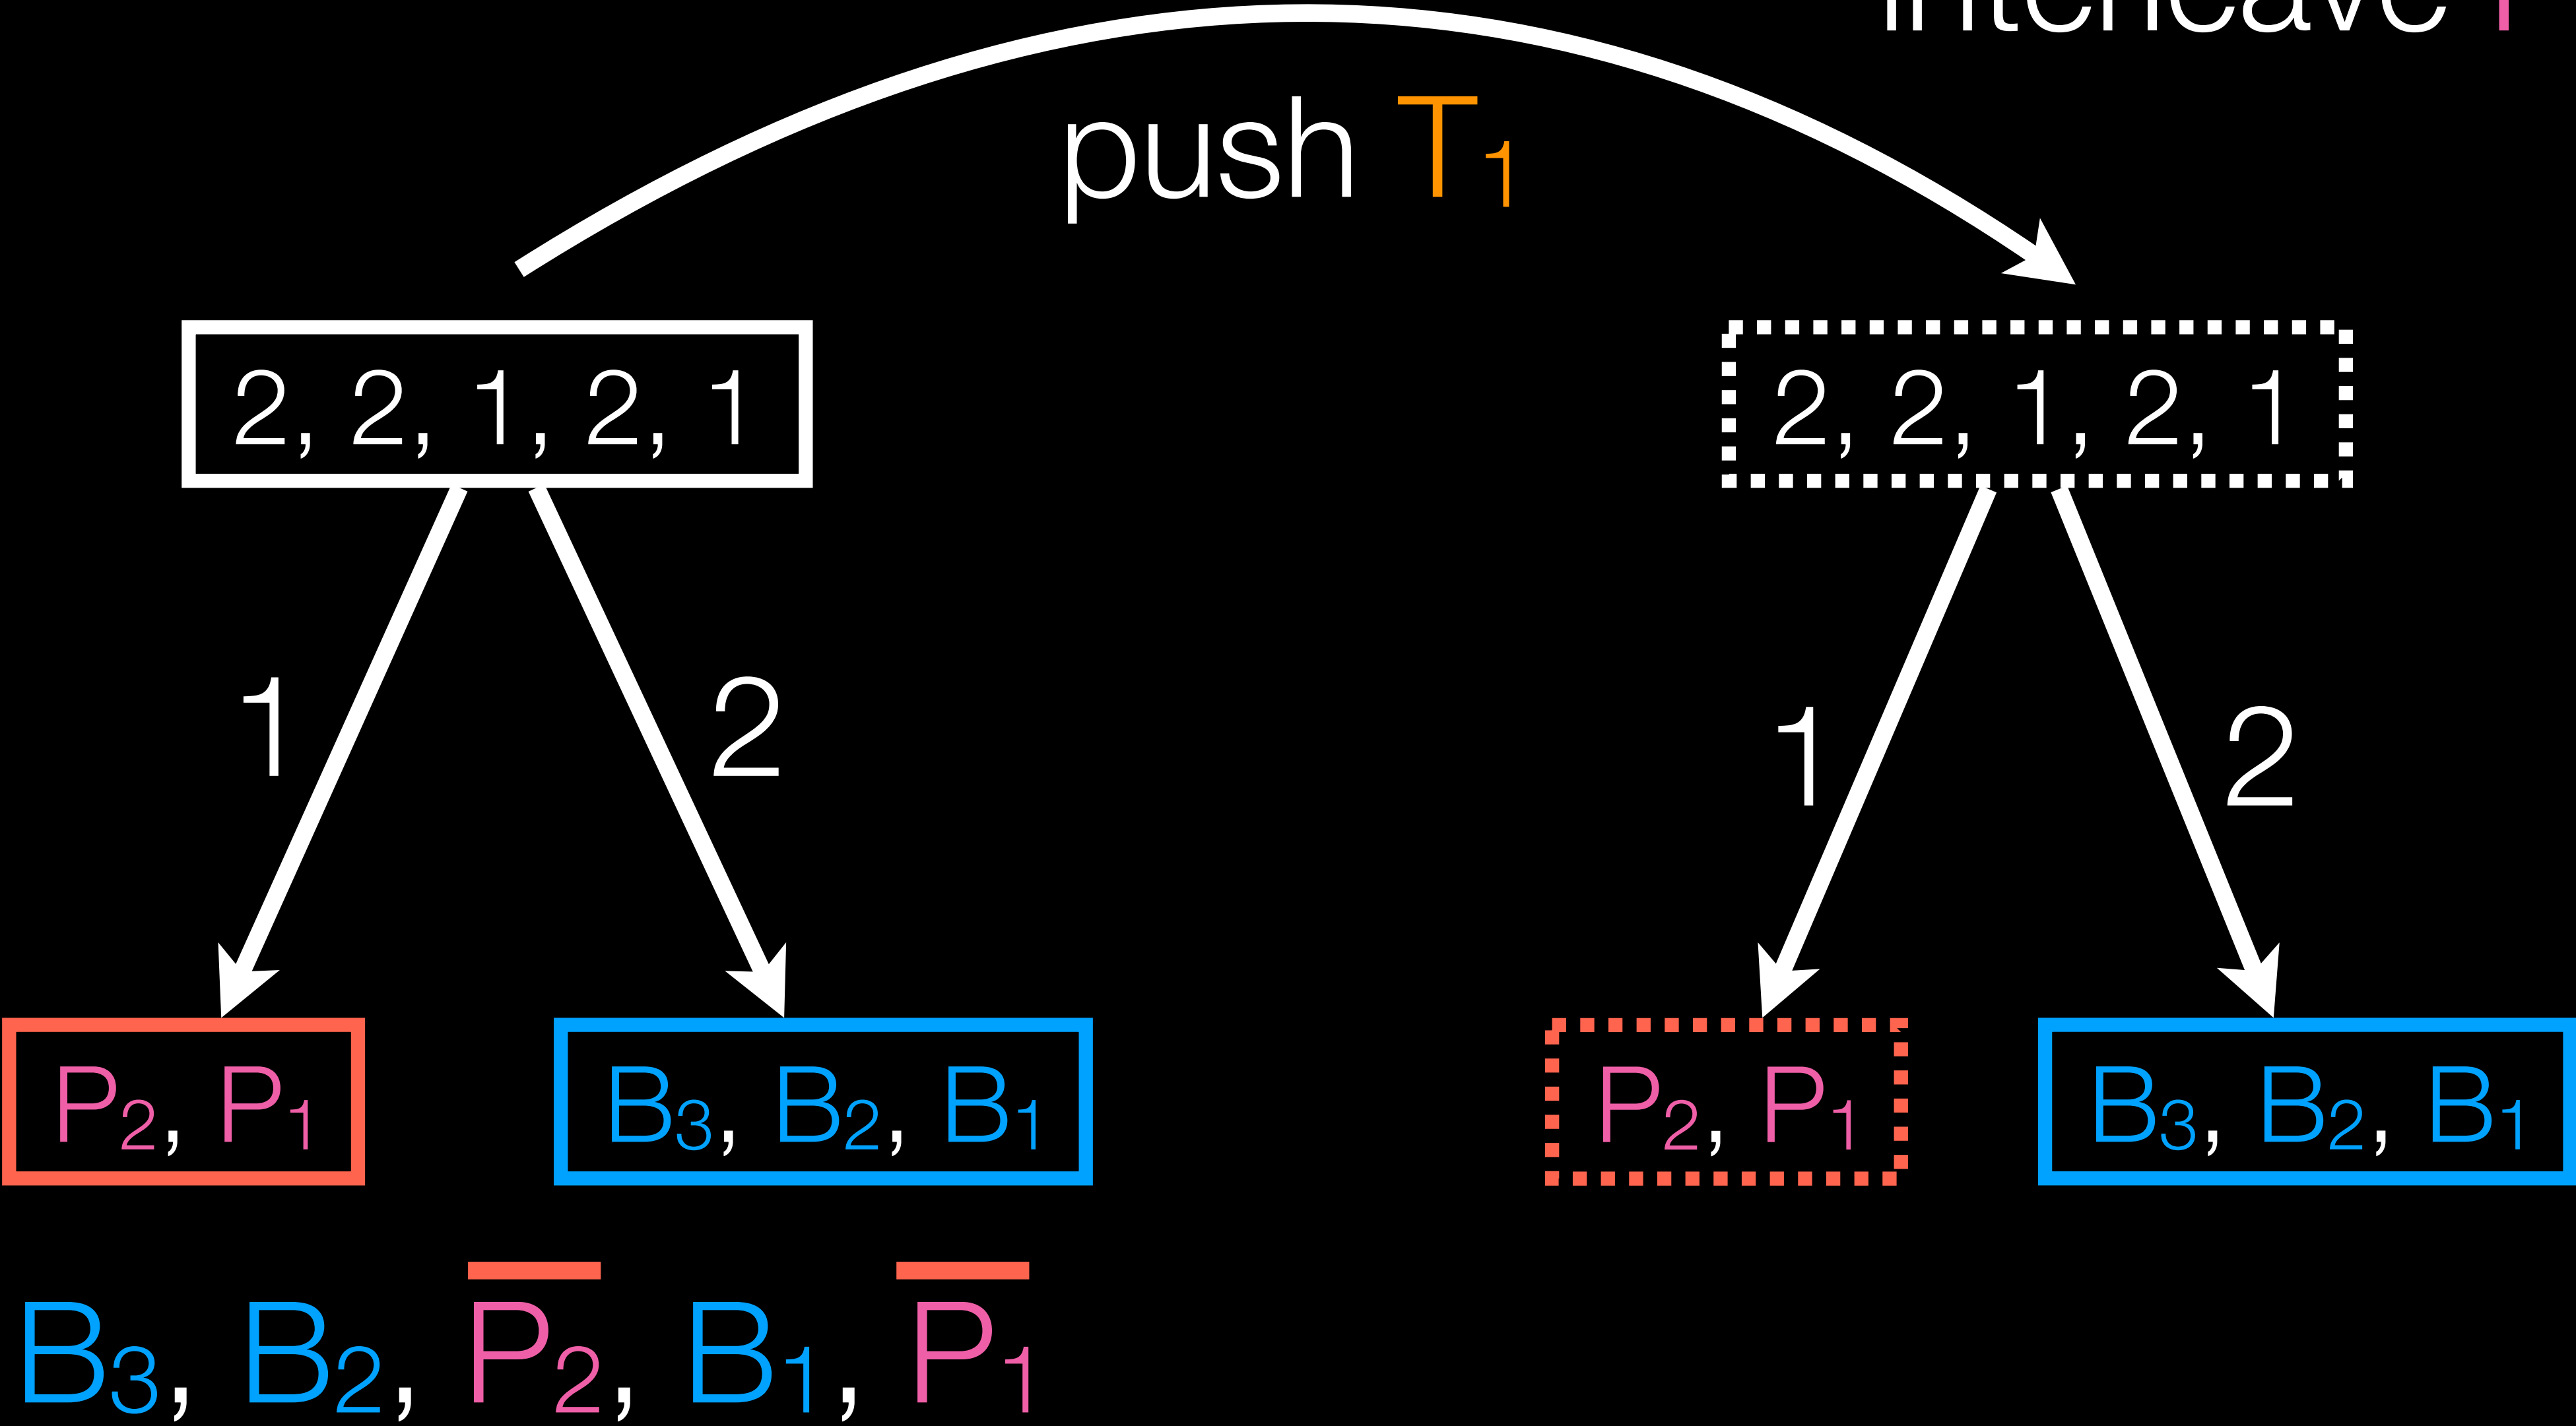
interleave **R** and **B**;
interleave **P** and **T**.



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

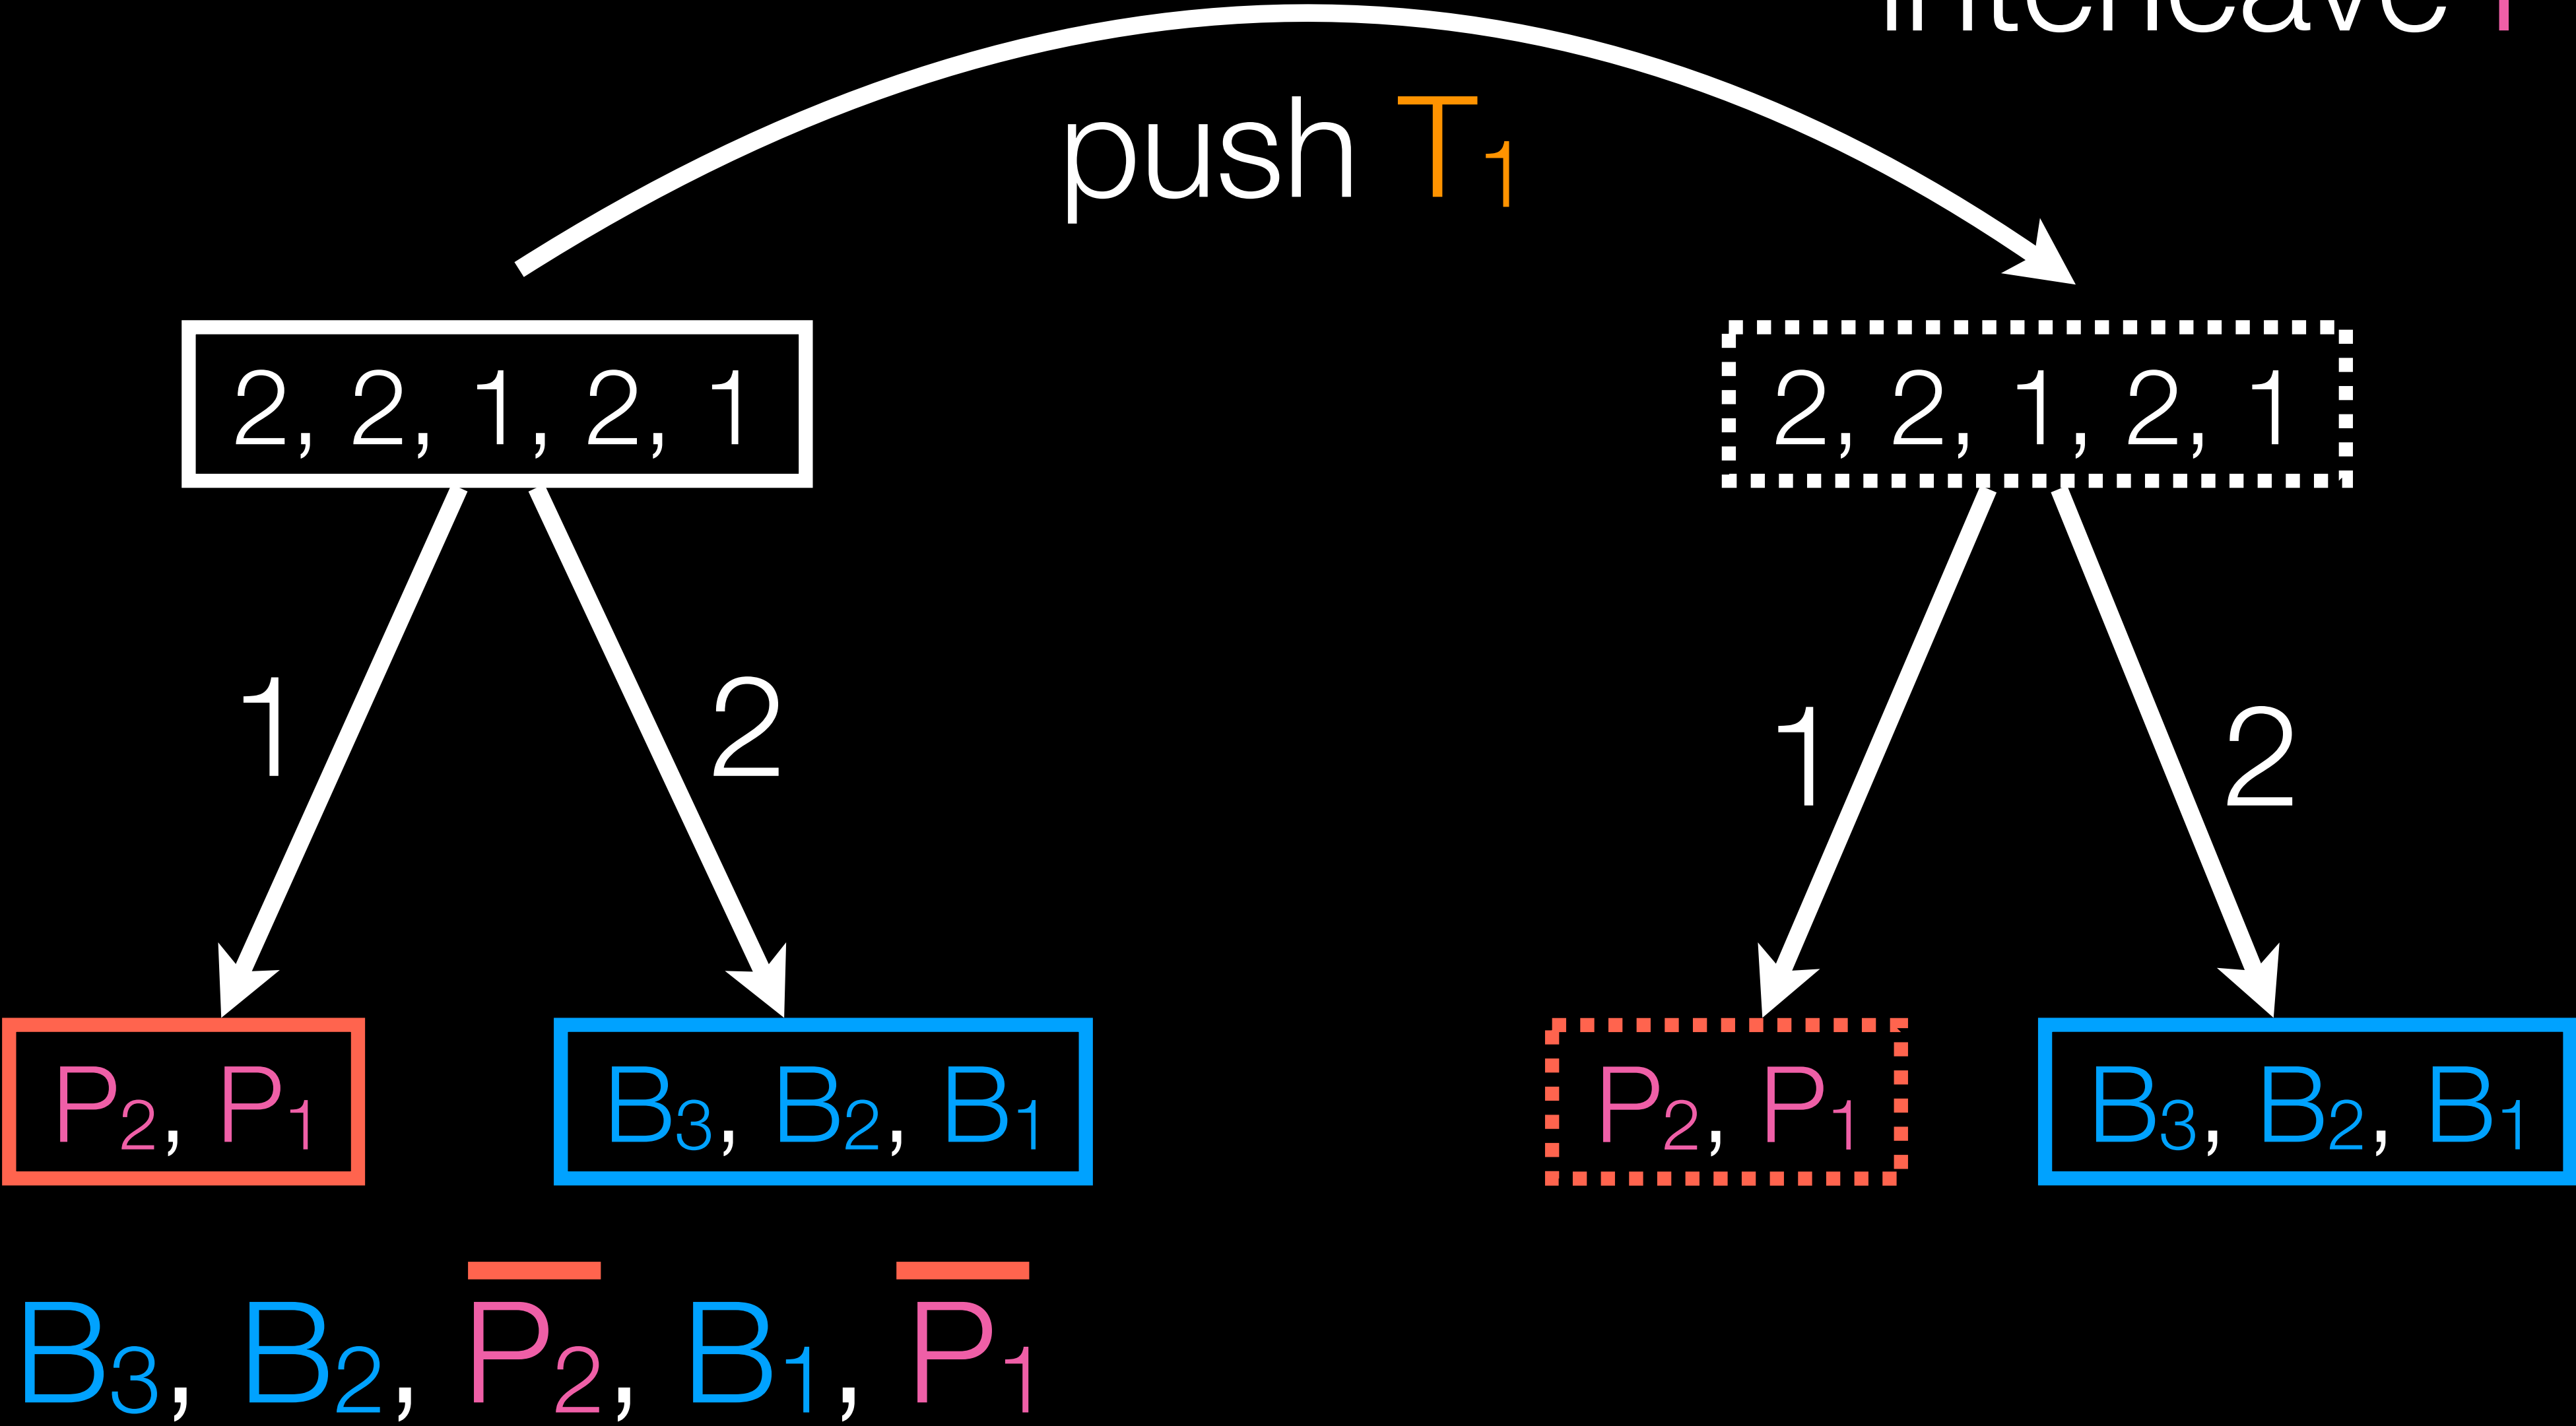
Aside: PIFO trees

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



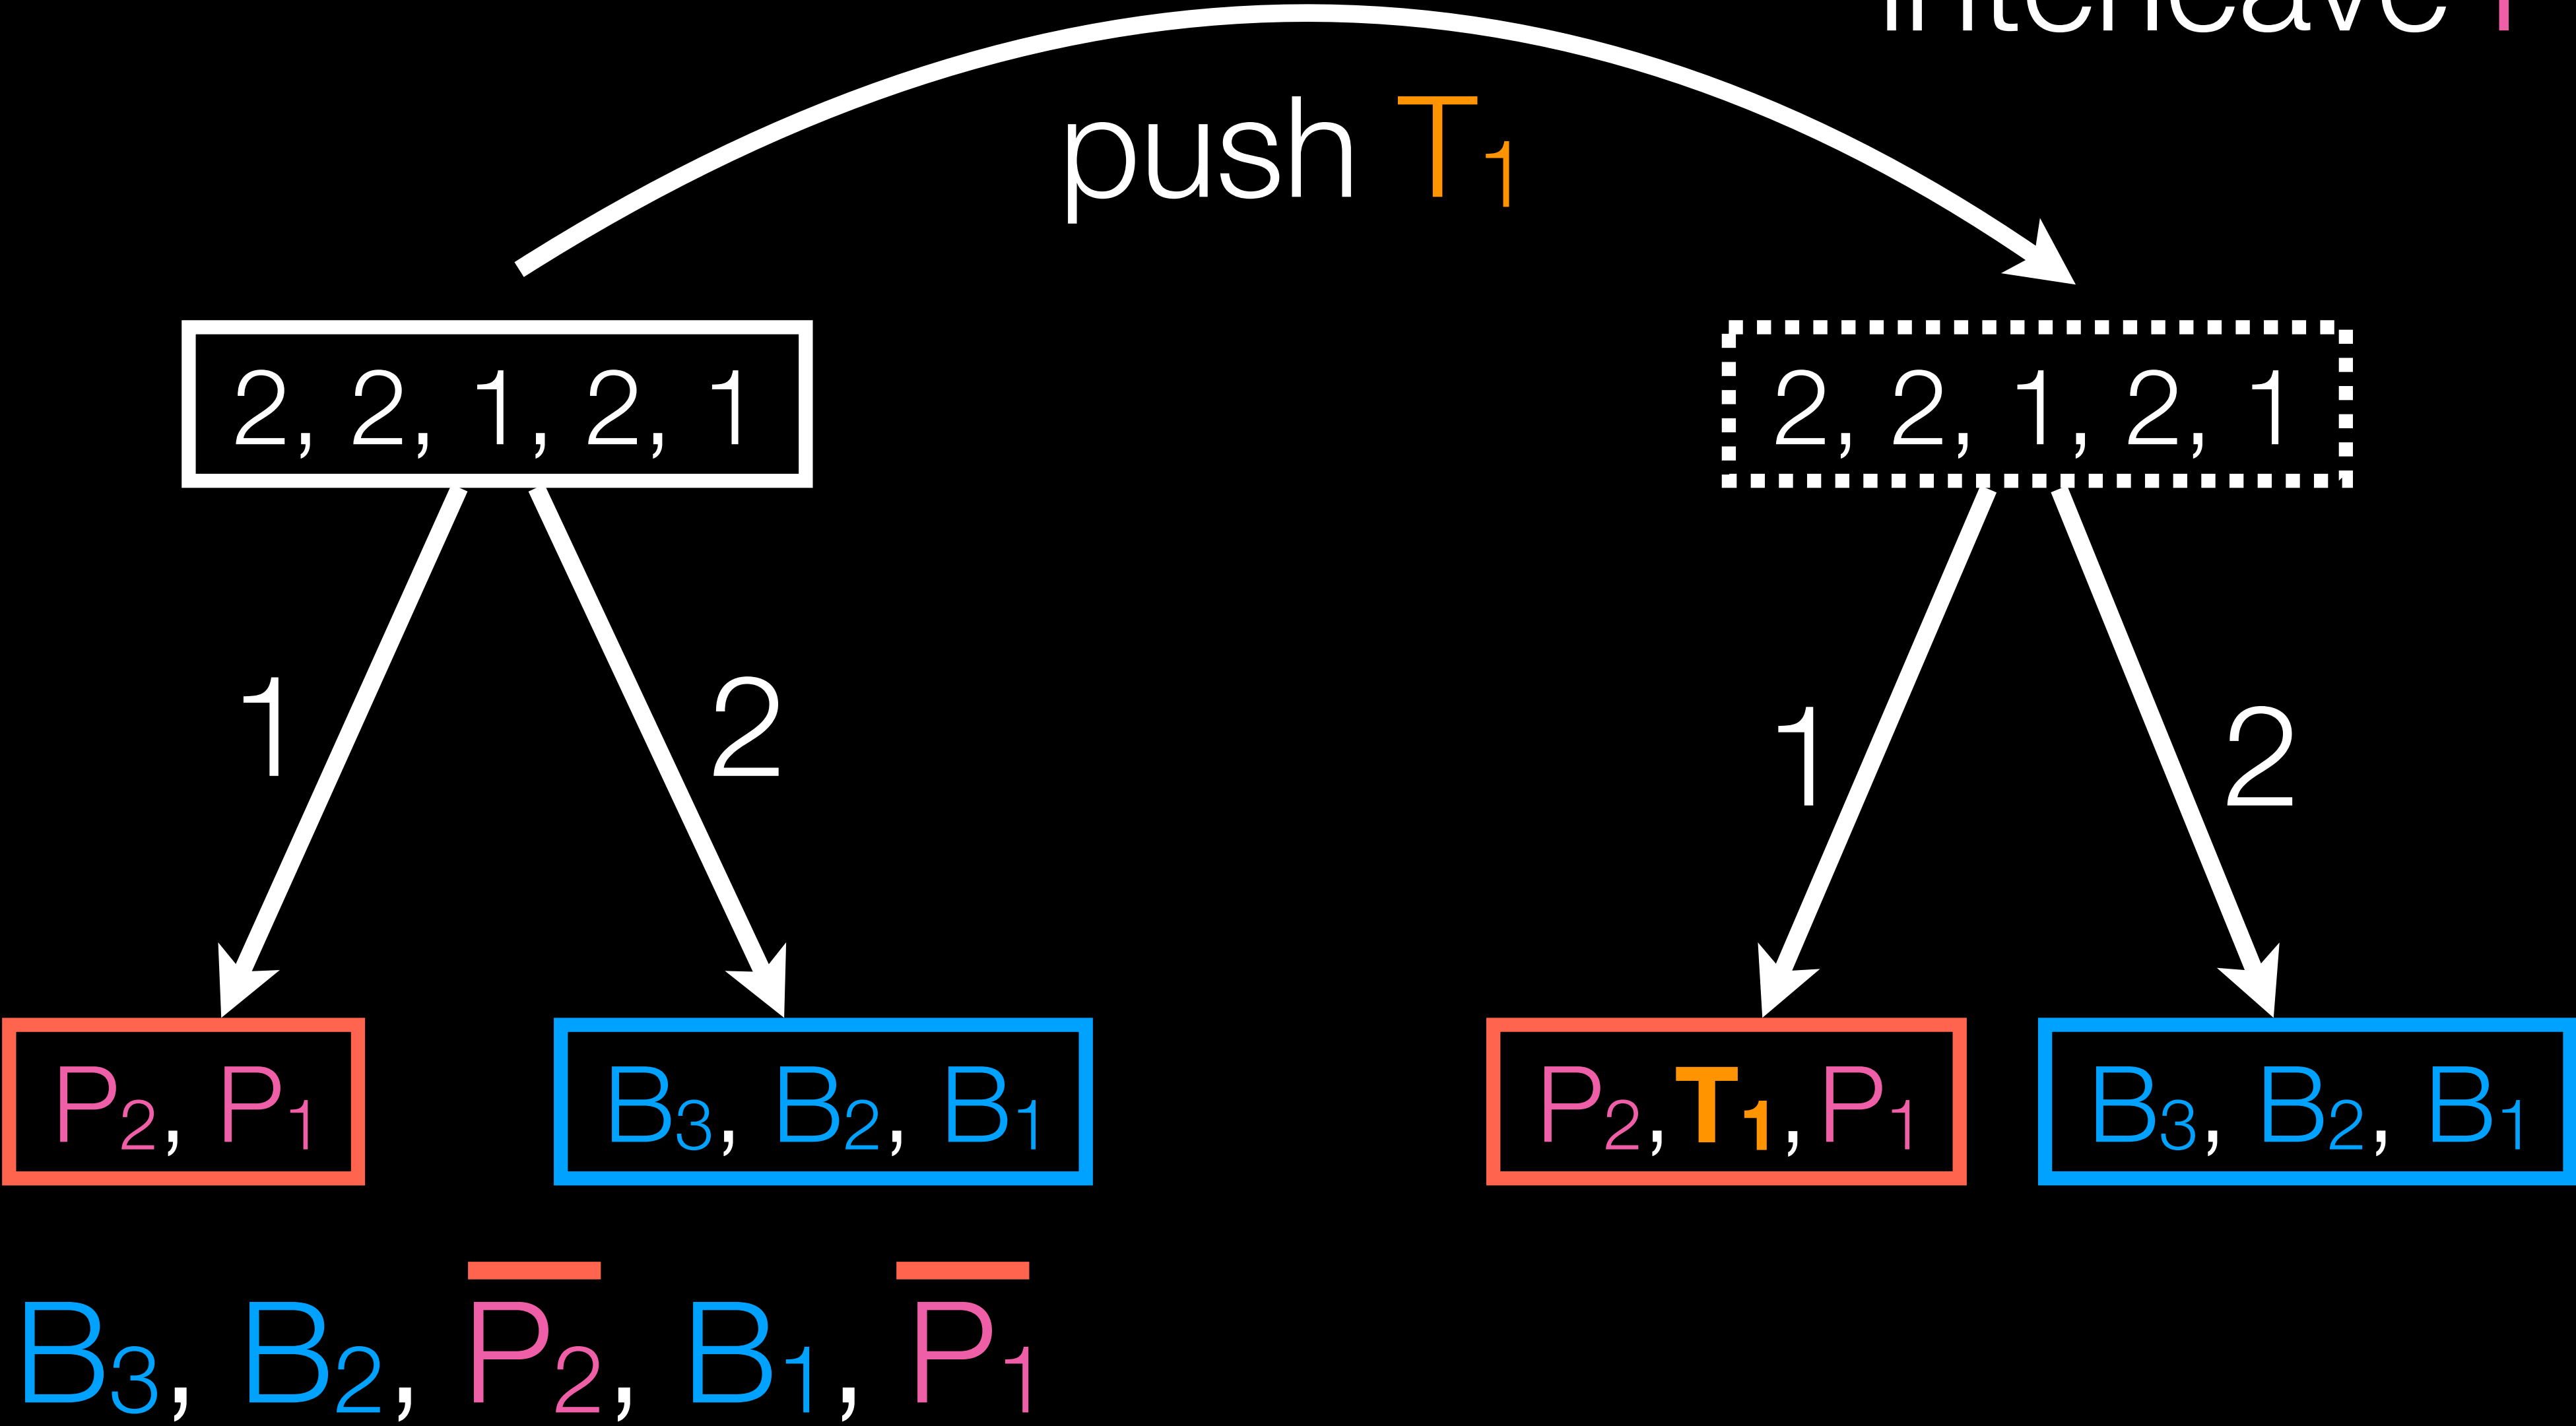
Aside: PIFO trees

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




Aside: PIFO trees

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




Aside: PIFO trees

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



Aside: PIFO trees


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



Aside: PIFO trees


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

push **T₁**



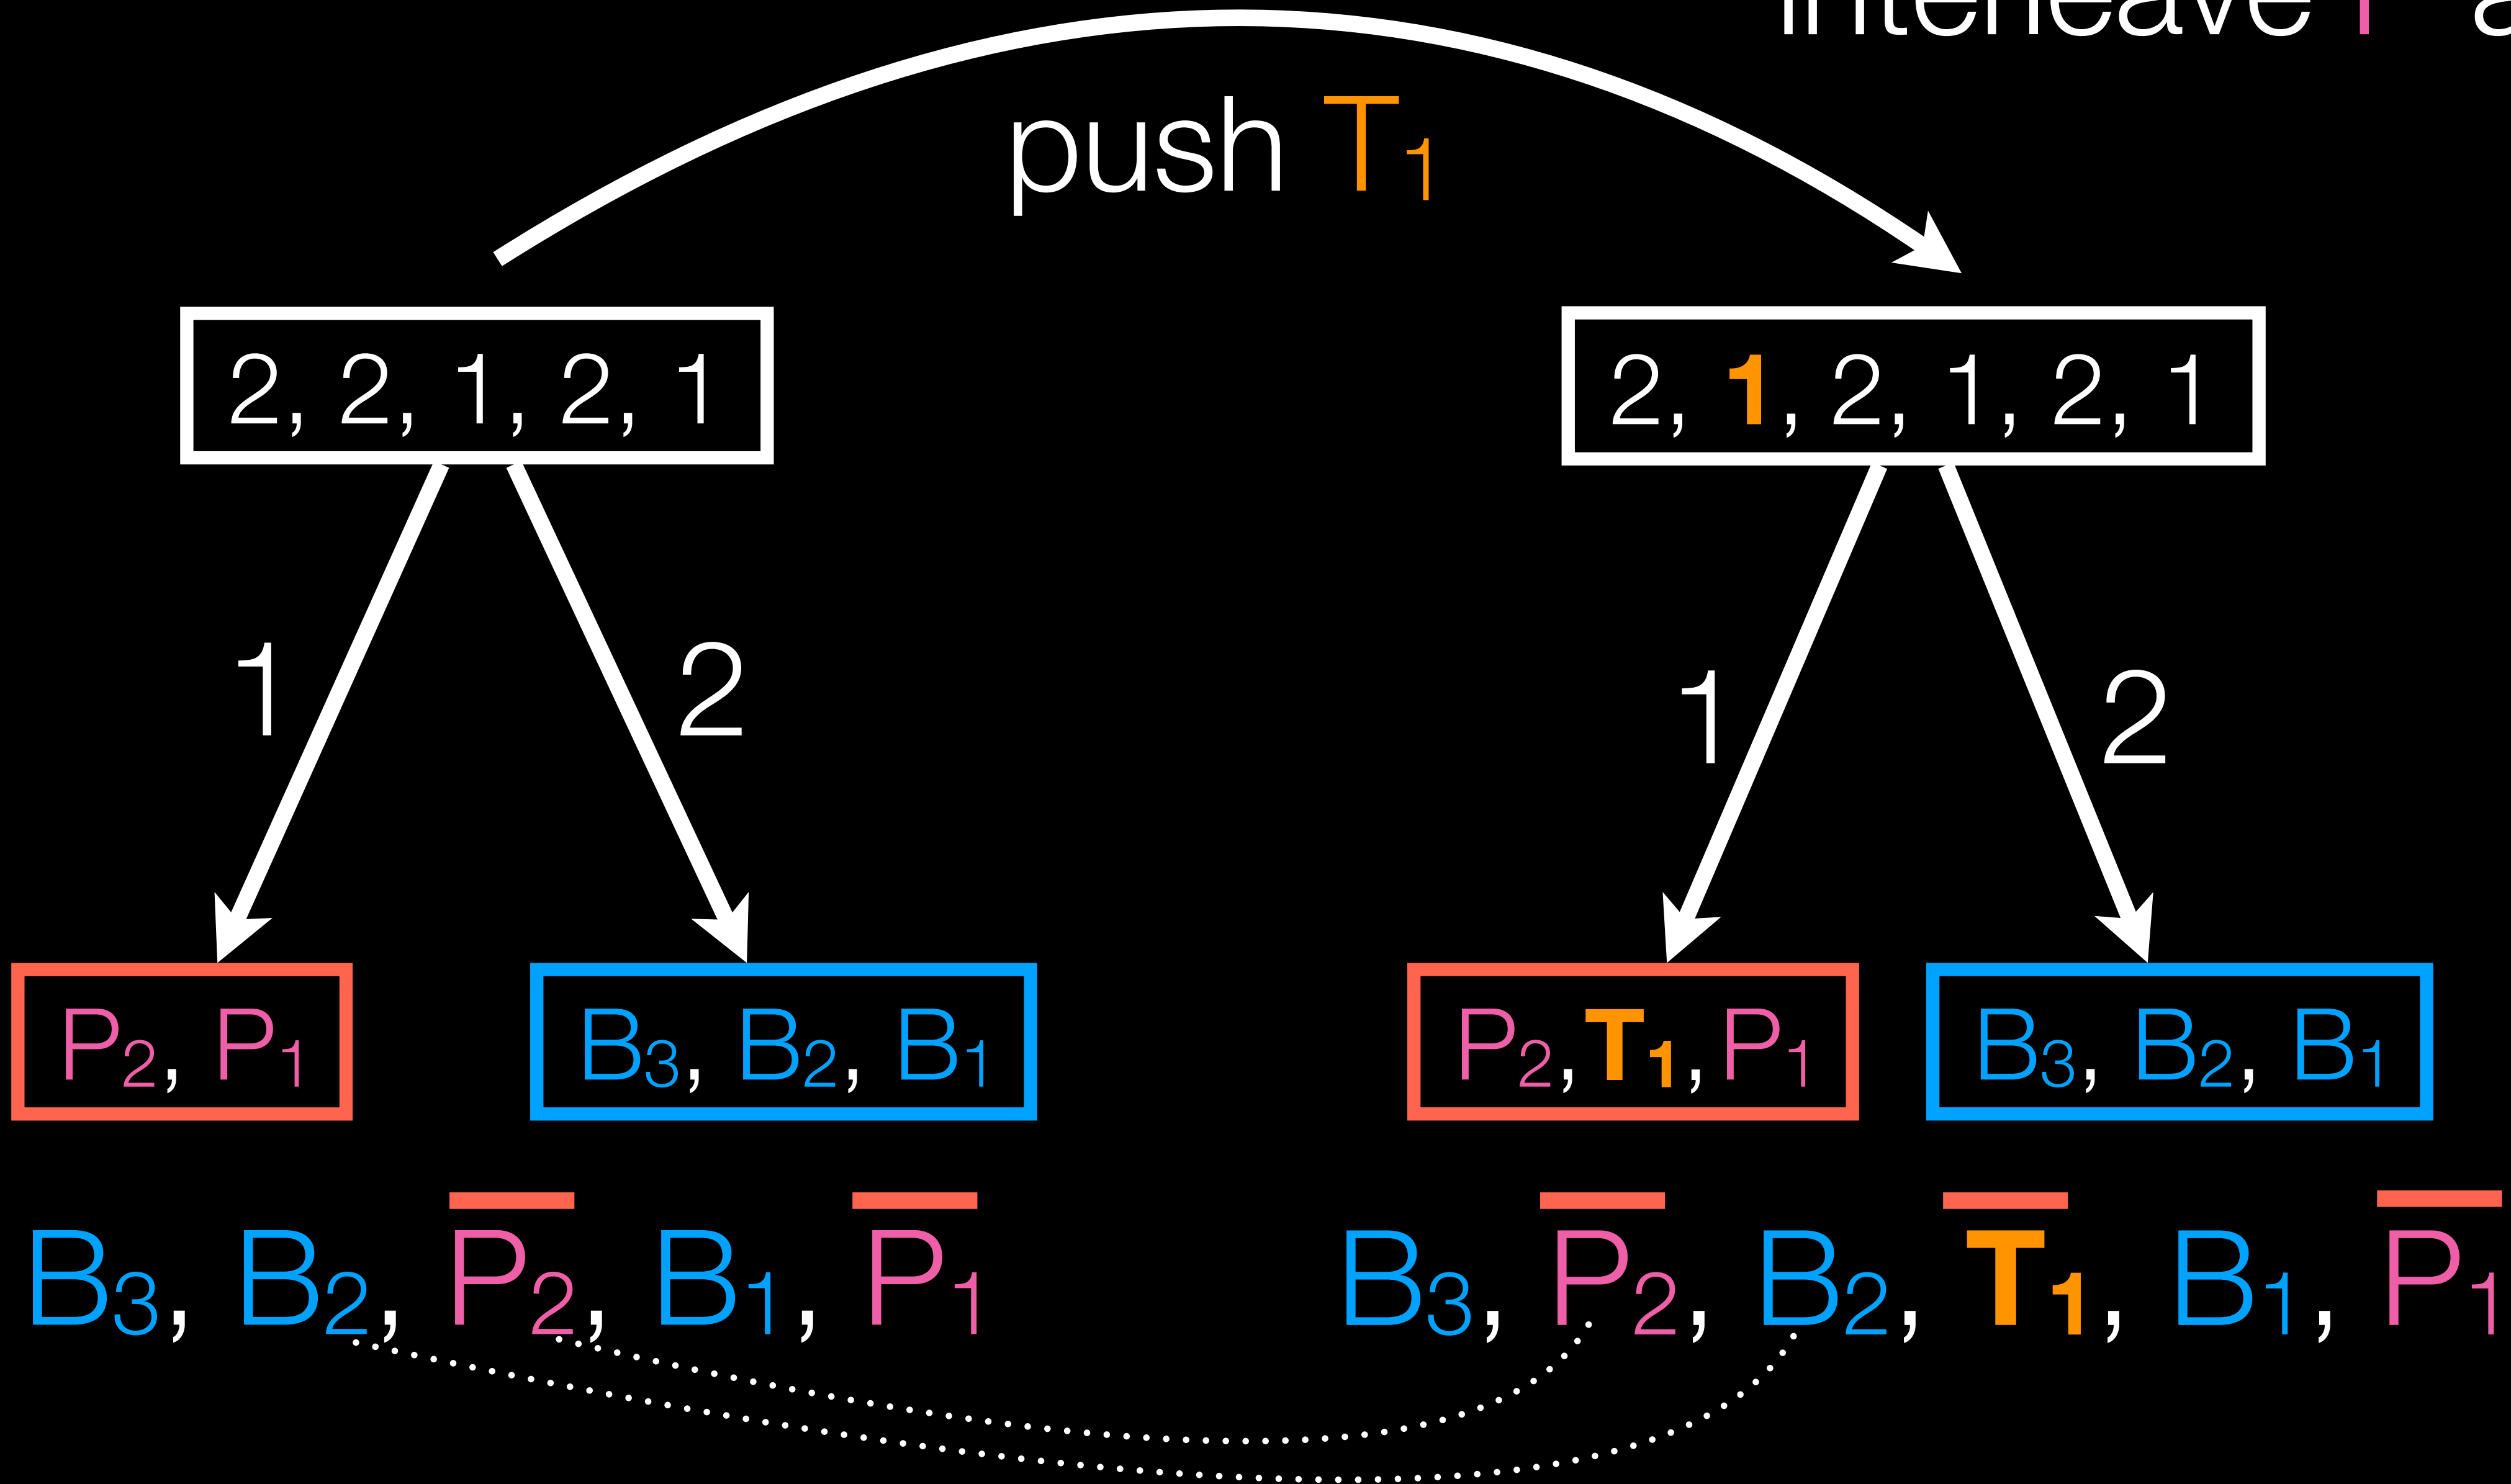
Aside: PIFO trees

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



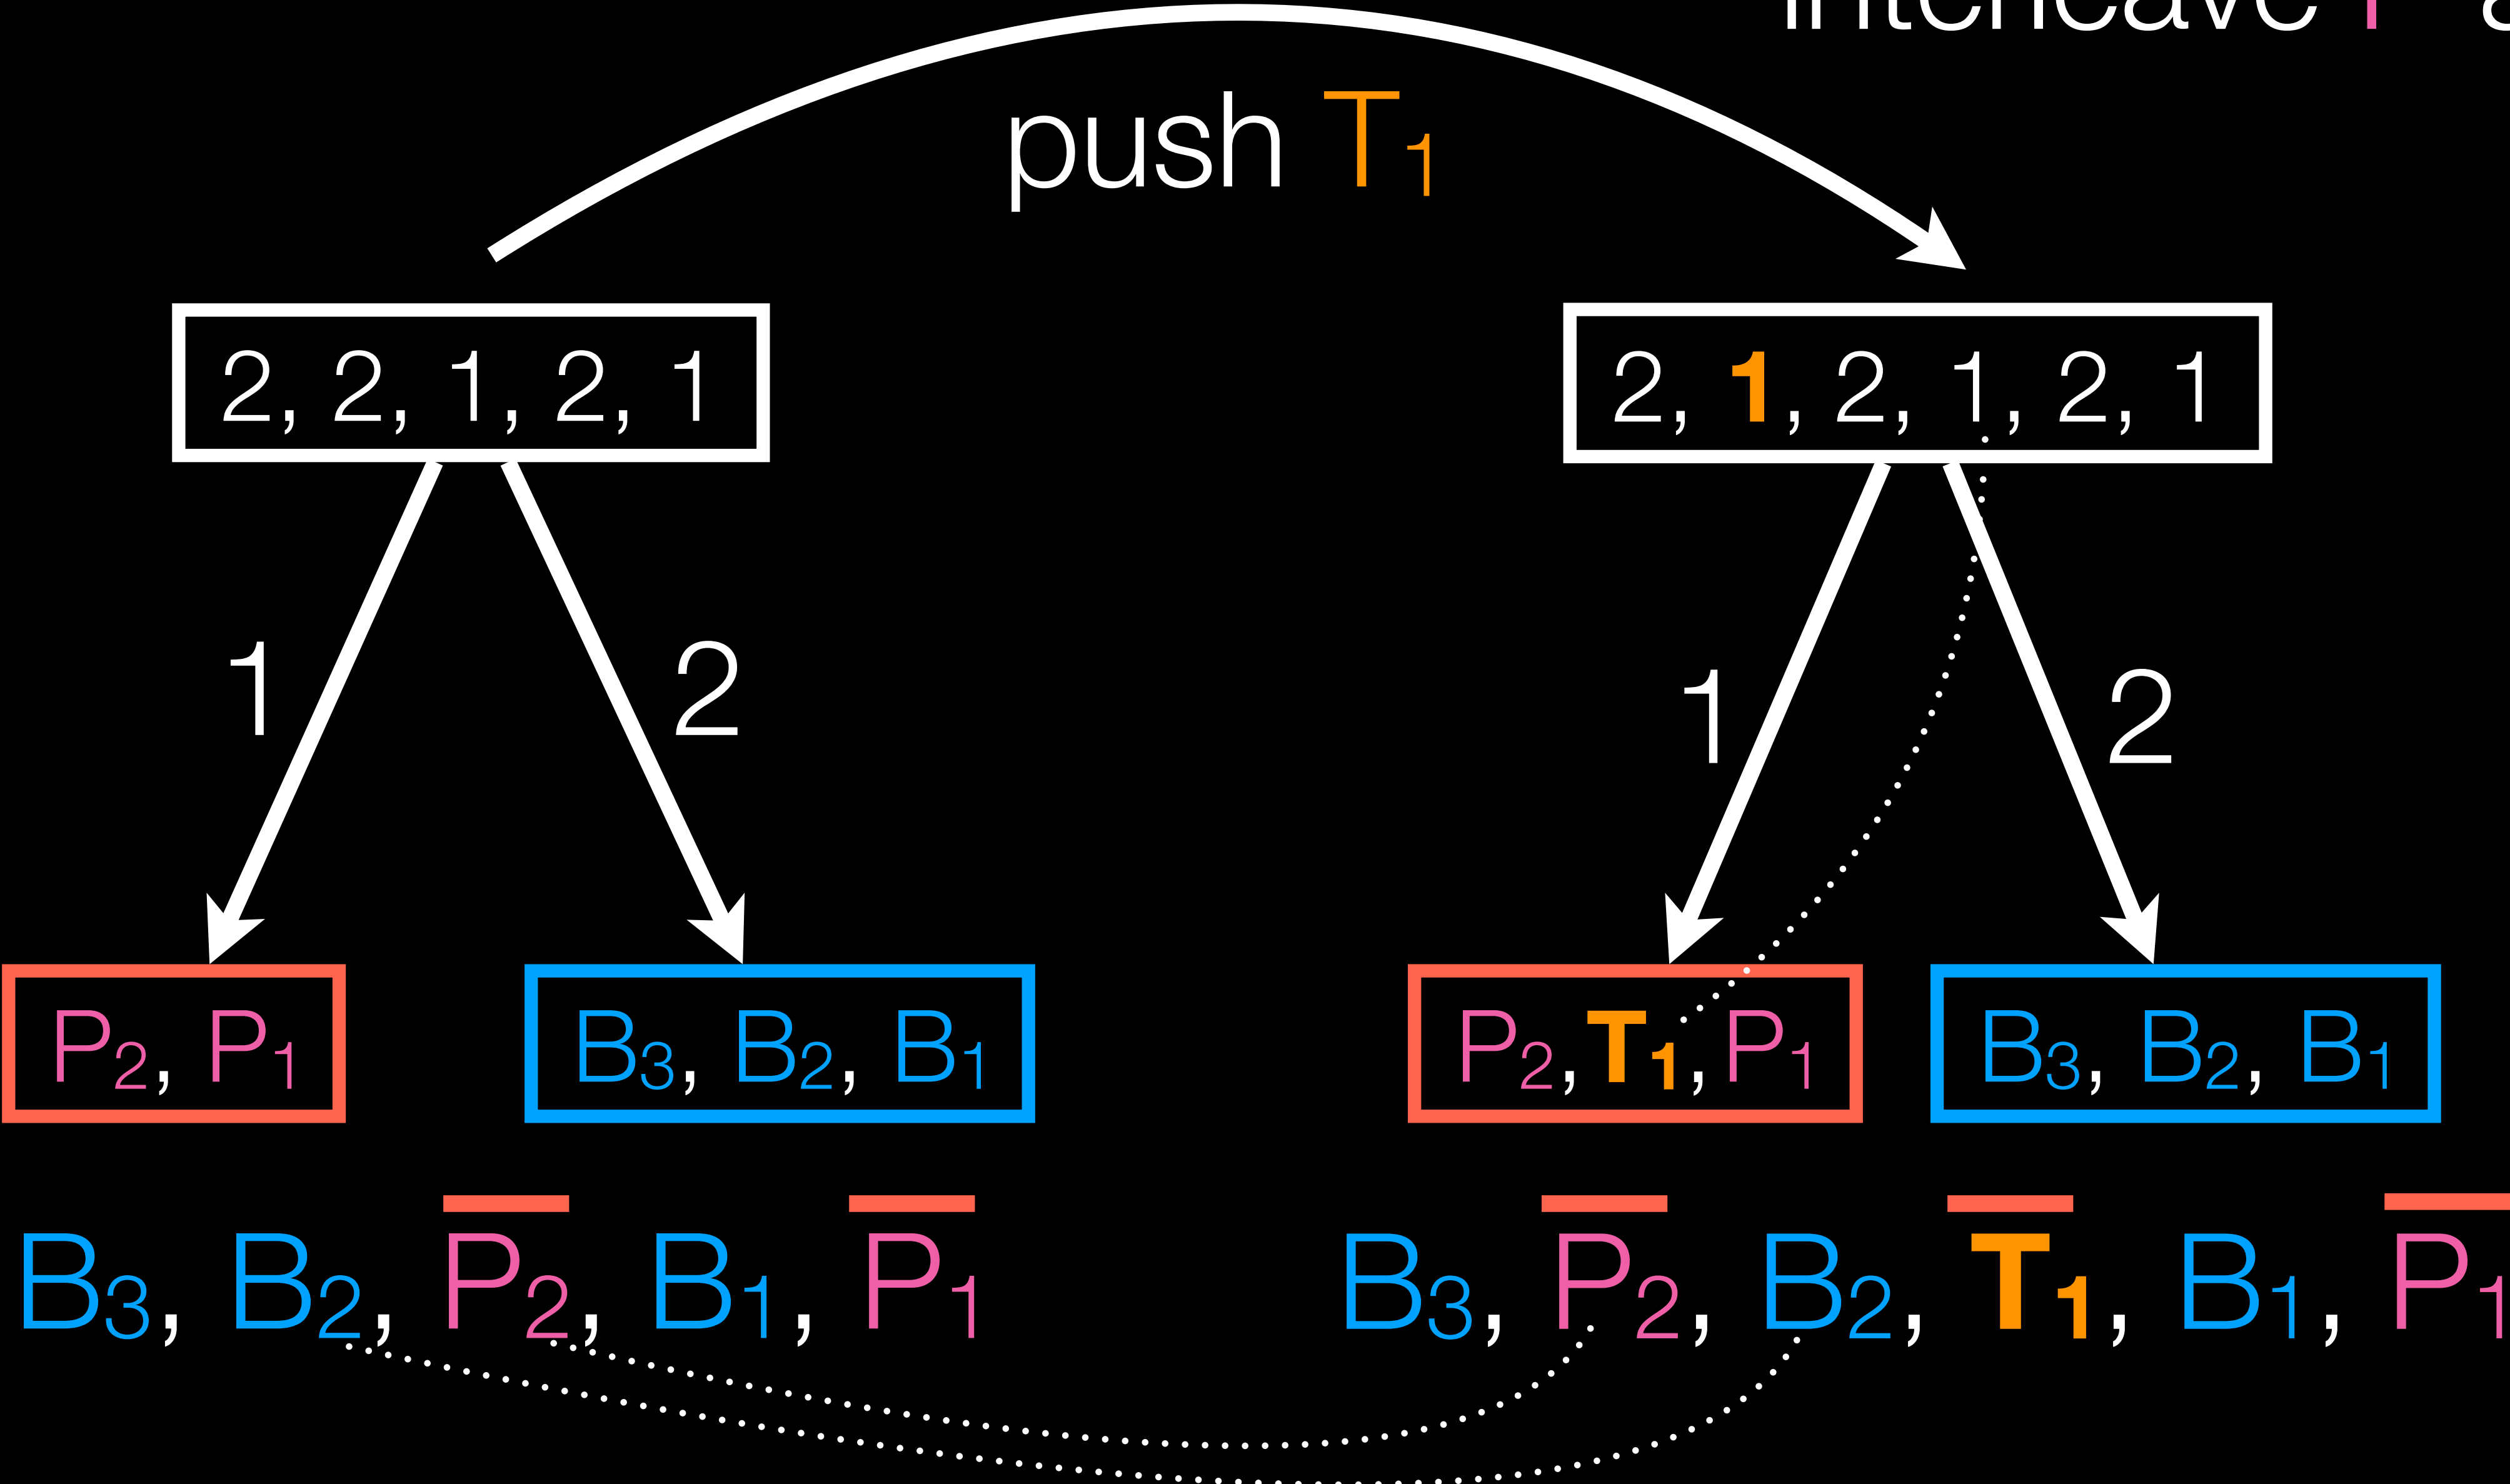
Aside: PIFO trees

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




Aside: PIFO trees

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

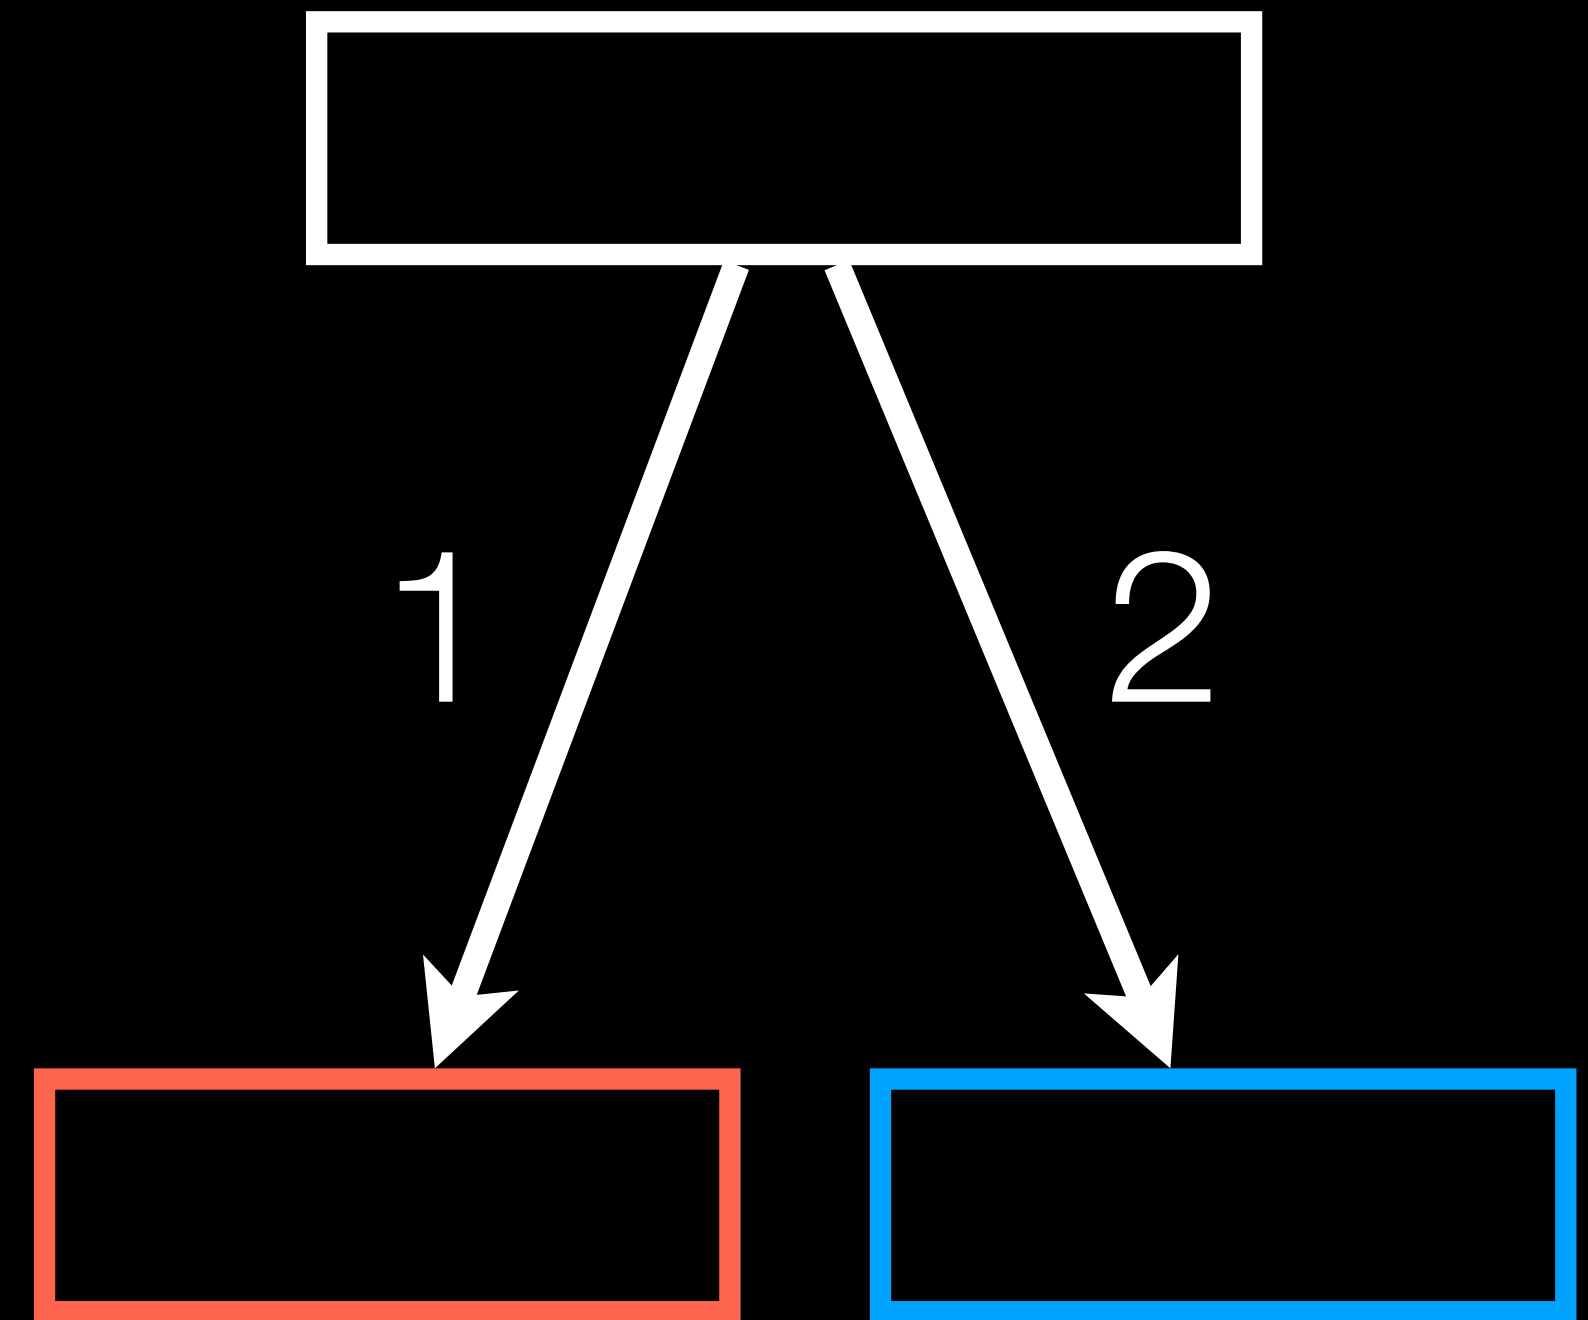


Aside: PIFO trees

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

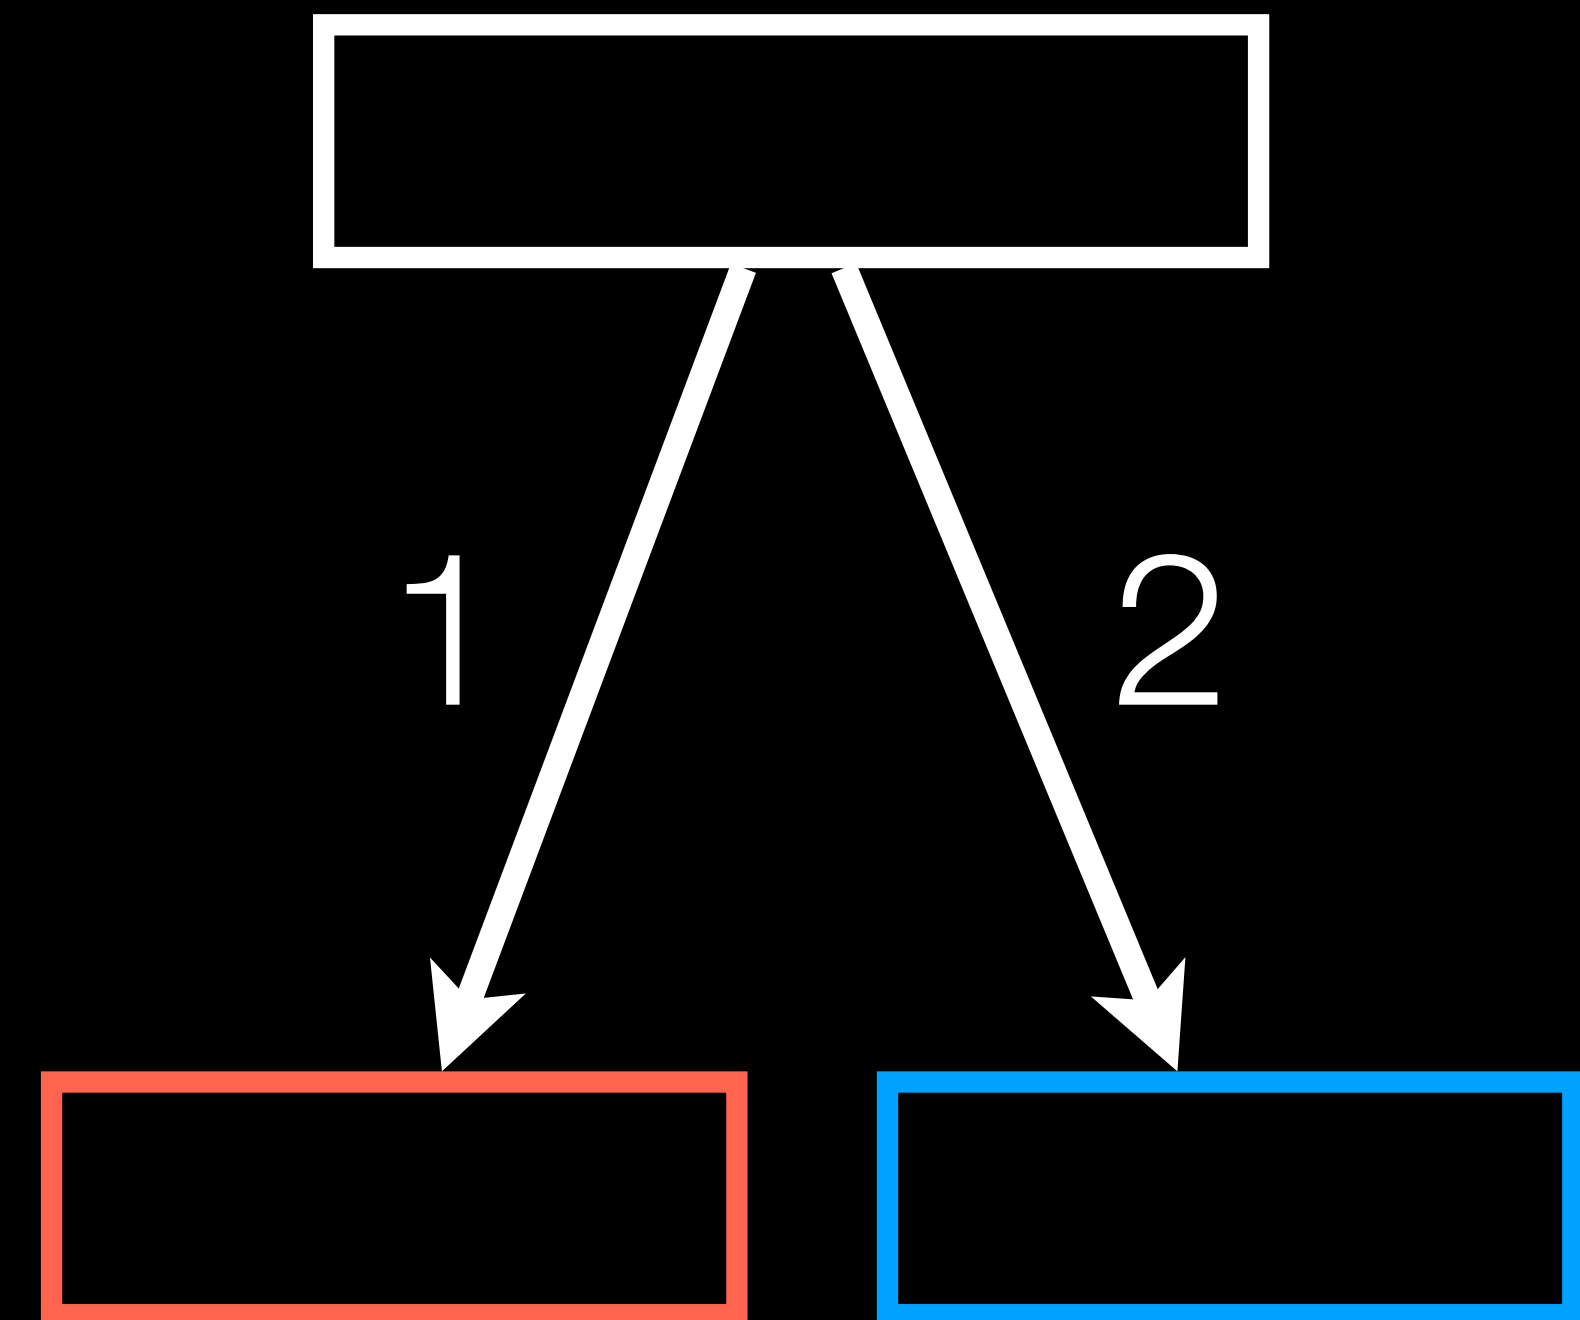


Key Insight



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



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.~~



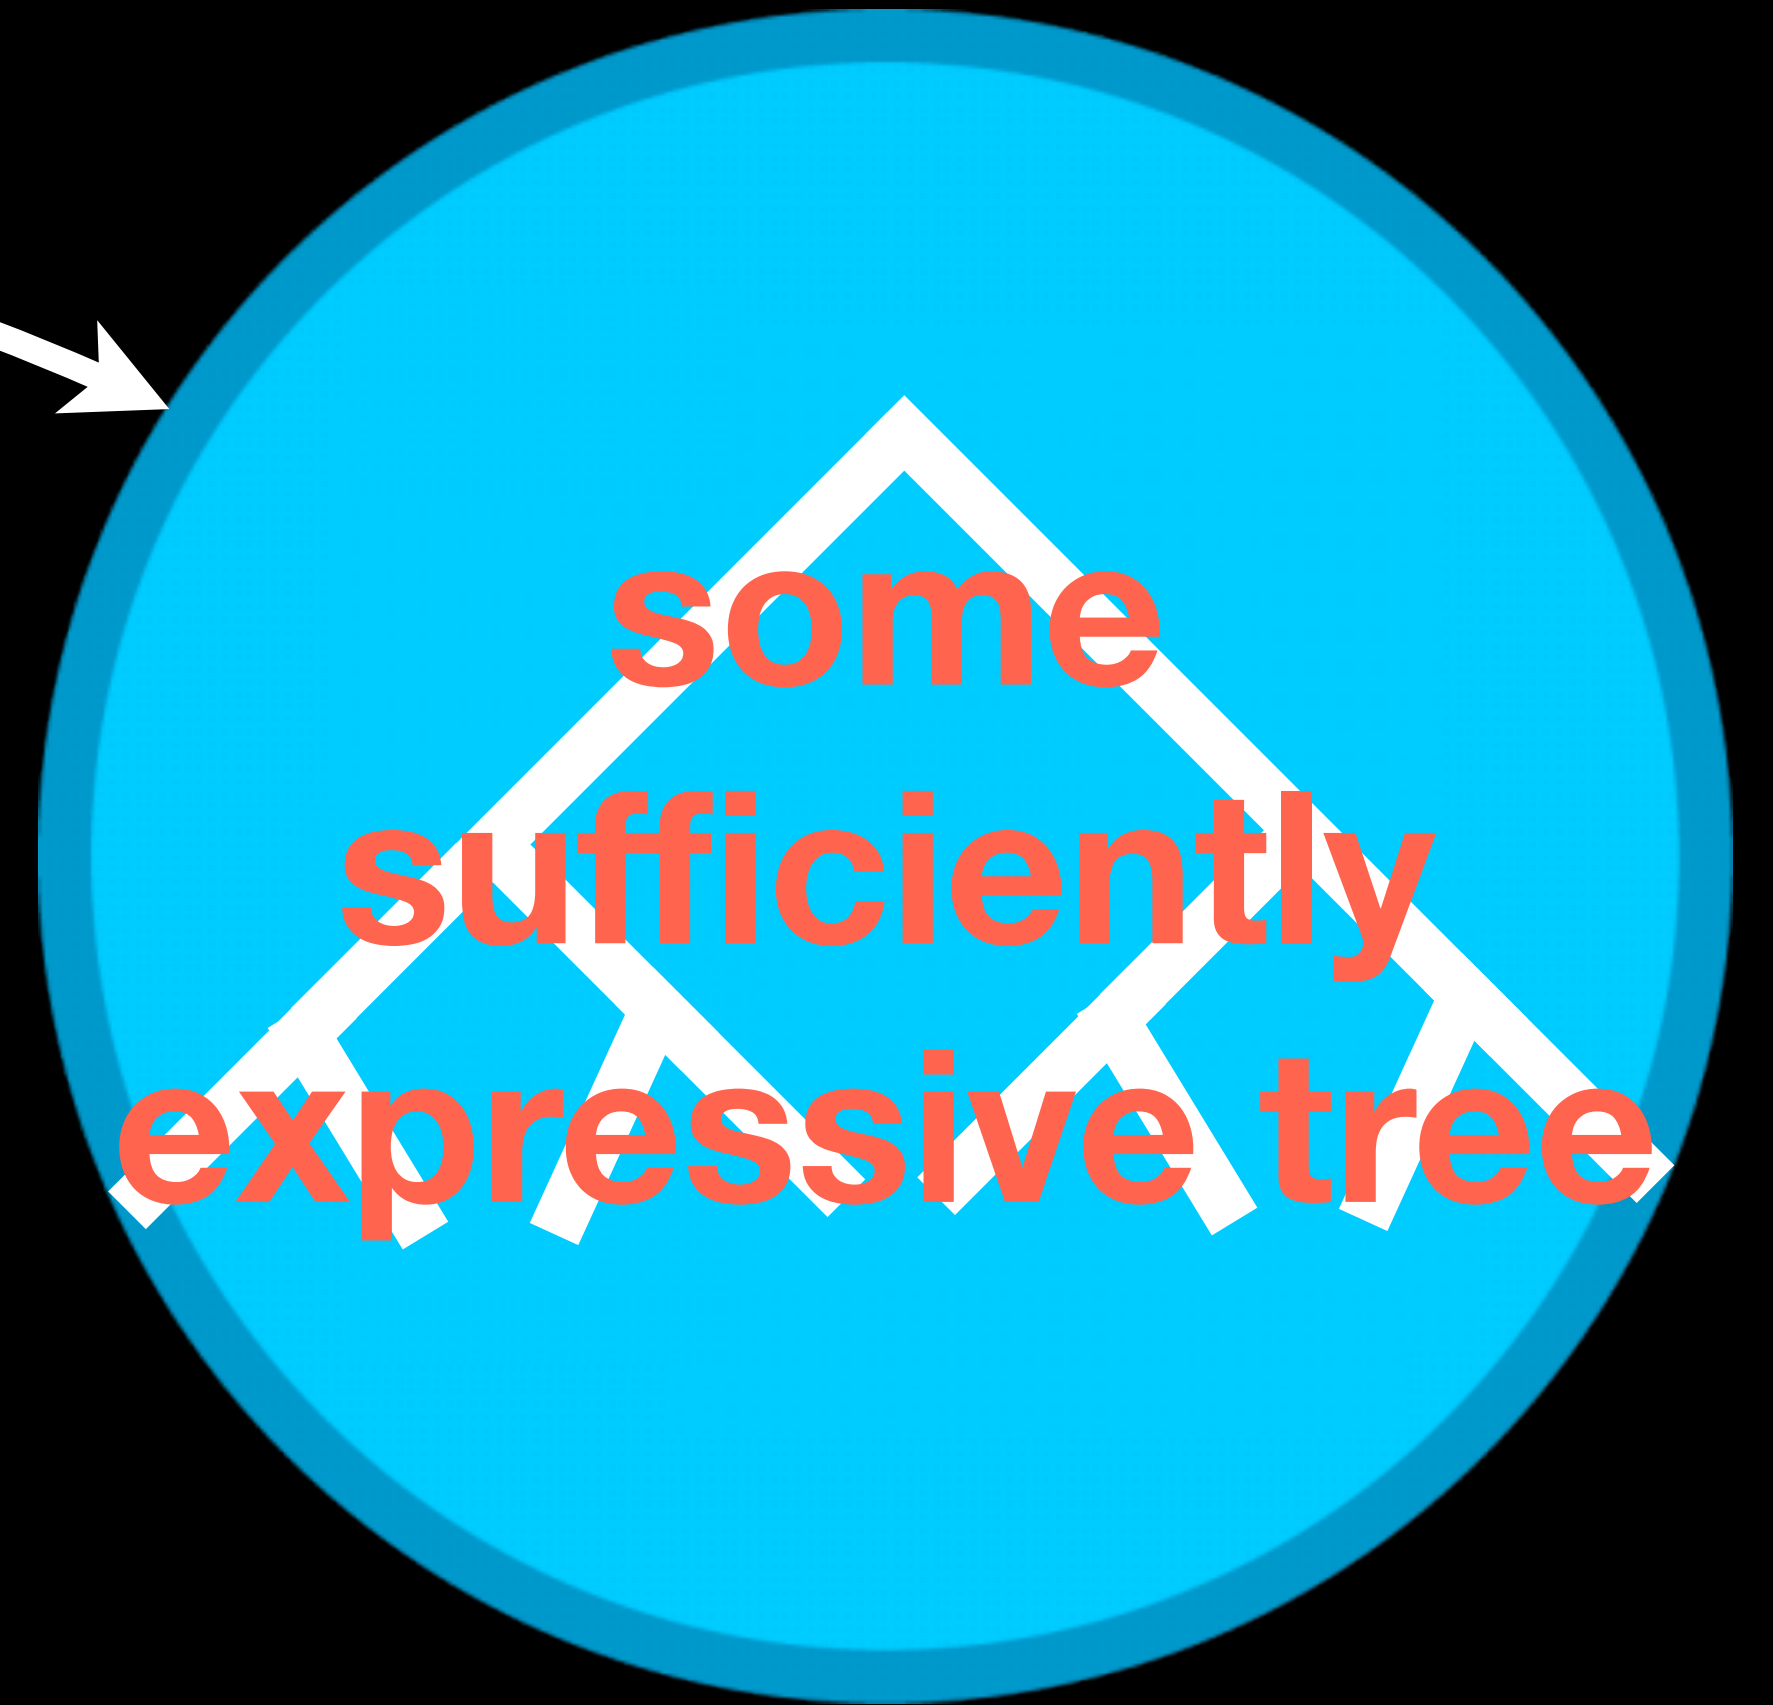
A human needs a *range* of trees.

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




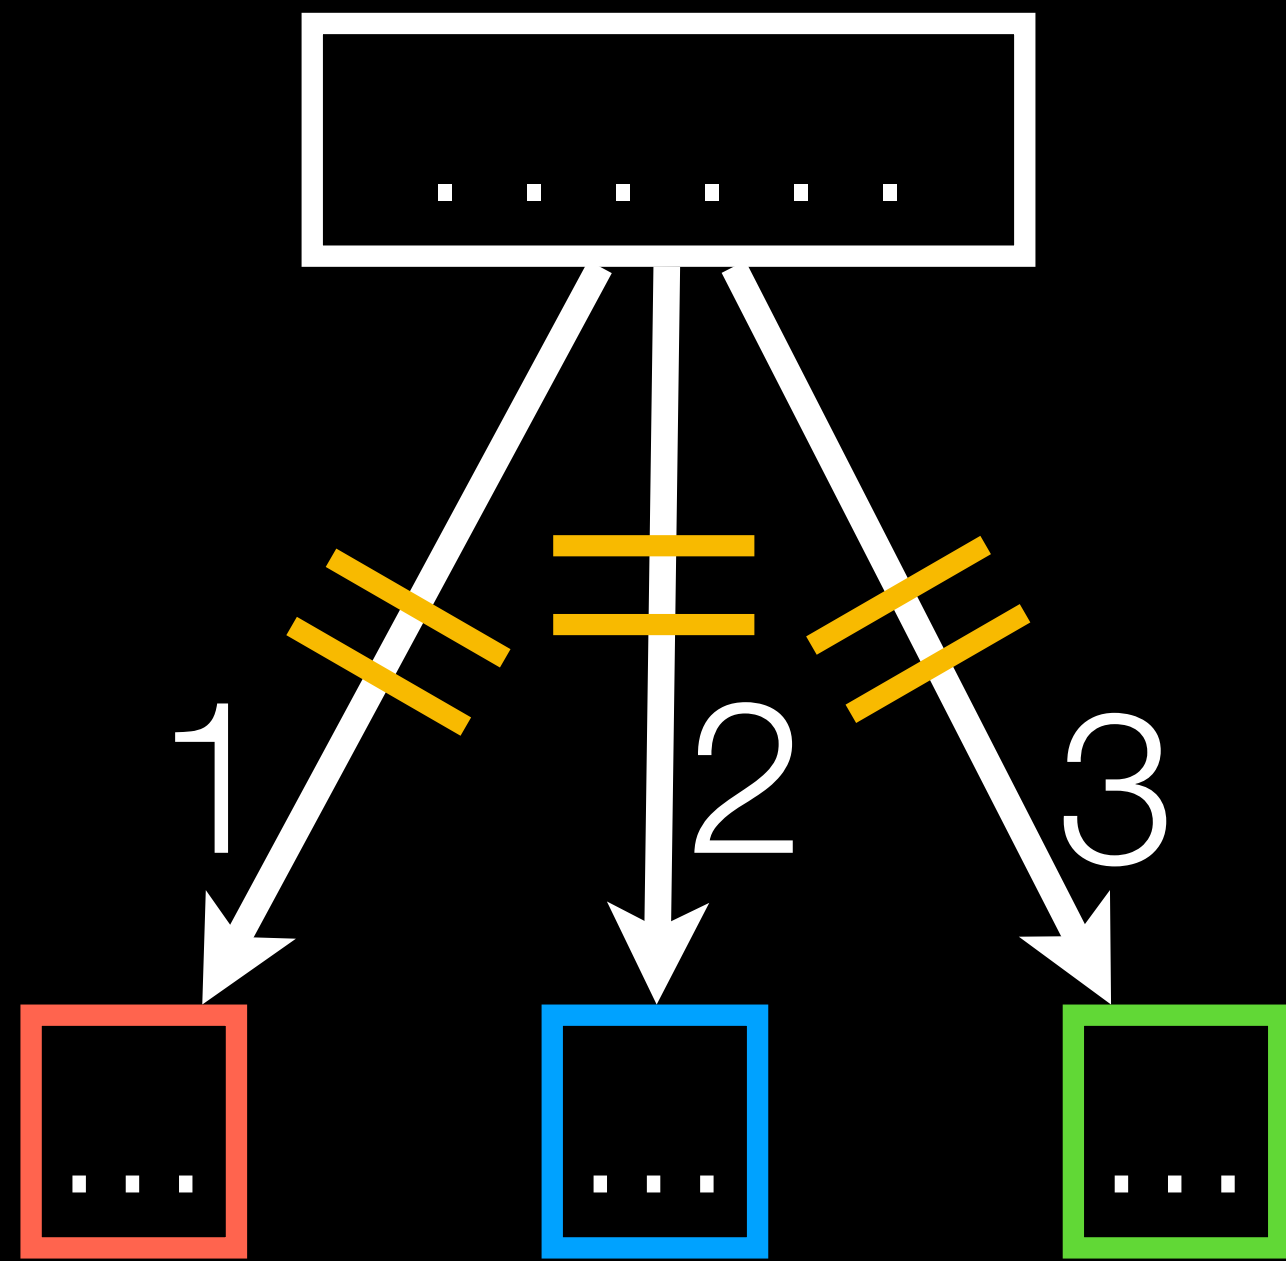
Compiling programs



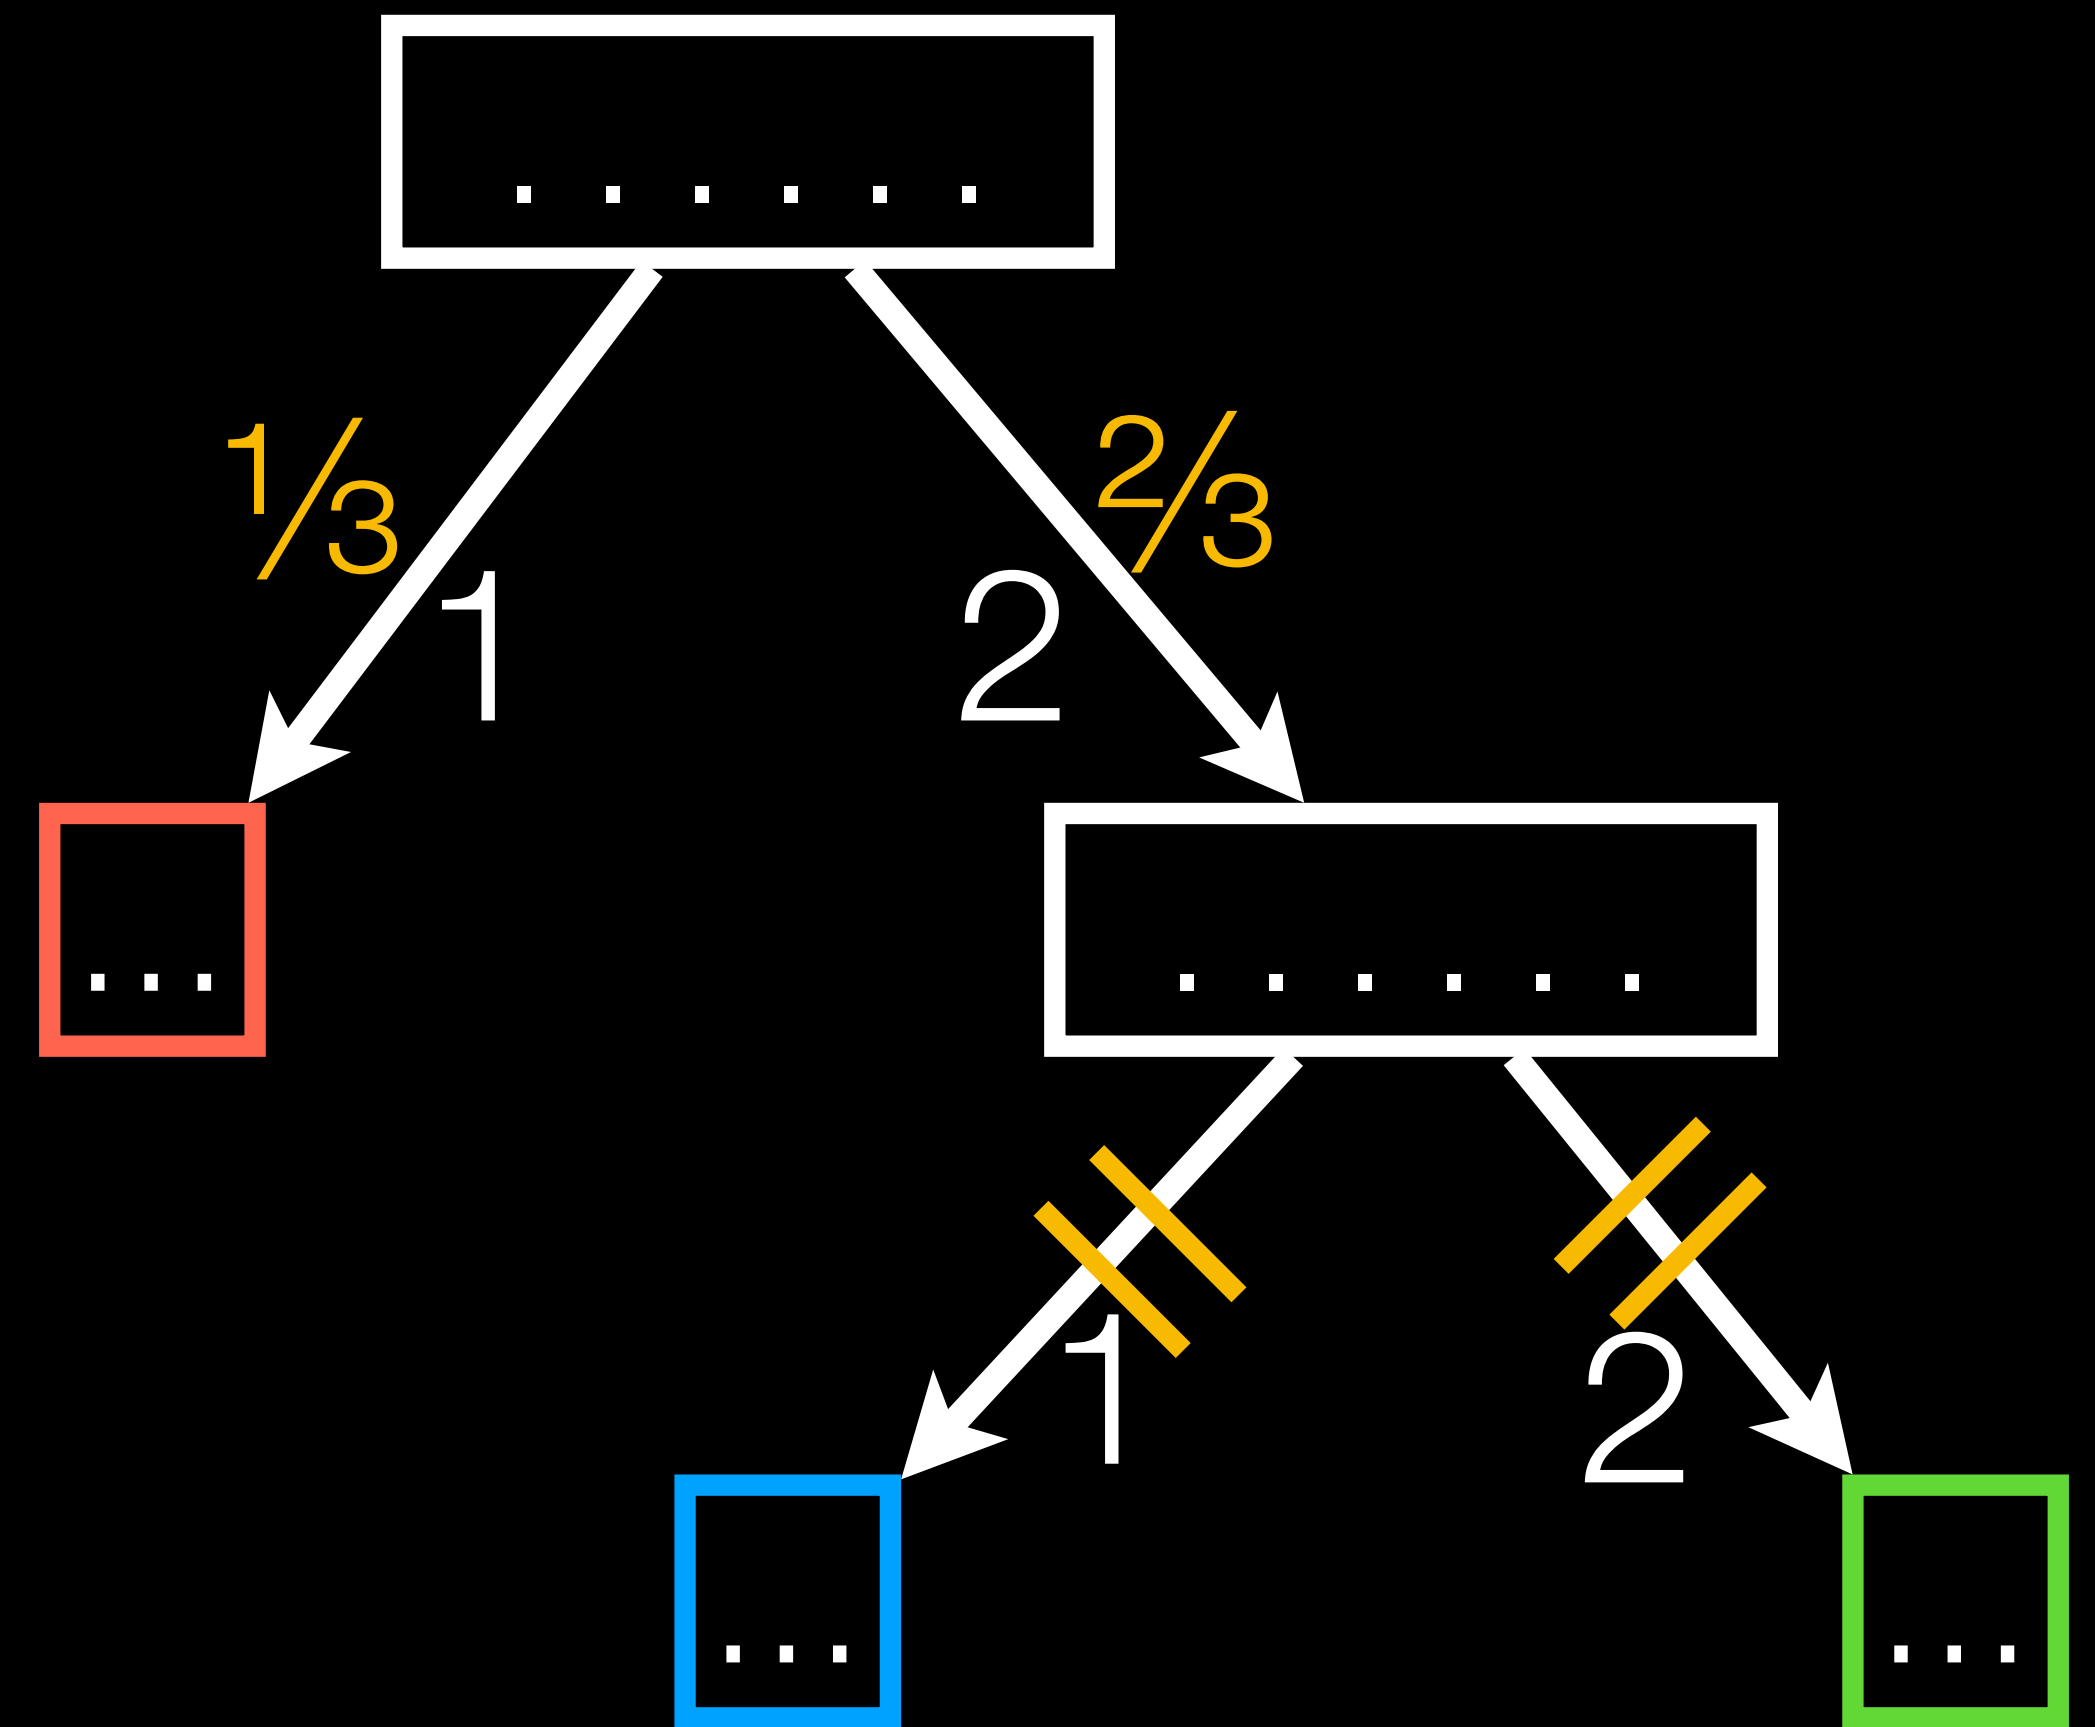

Compiling programs




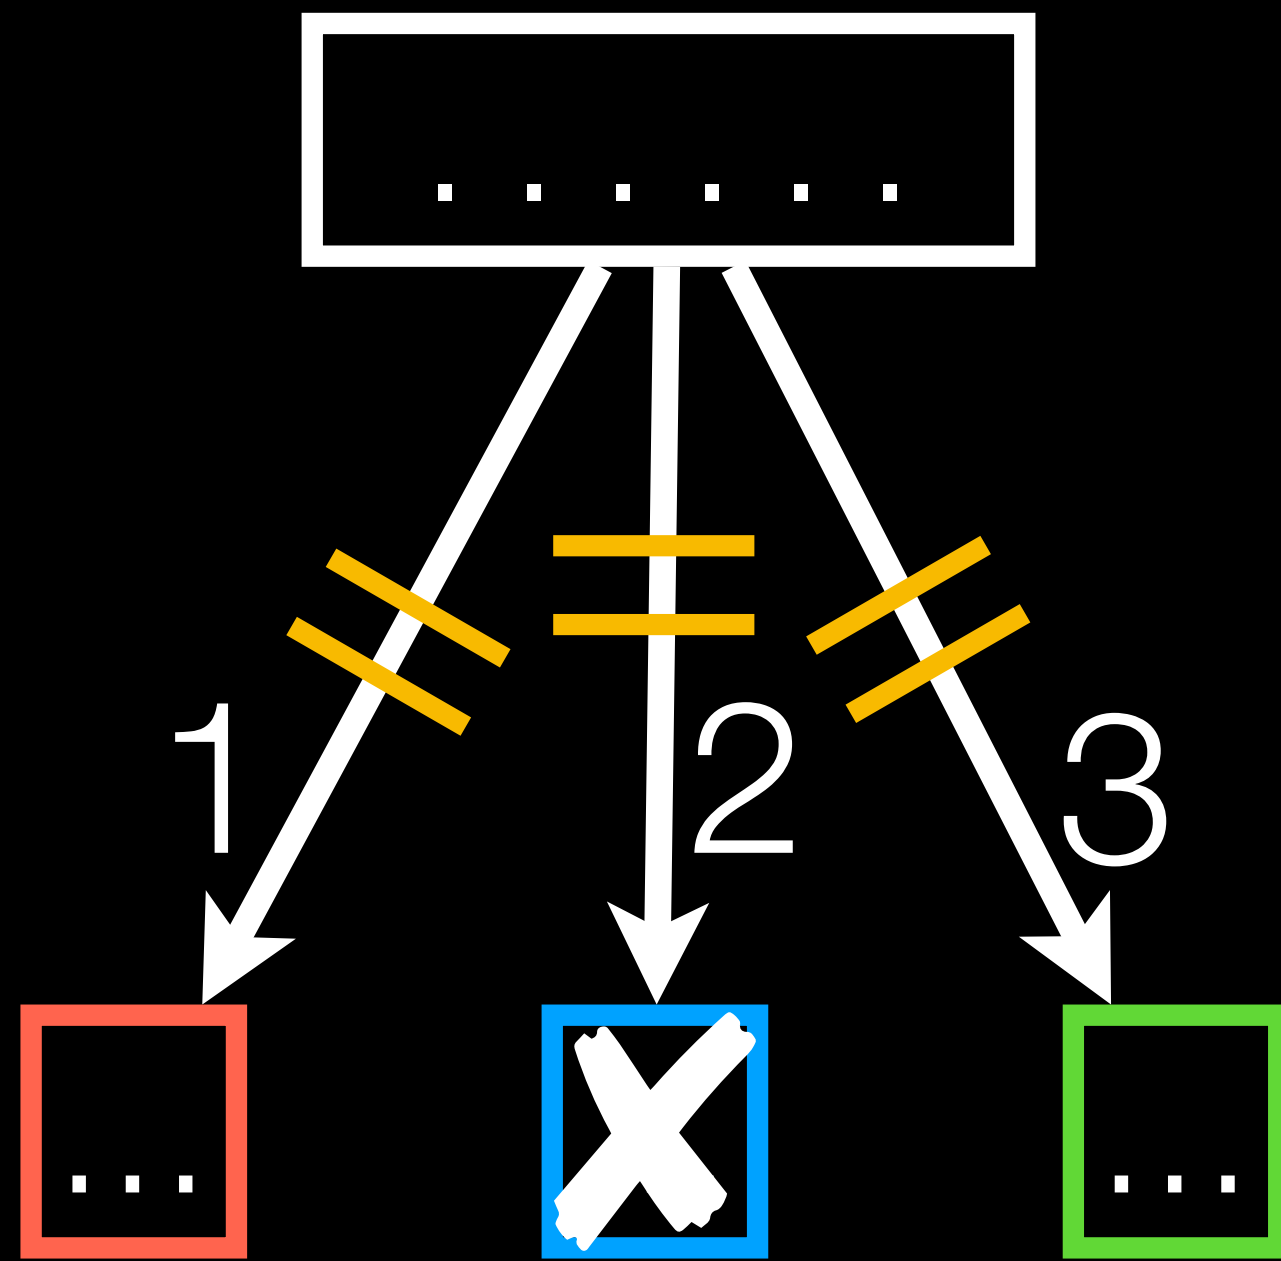
Compiling programs



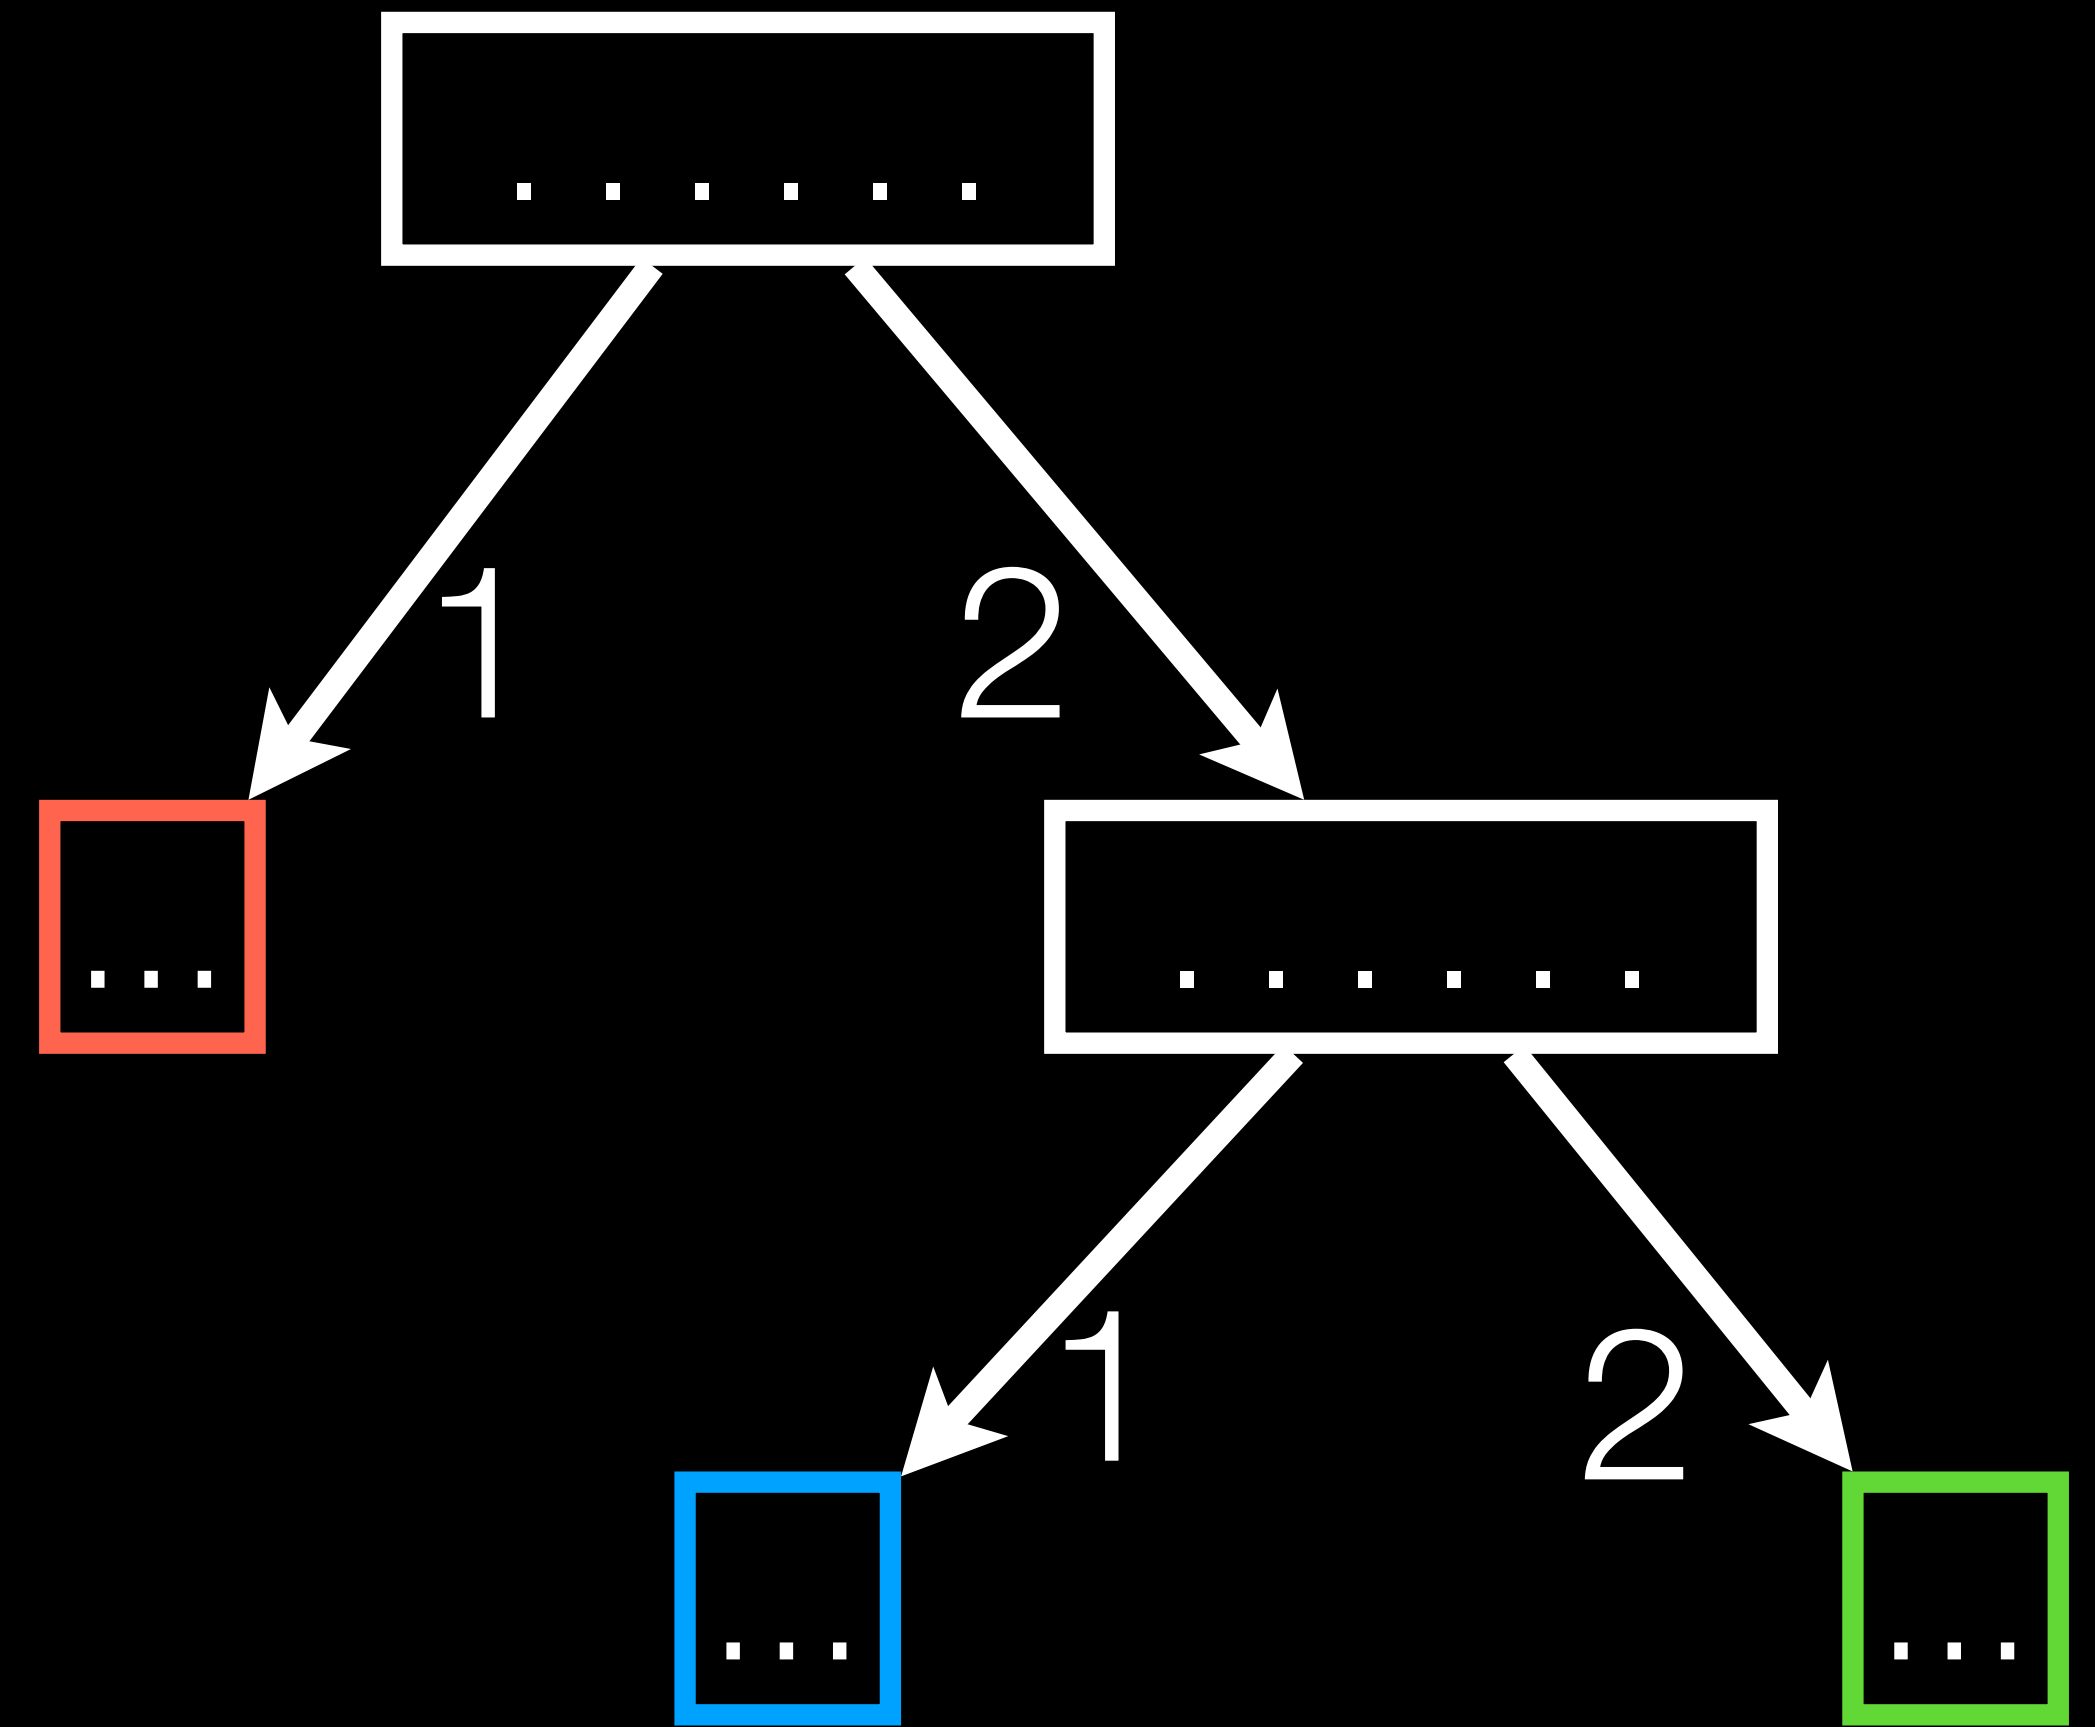
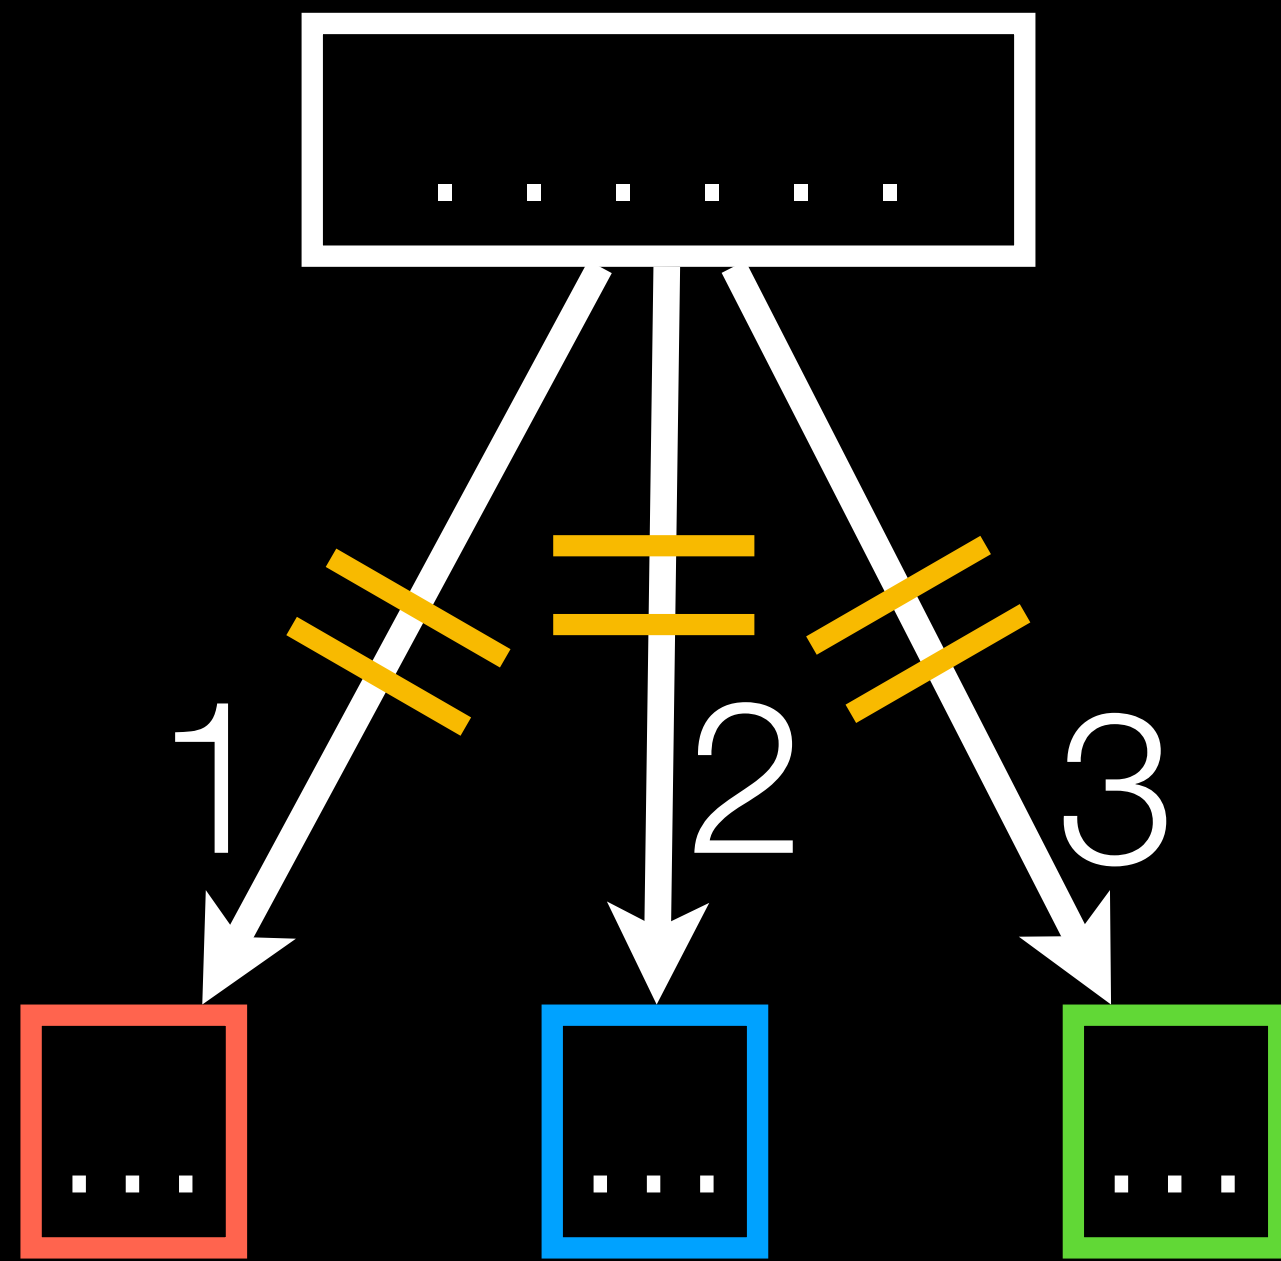
Compiling programs



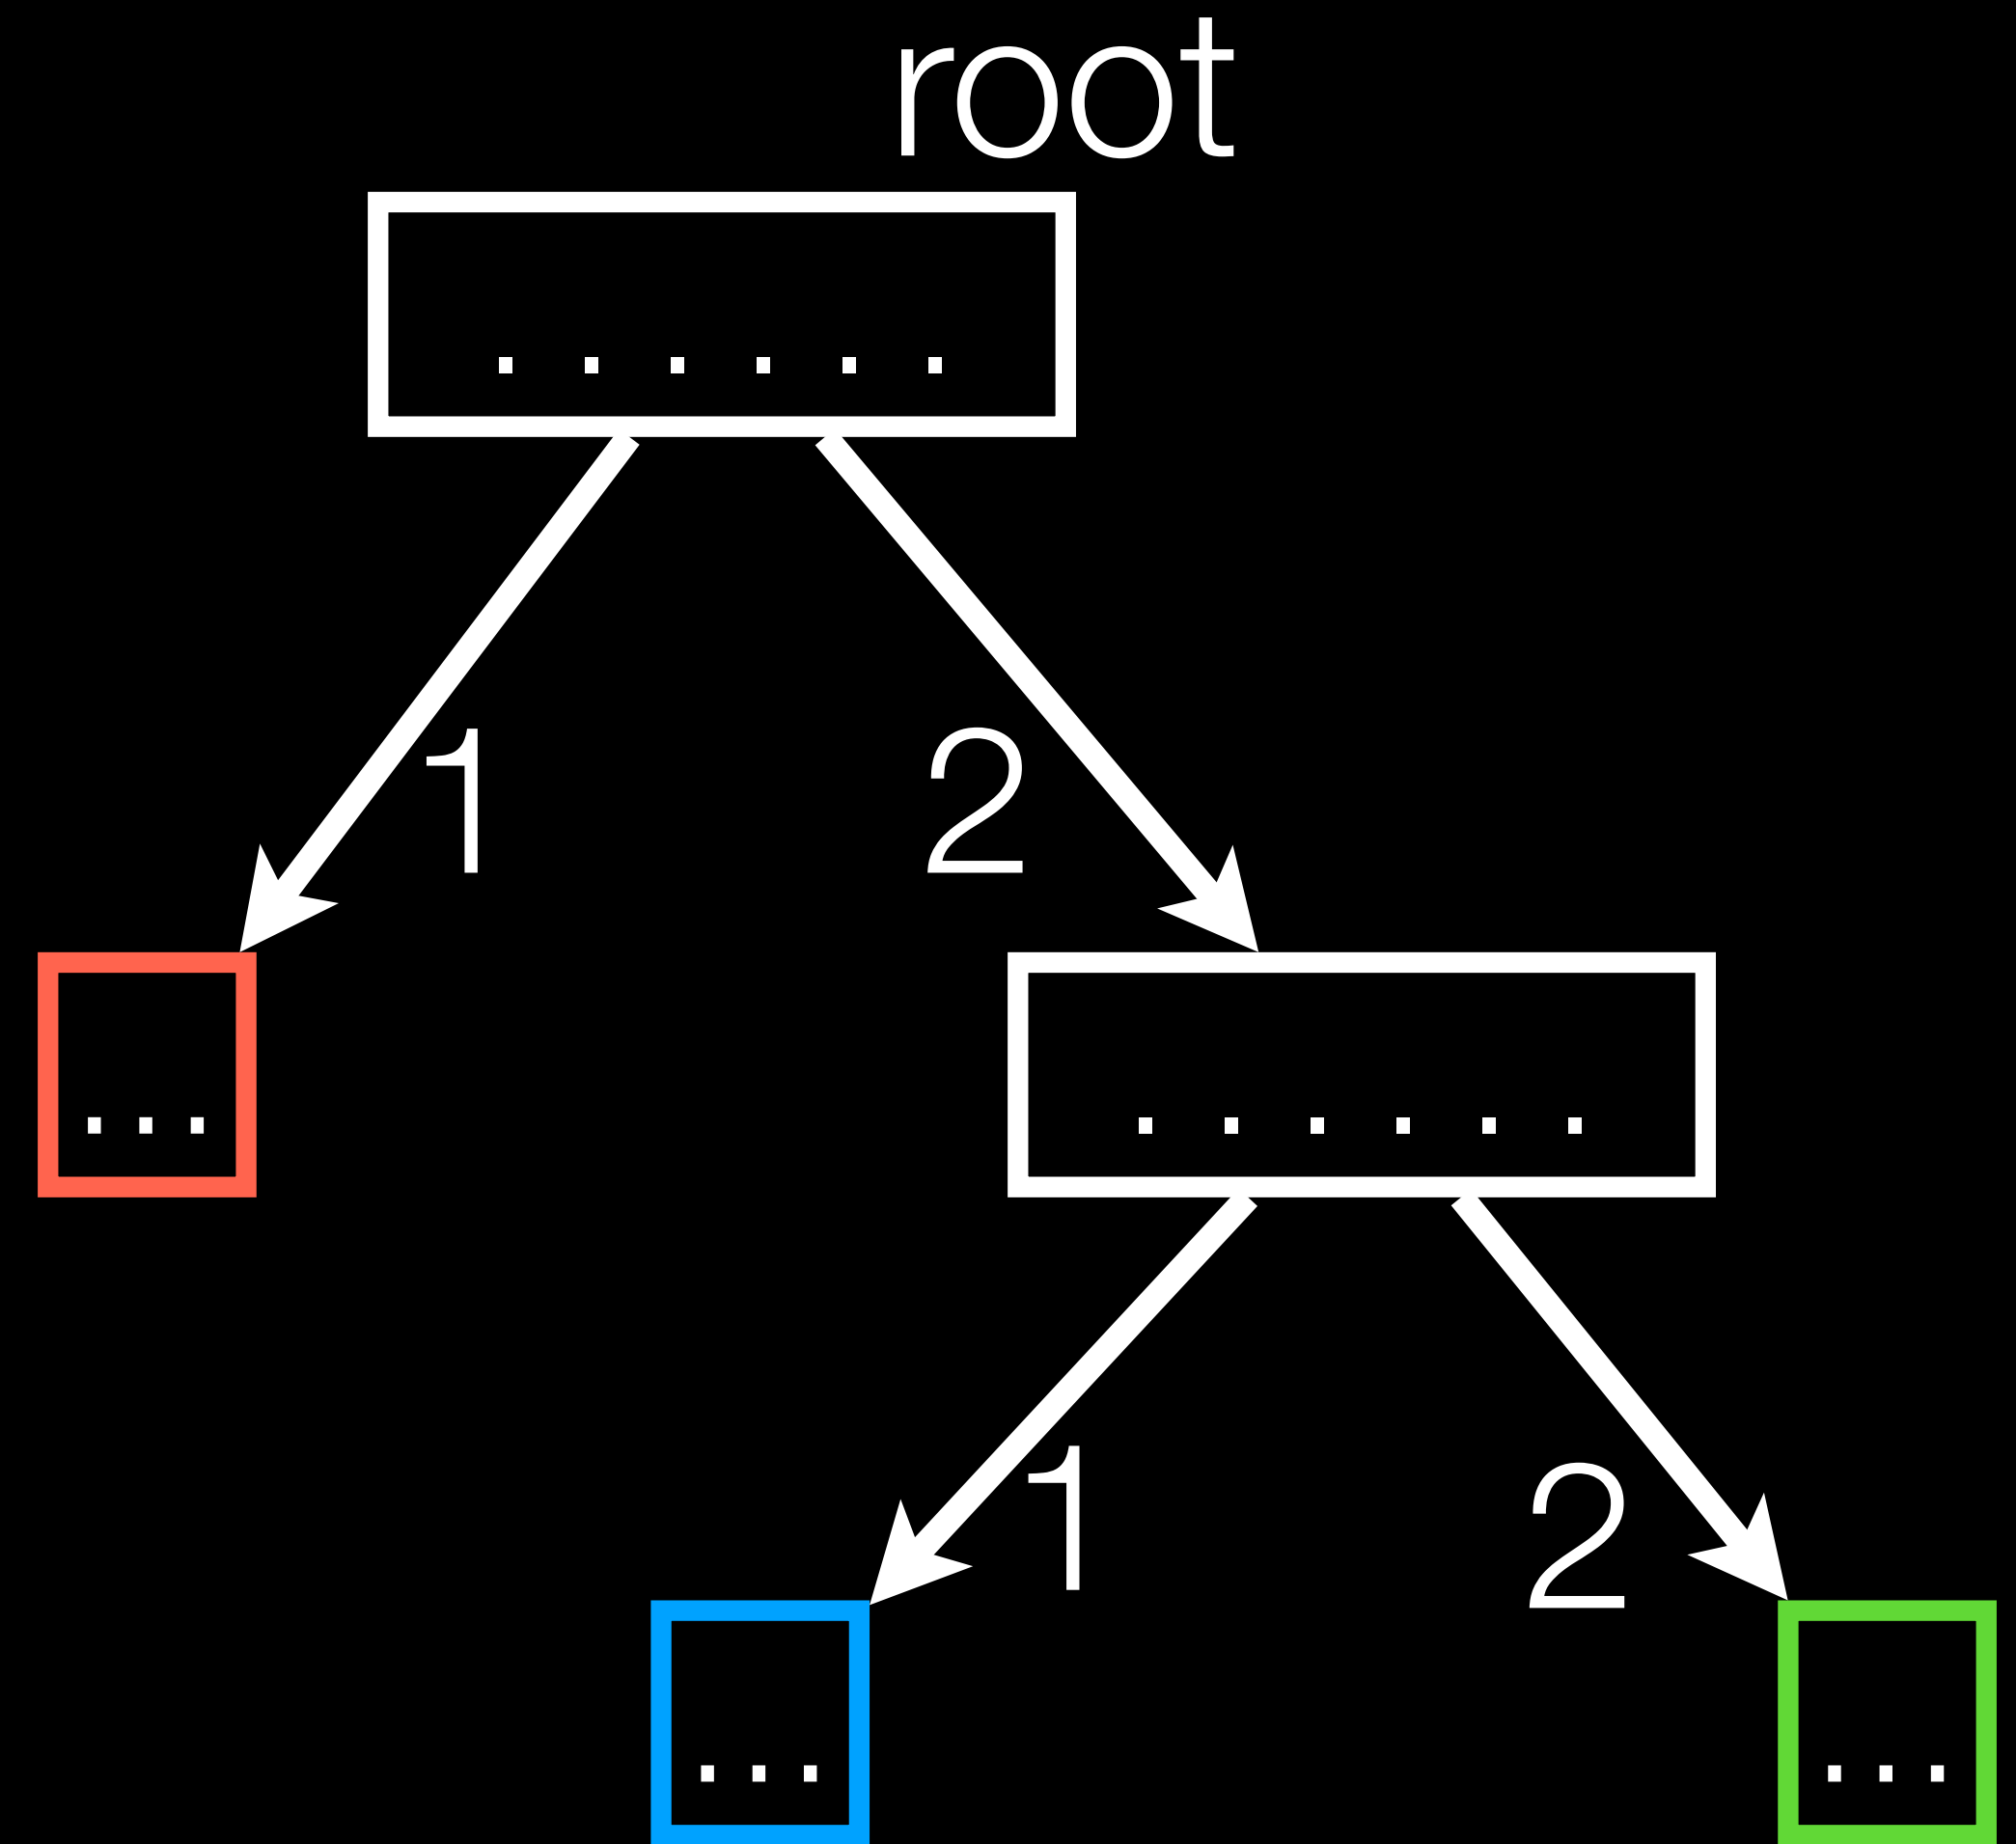
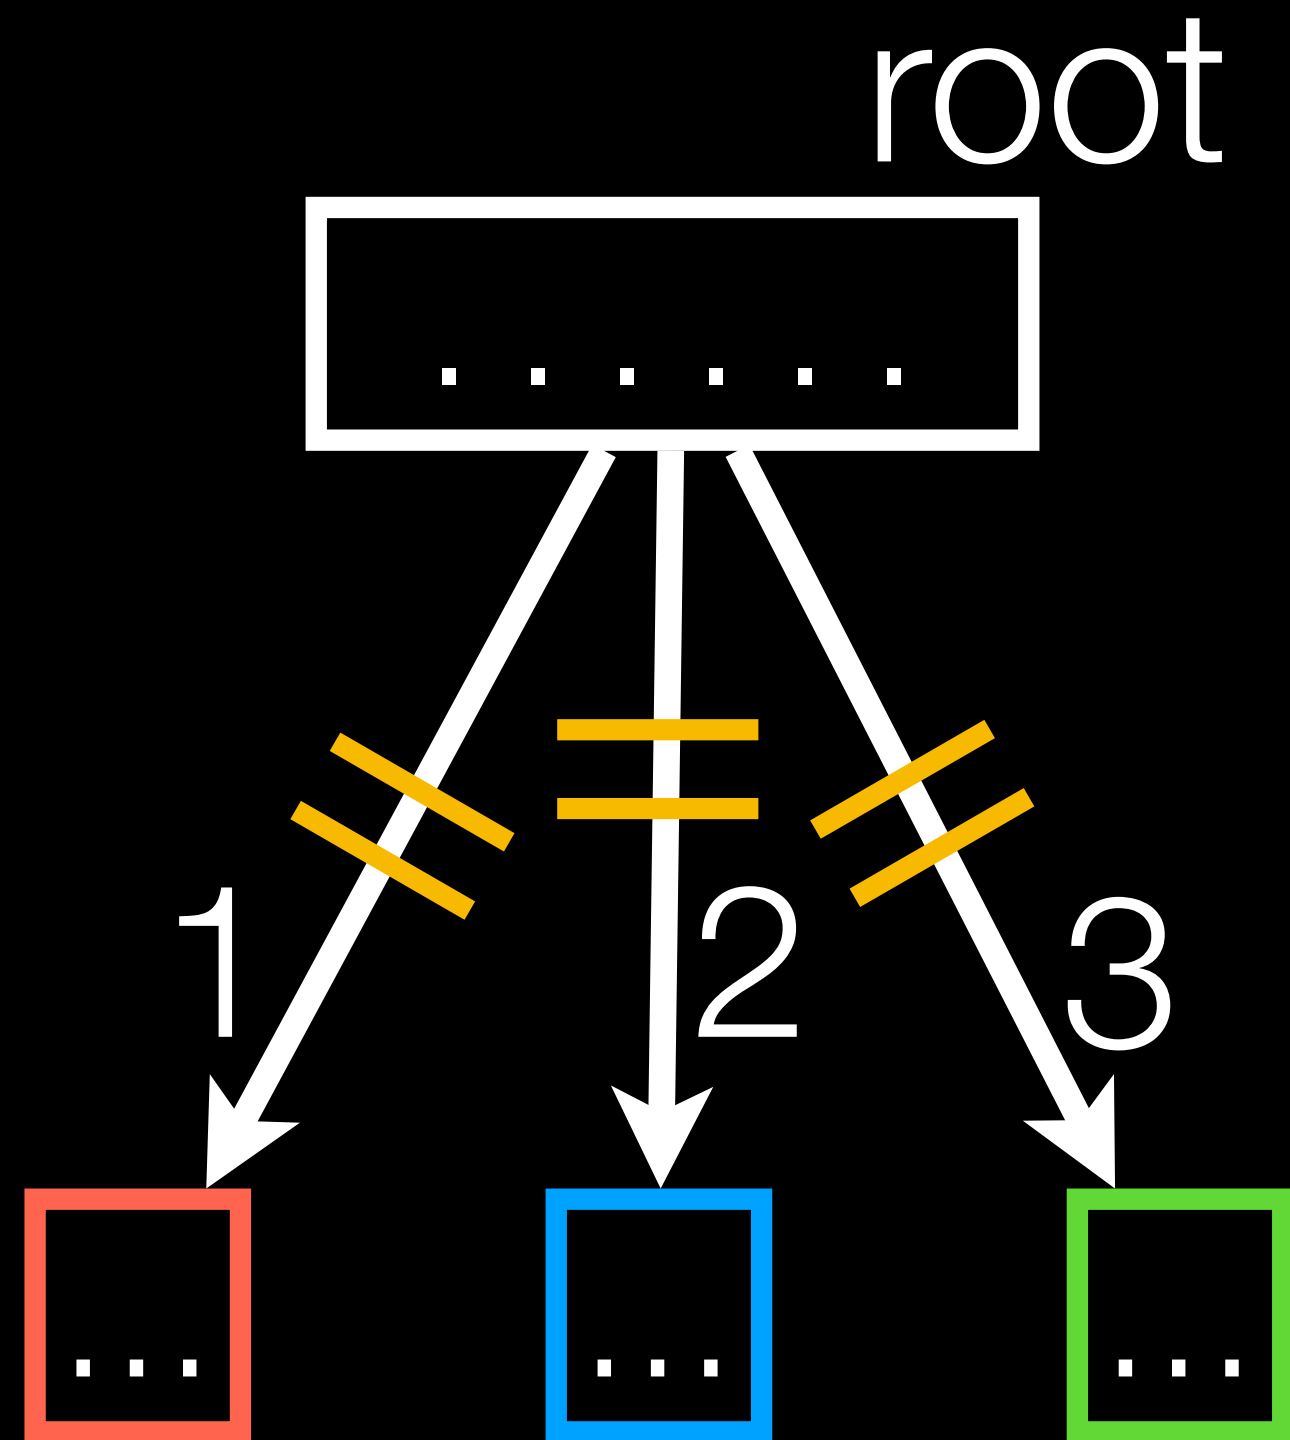
Compiling programs



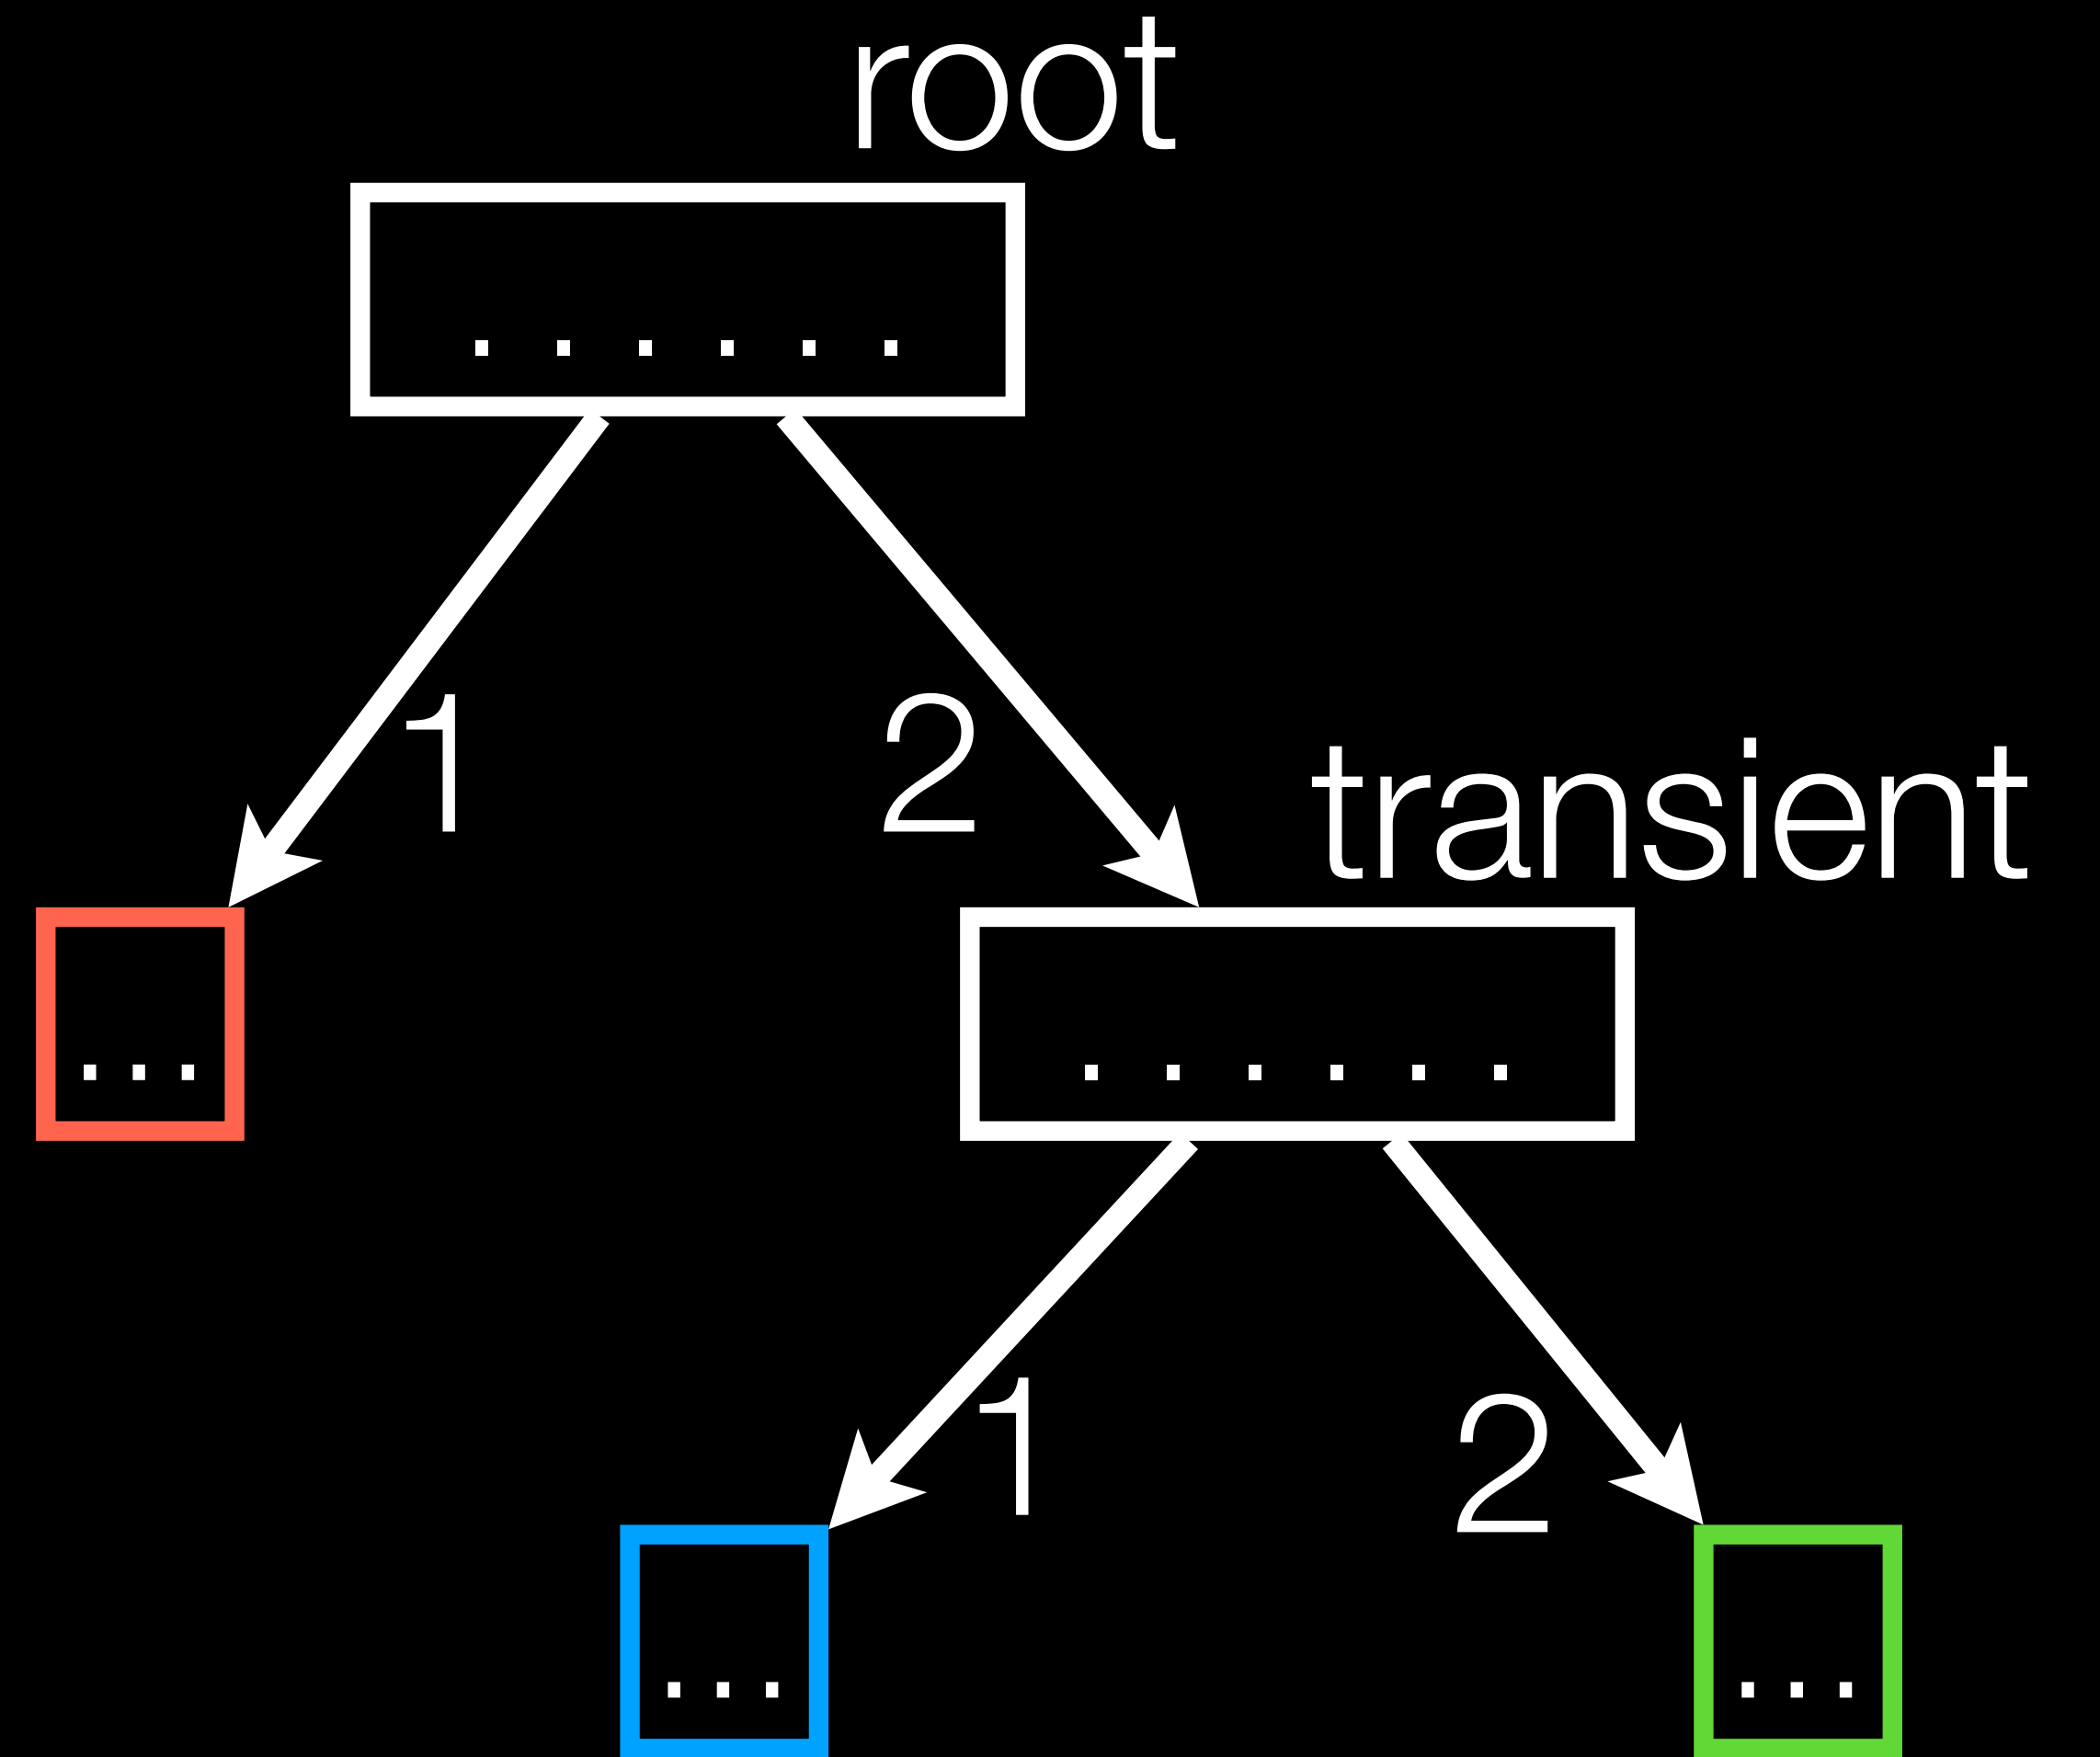
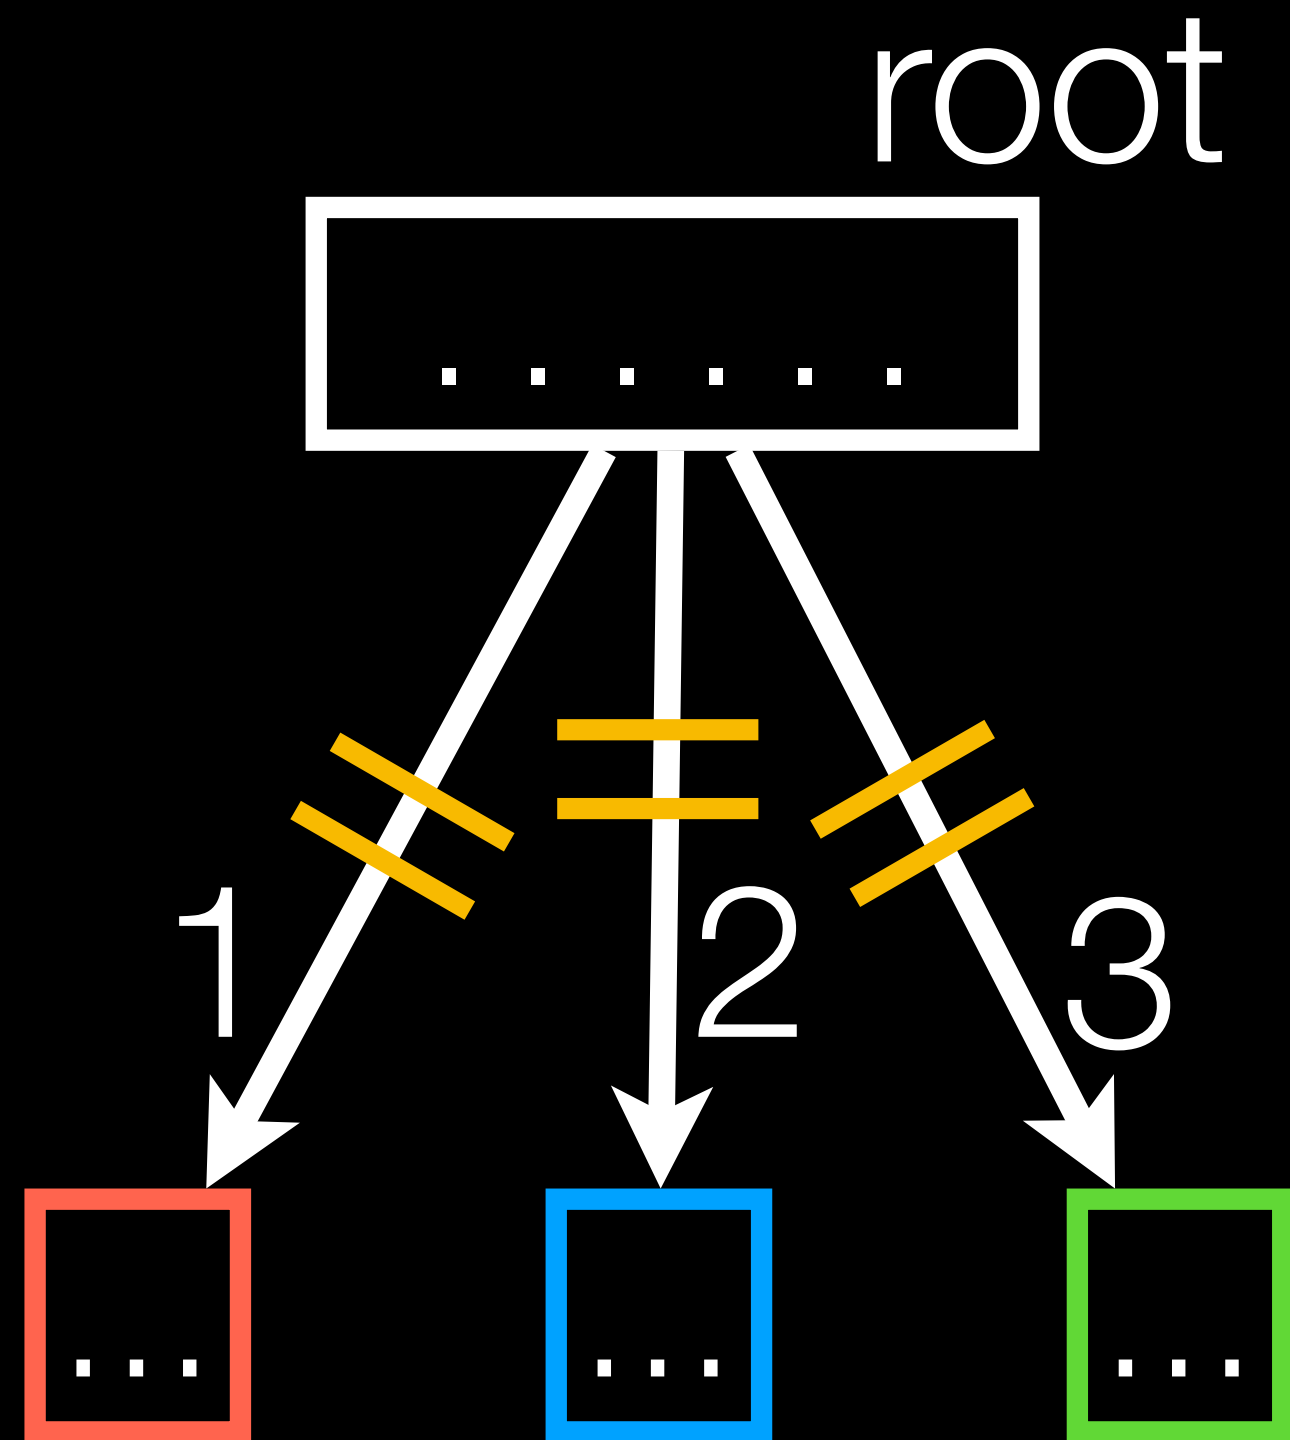
Compiling programs





Compiling programs



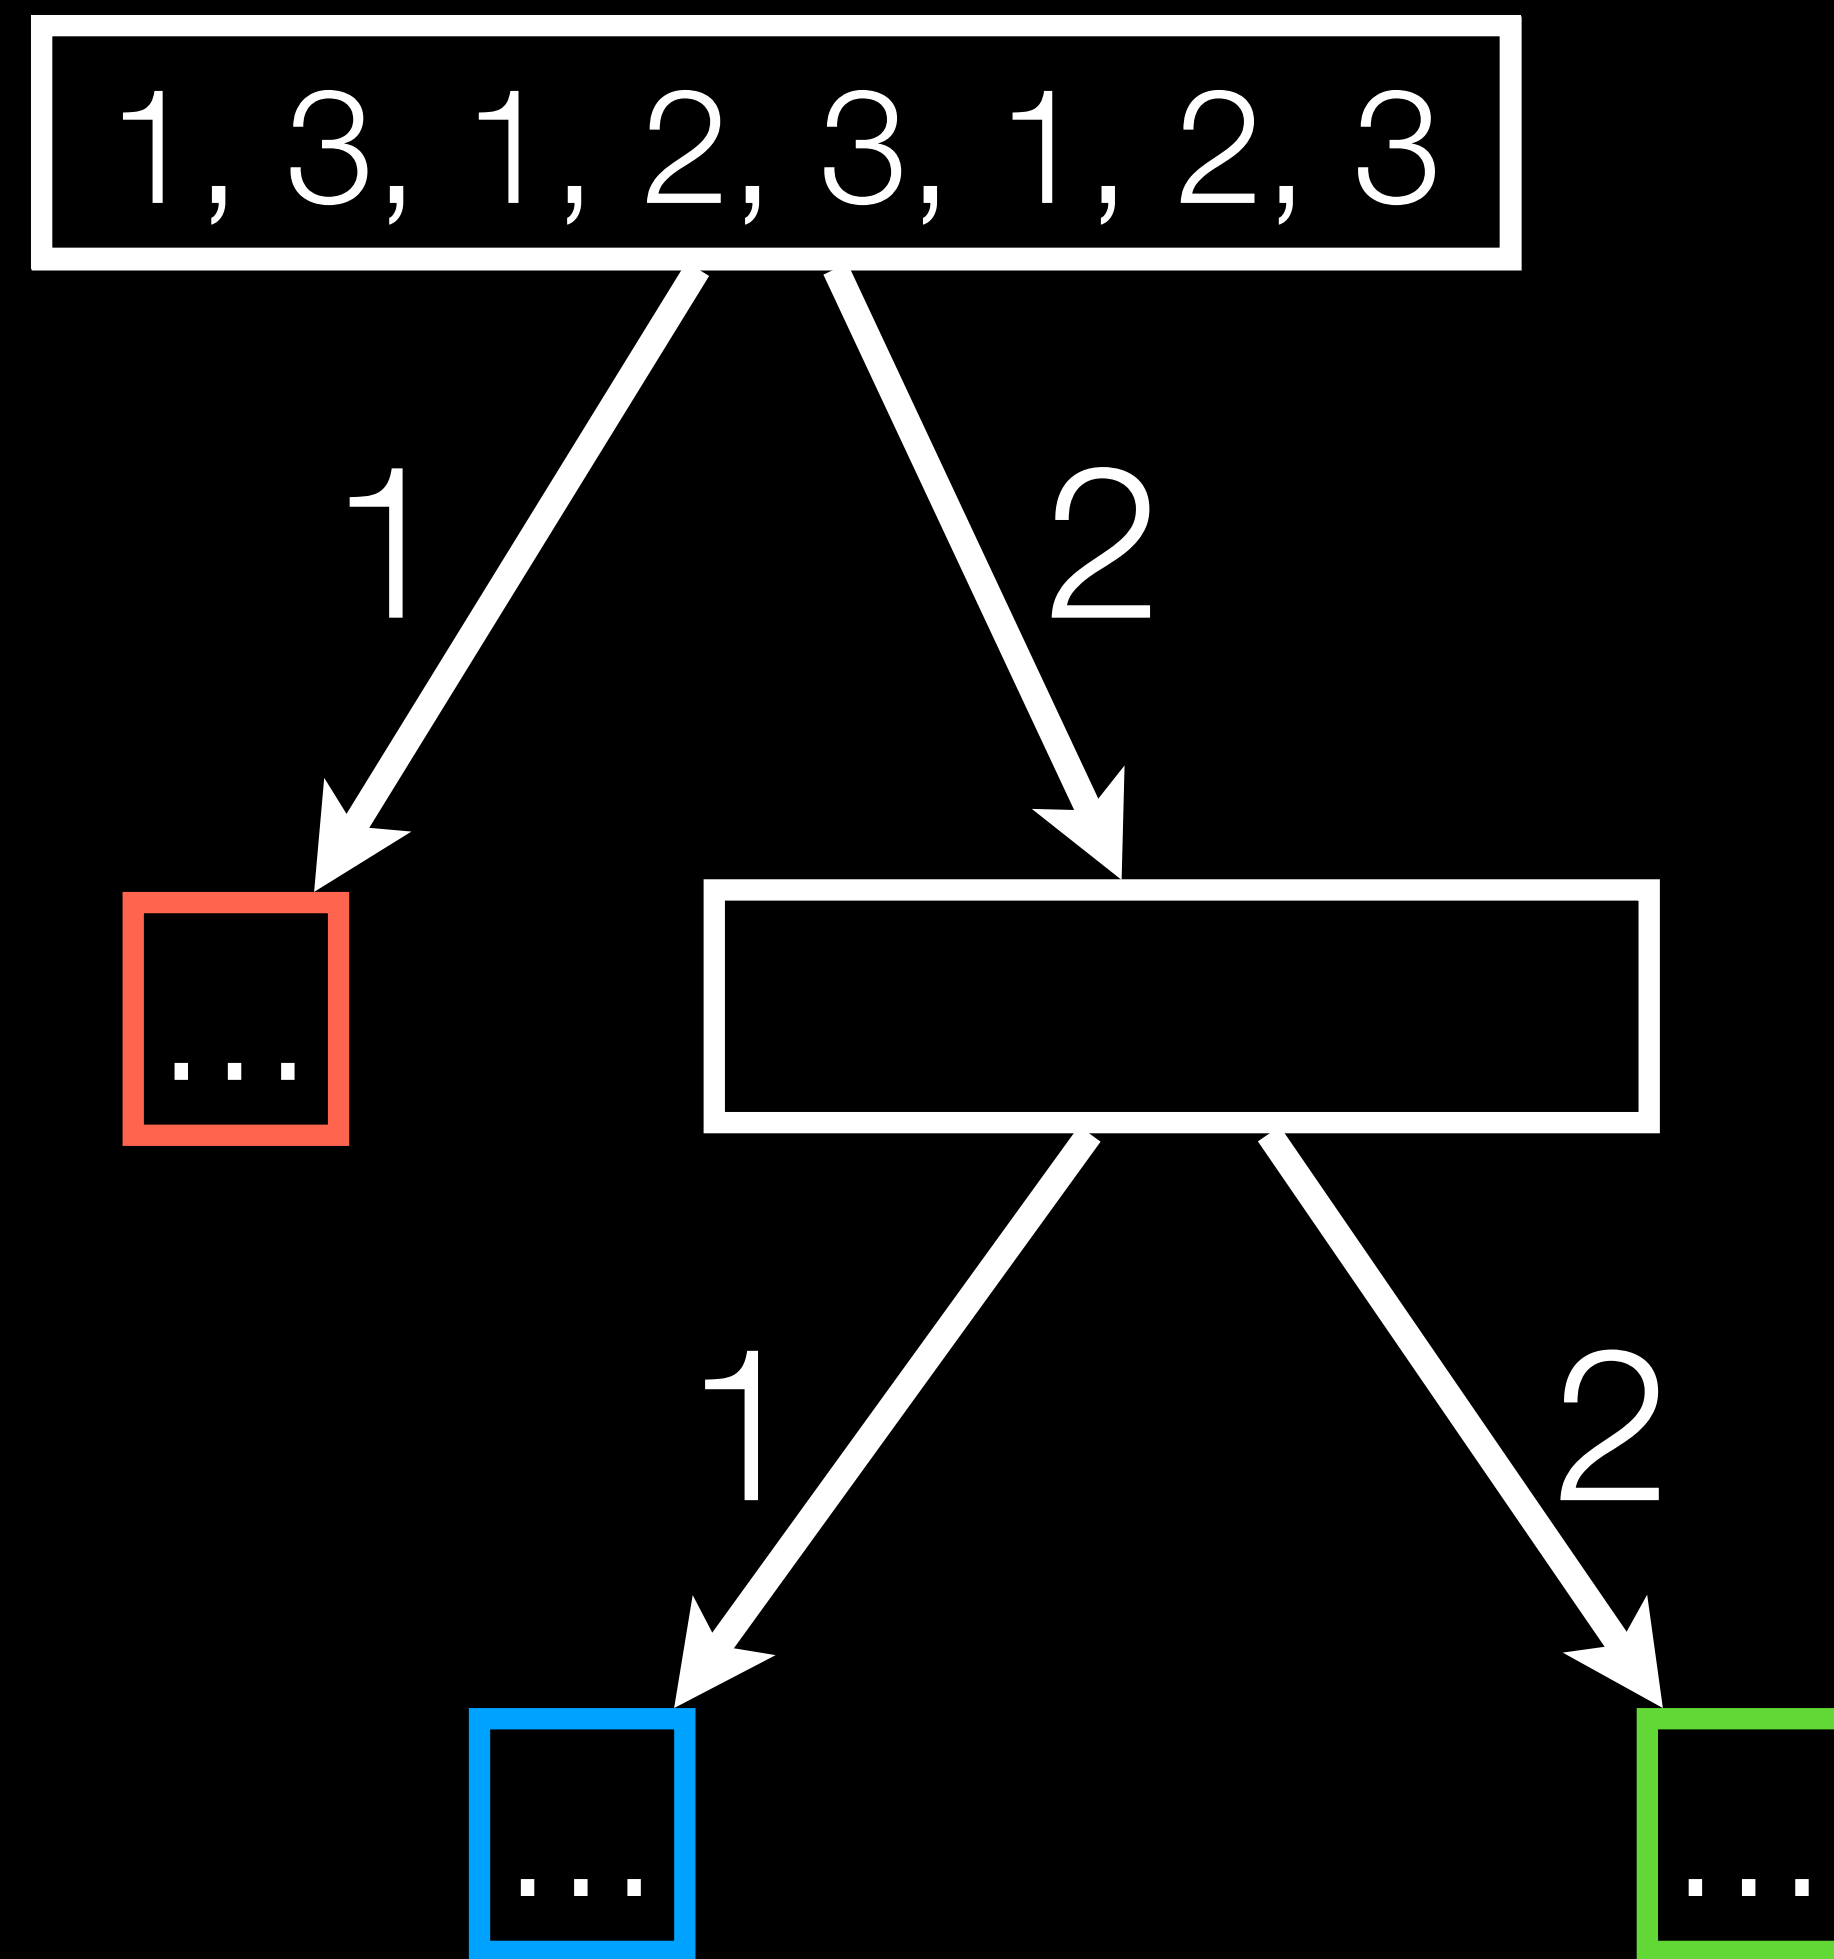
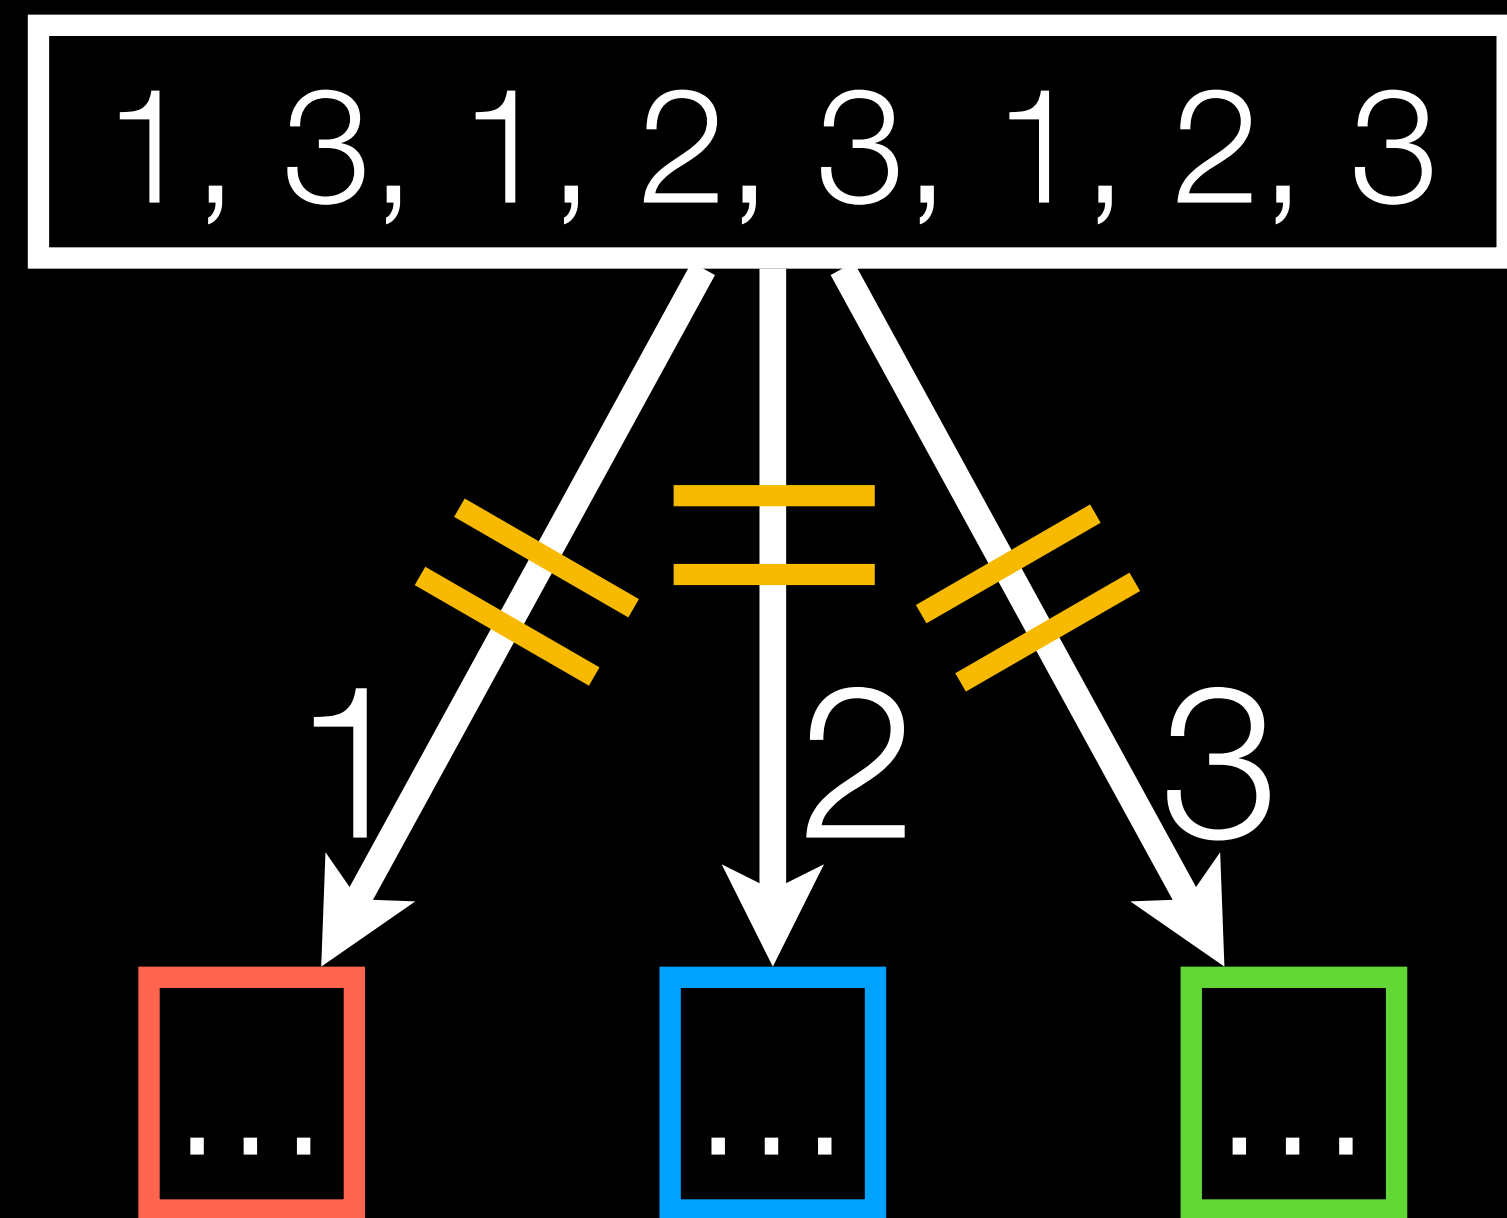
Compiling programs





Compiling programs





Compiling programs





Compiling programs





Compiling programs





Compiling programs





Compiling programs




Compiling programs



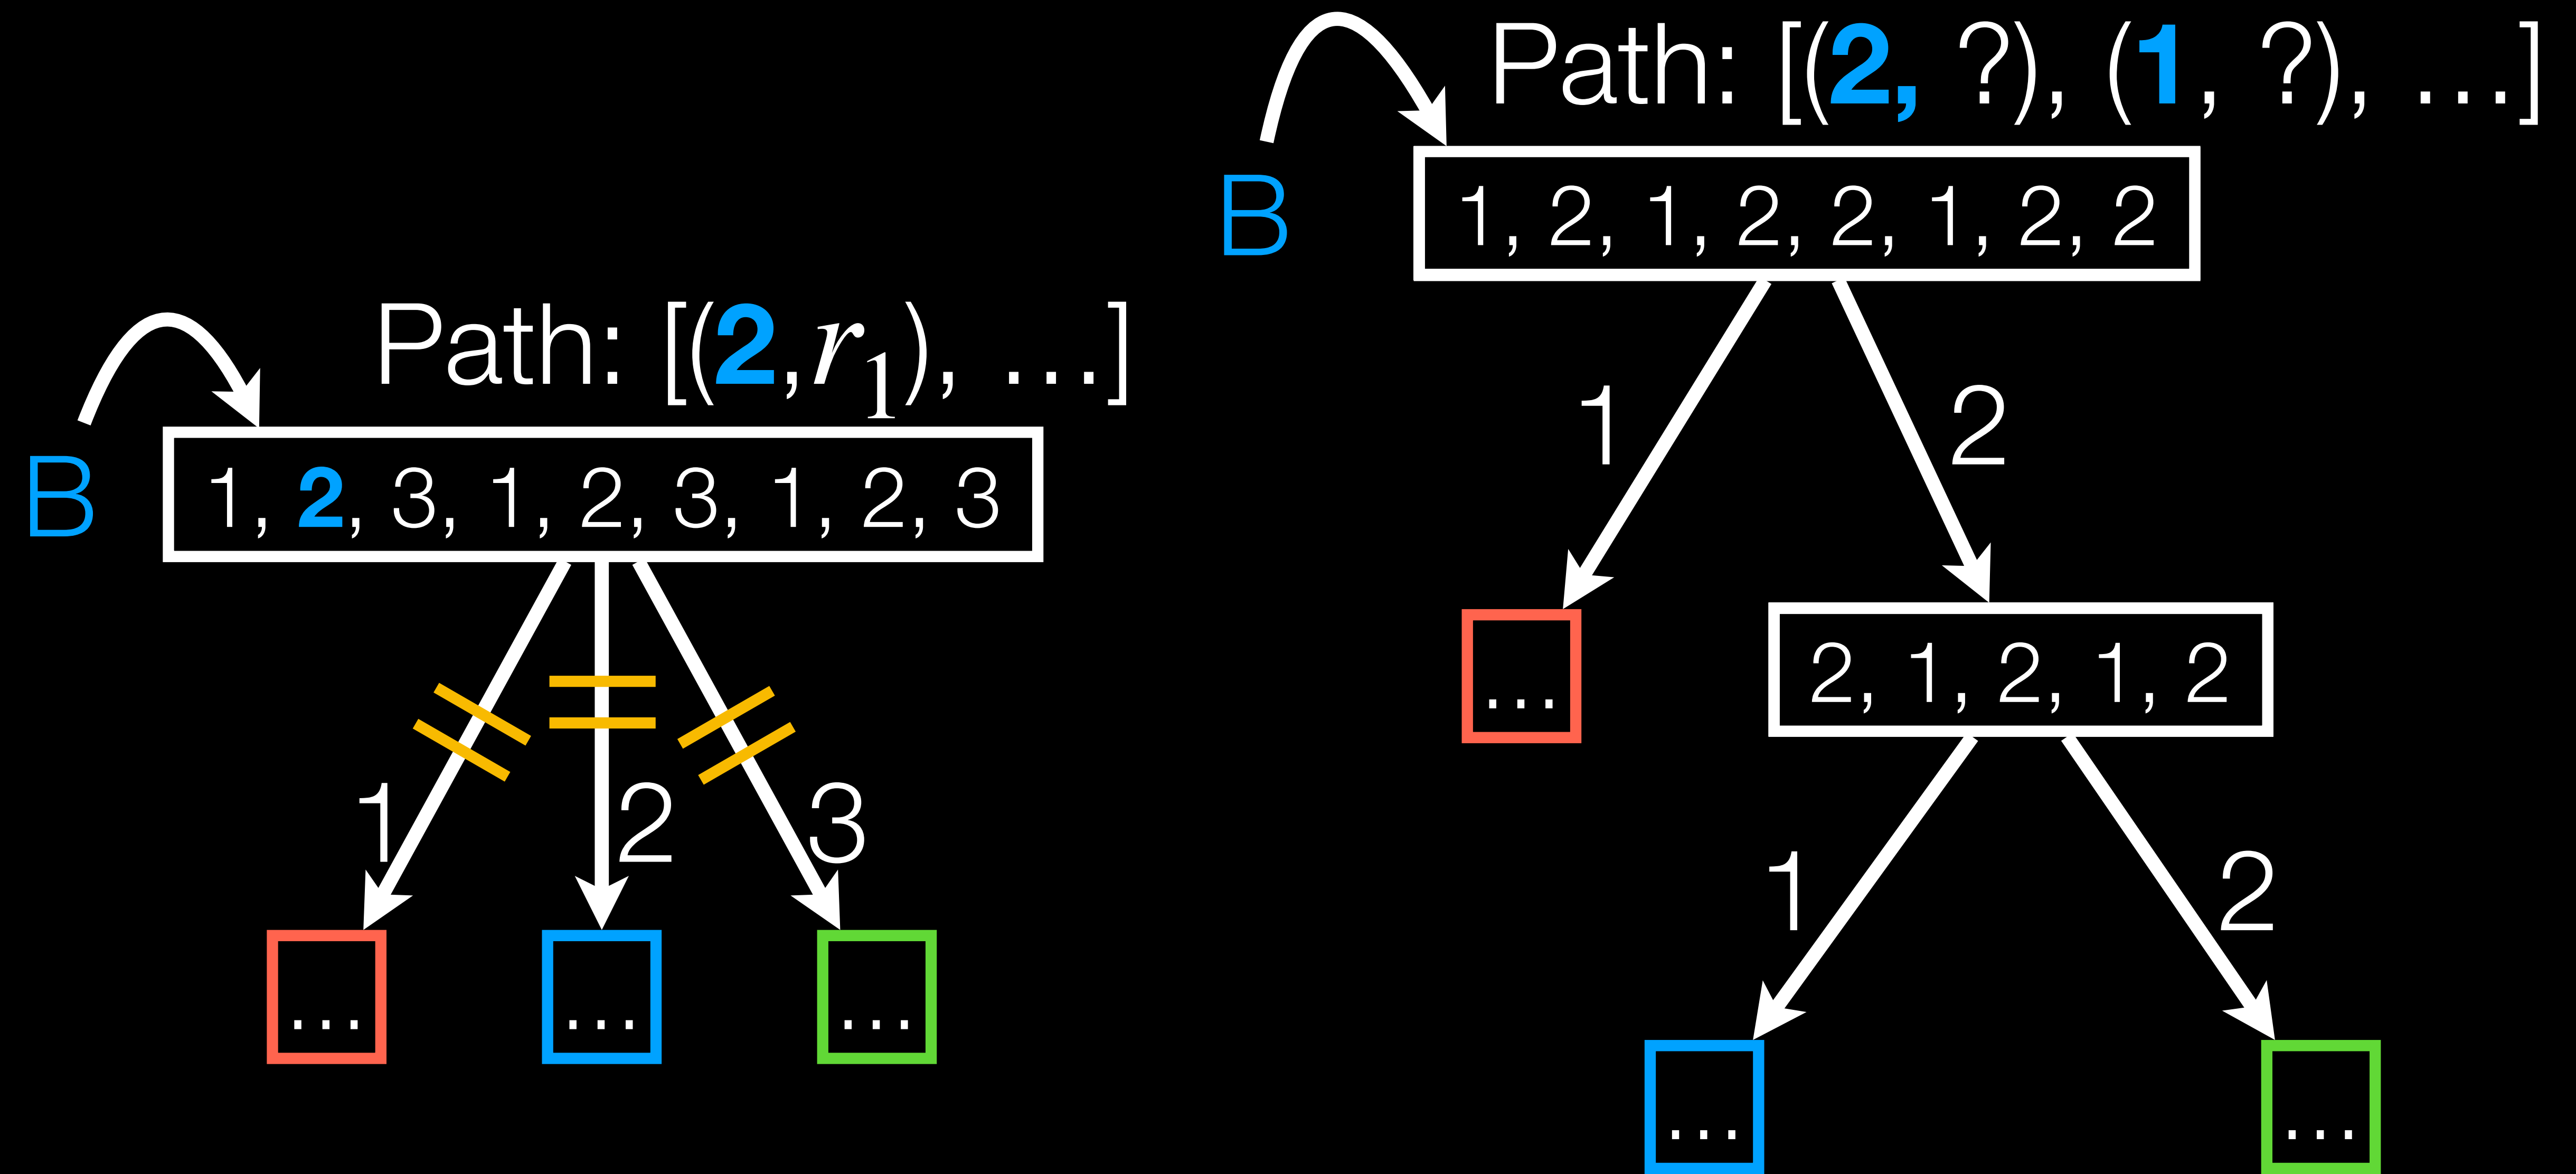
Compiling programs



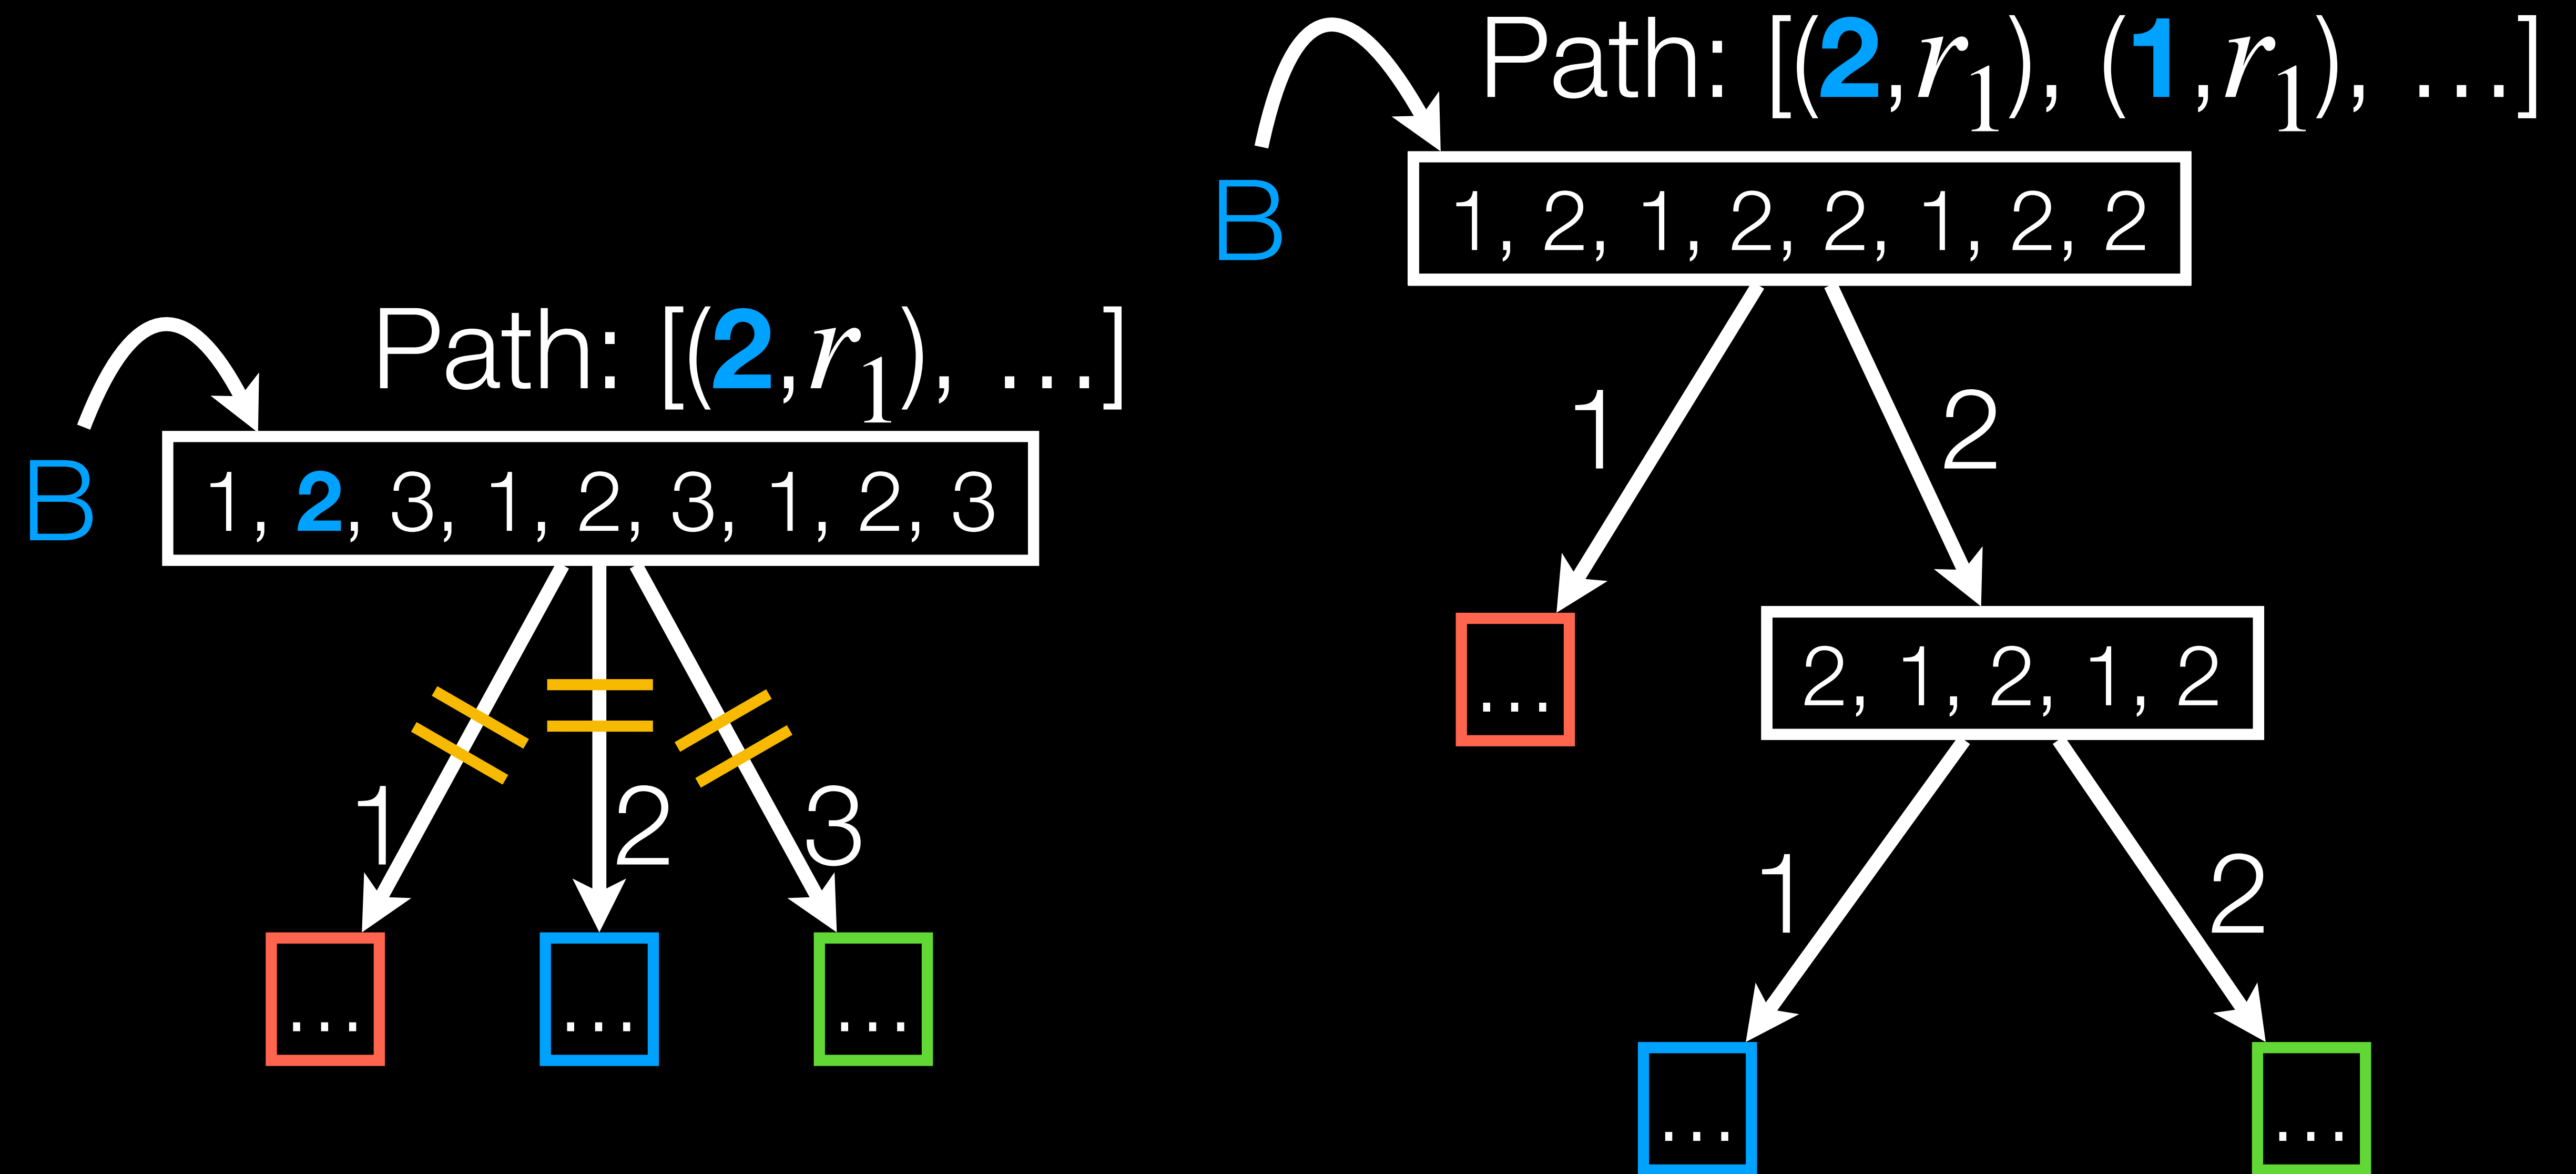
Compiling programs



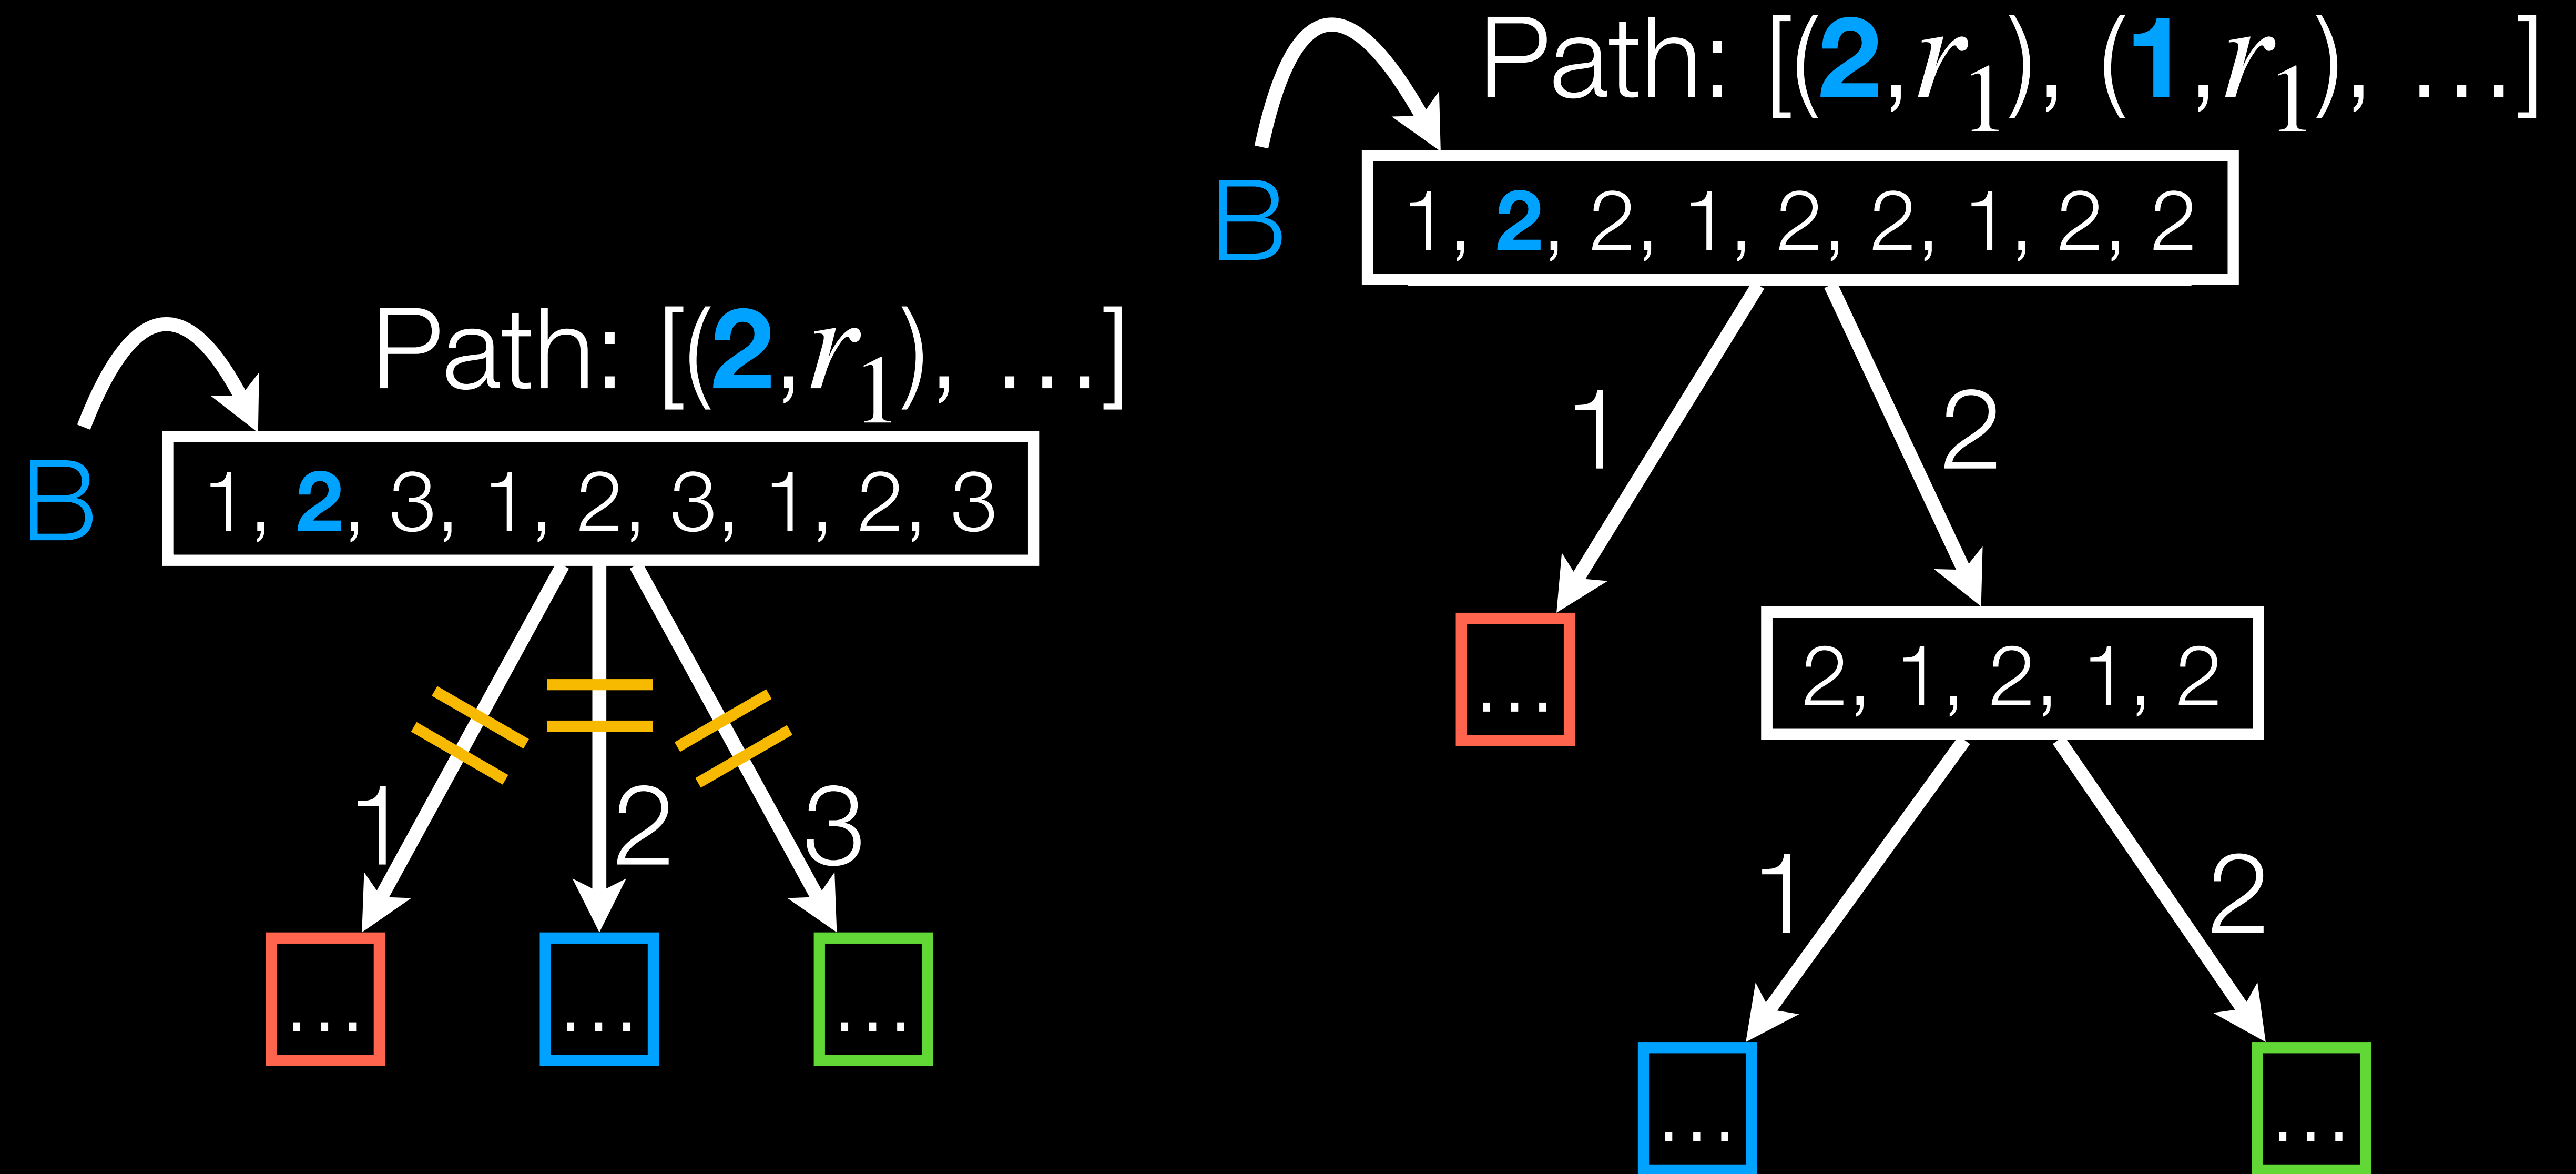
Compiling programs



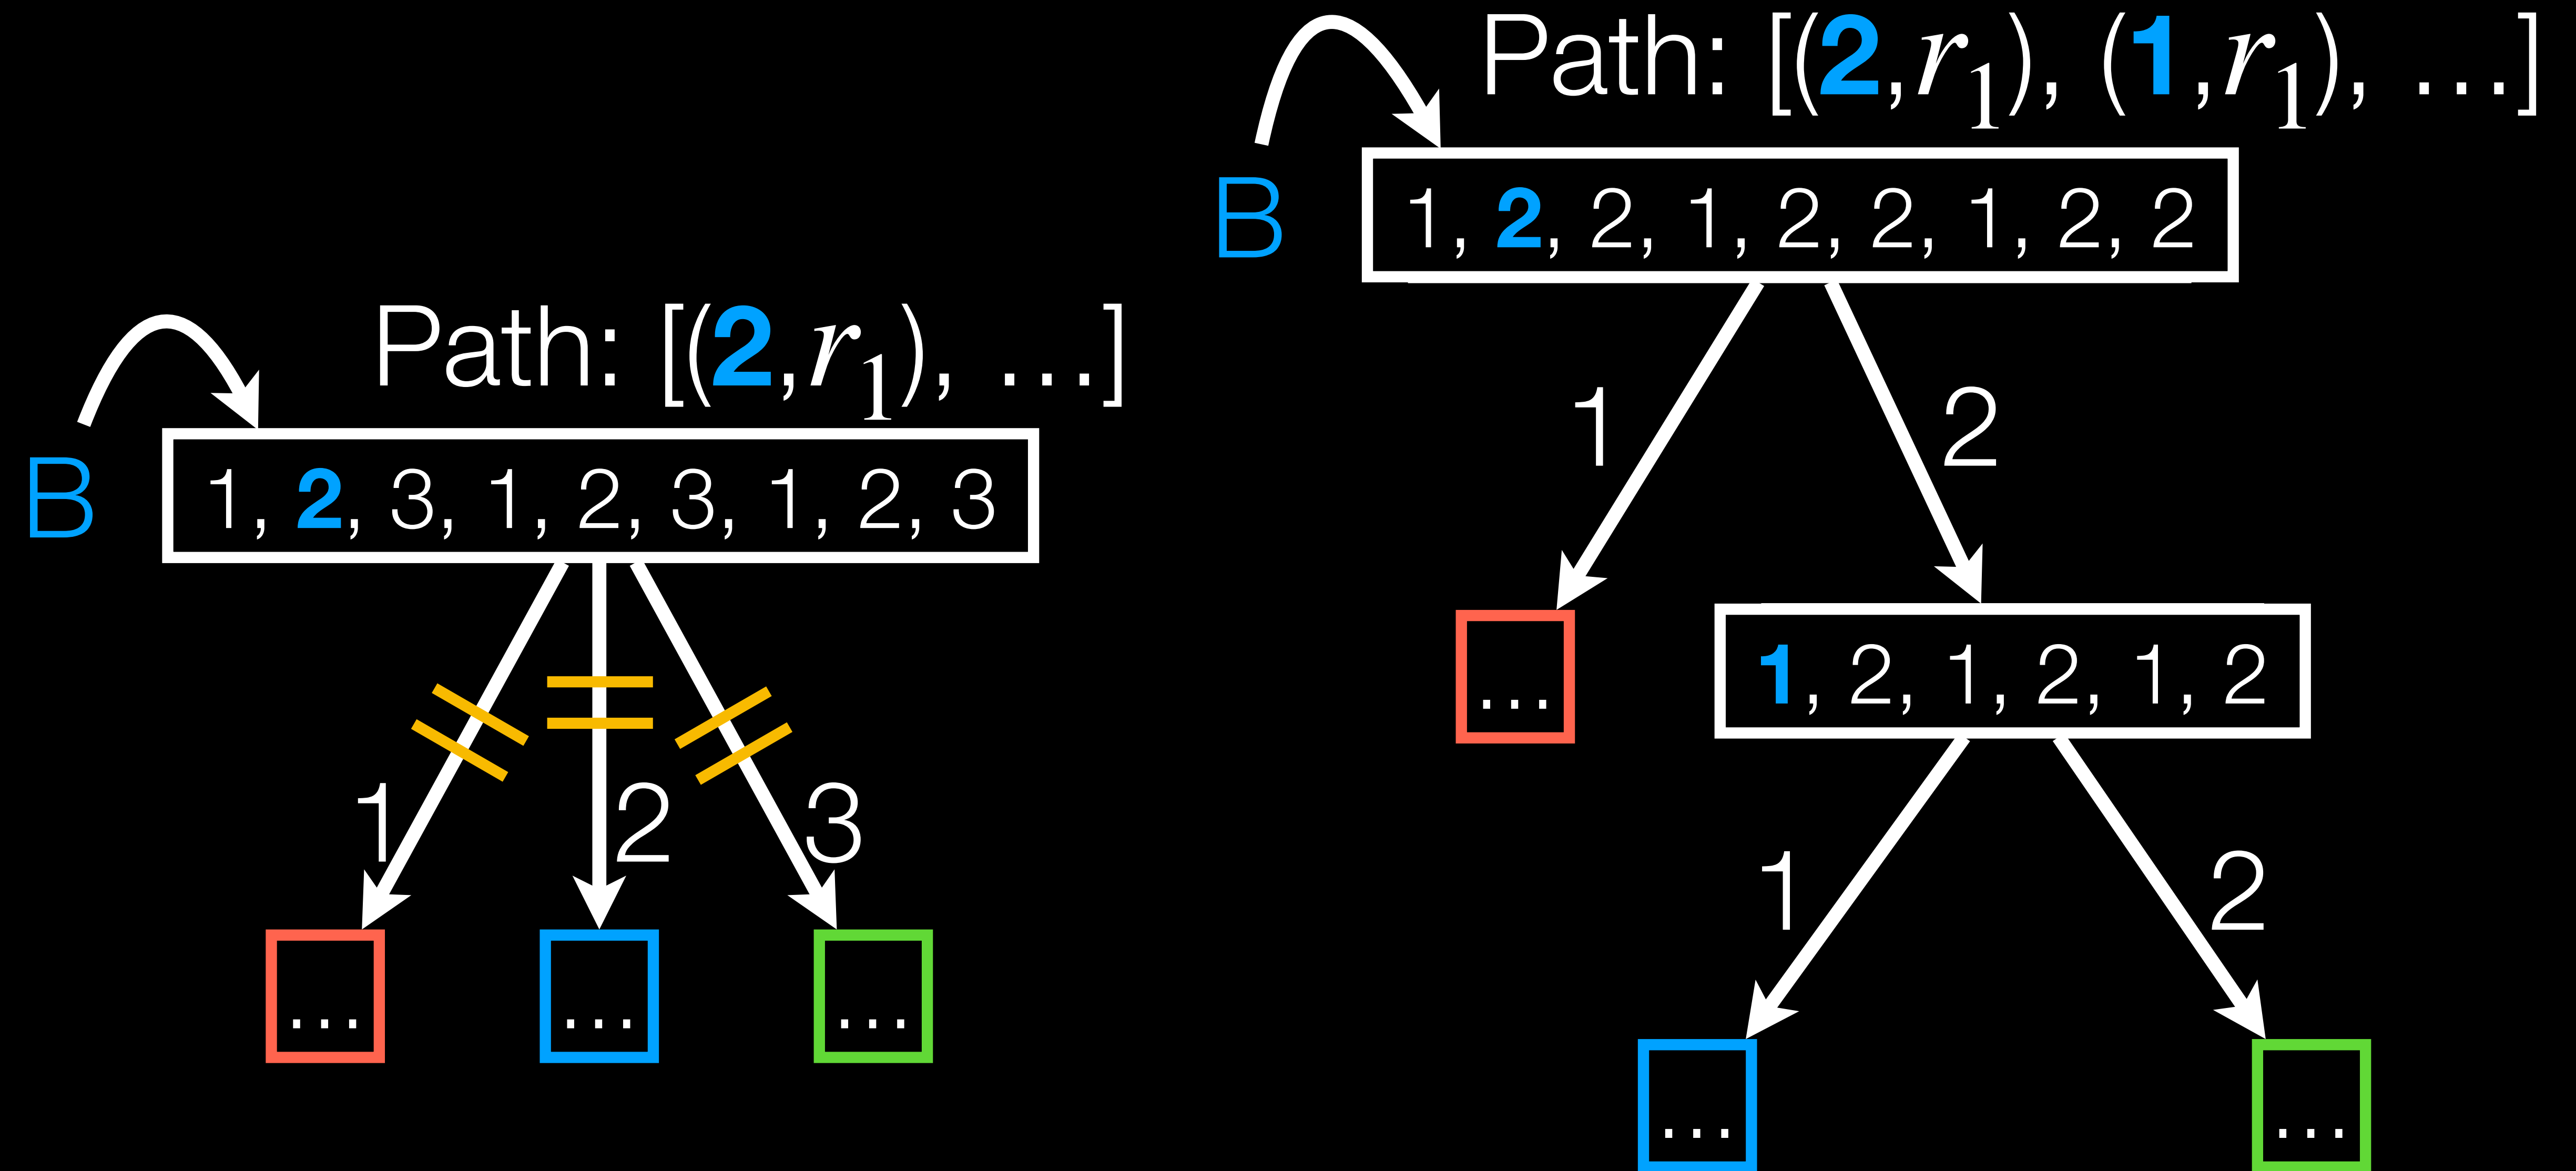
Compiling programs



Compiling programs



Compiling programs





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

Path: [**2**, r_1), (**1**, r_1), ...]

B 1, **2**, 2, 1, 2, 2, 1, 2, 2

Path: [**2**, r_1), ...]

B 1, **2**, 3, 1, 2, 3, 1, 2, 3

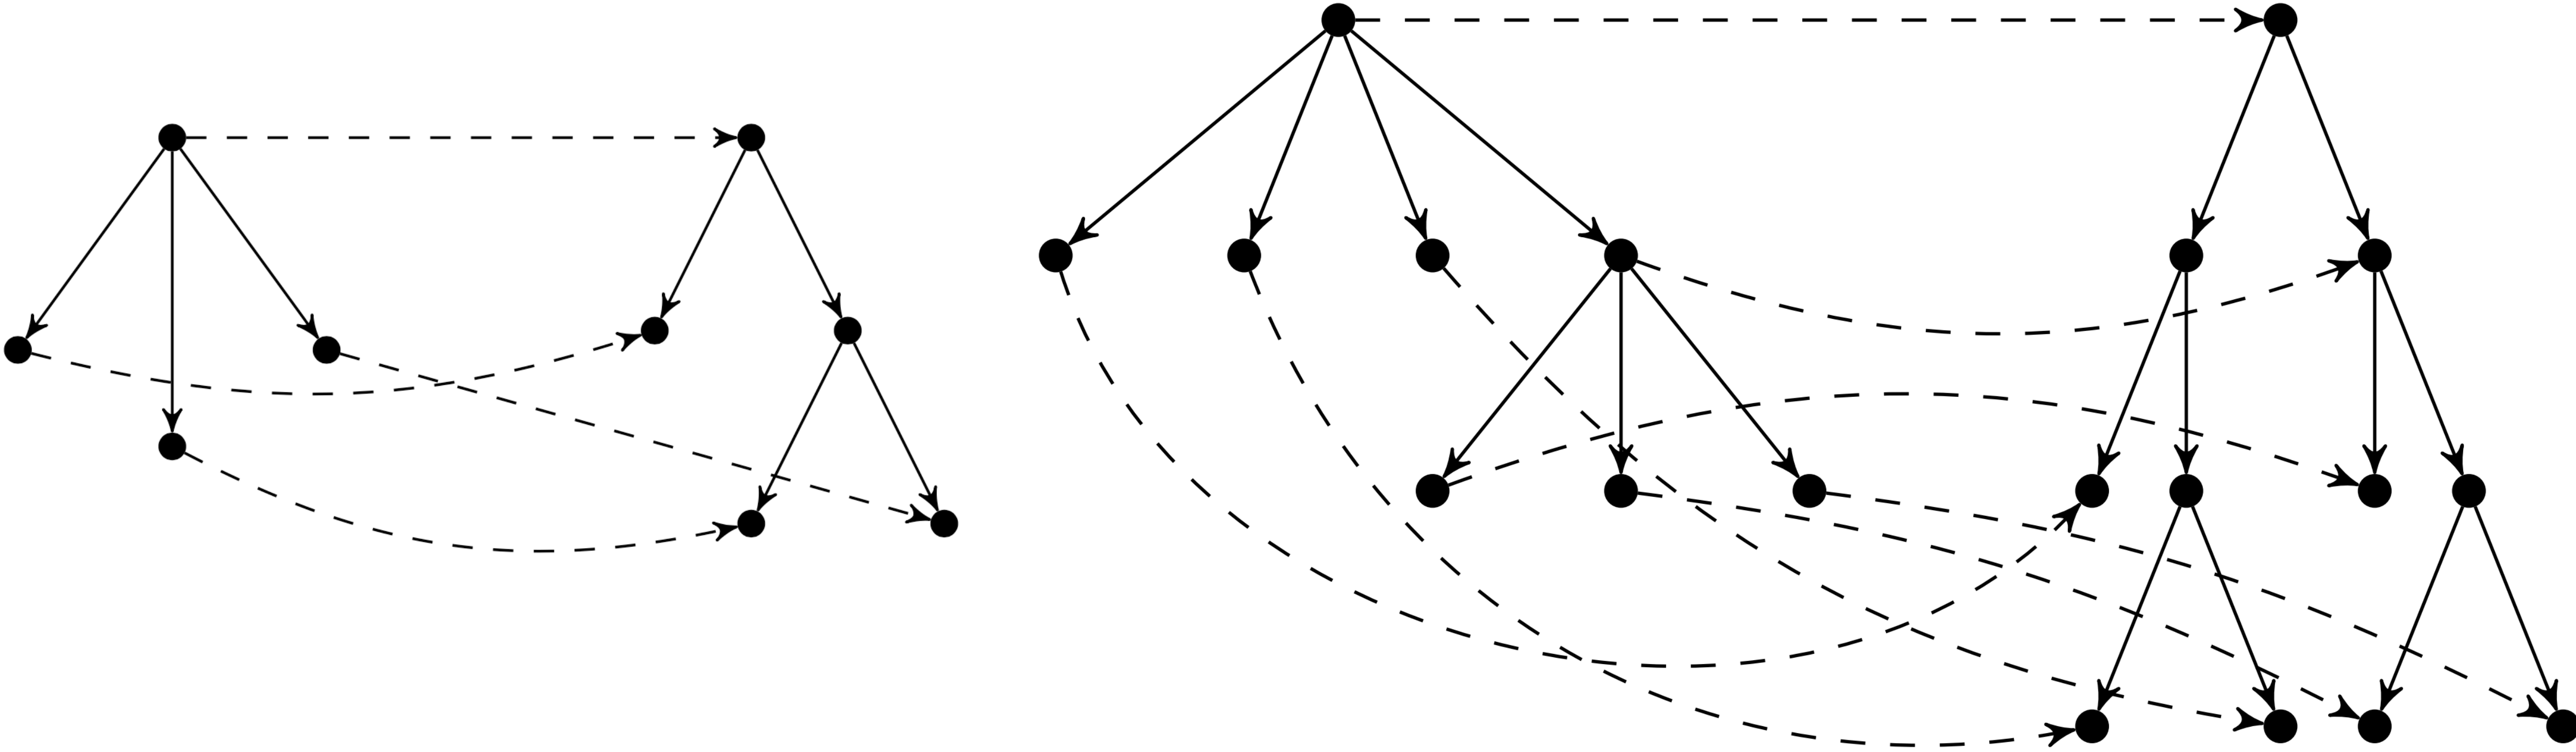


Generating embeddings automatically!

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.

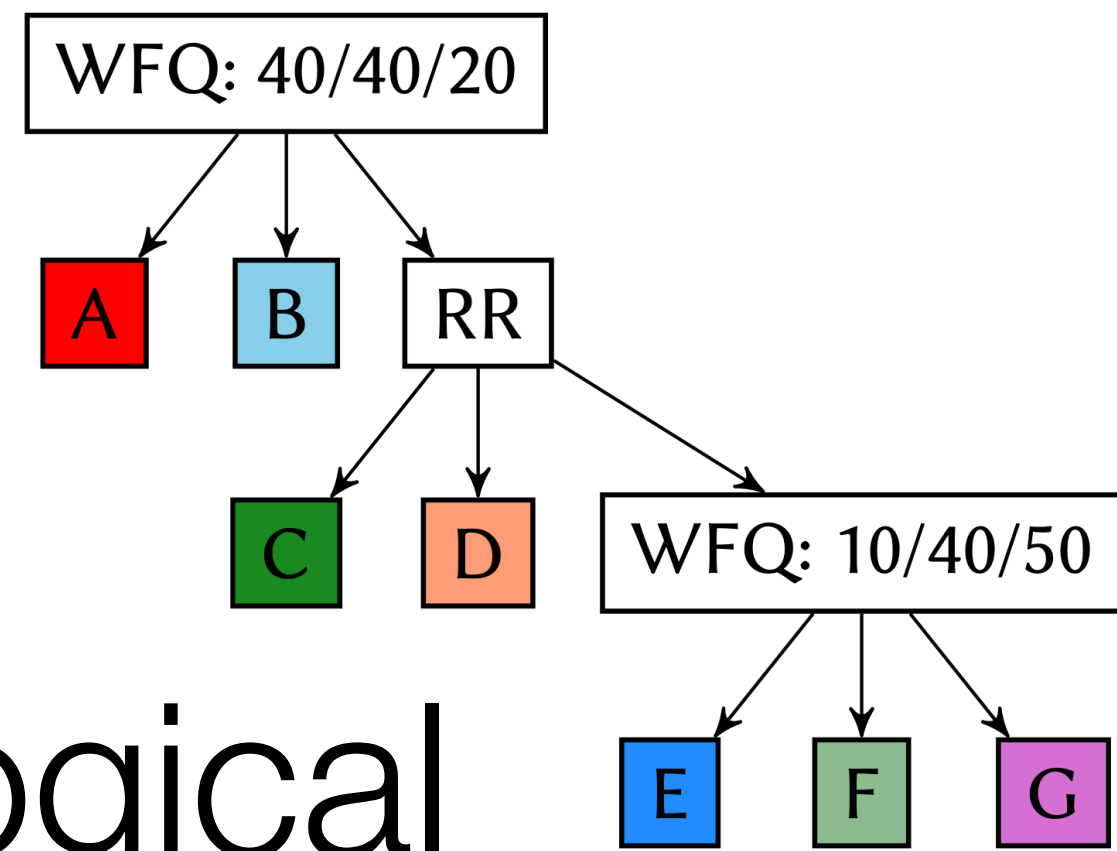
Two new algorithms,
both starting with heterogeneous source trees.

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

Workflow




Workflow



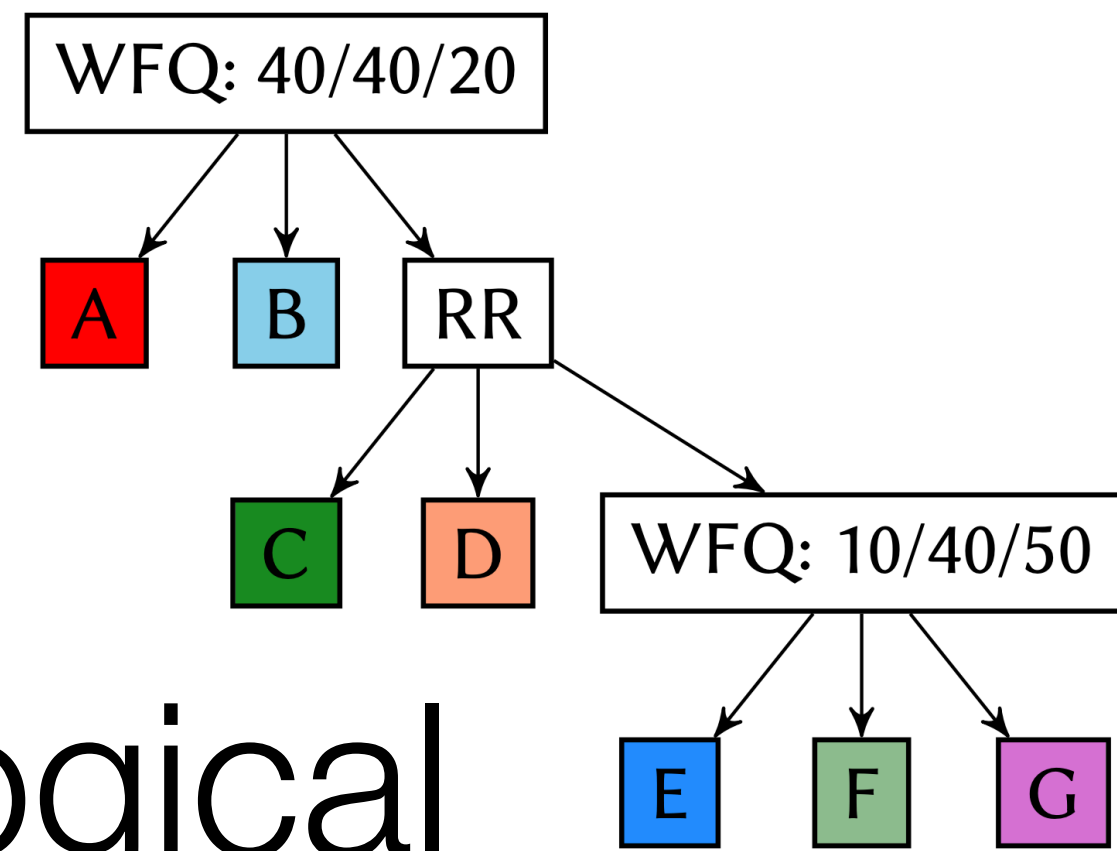
logical

Workflow



But the hardware supports
a regular-branching binary tree.

Workflow



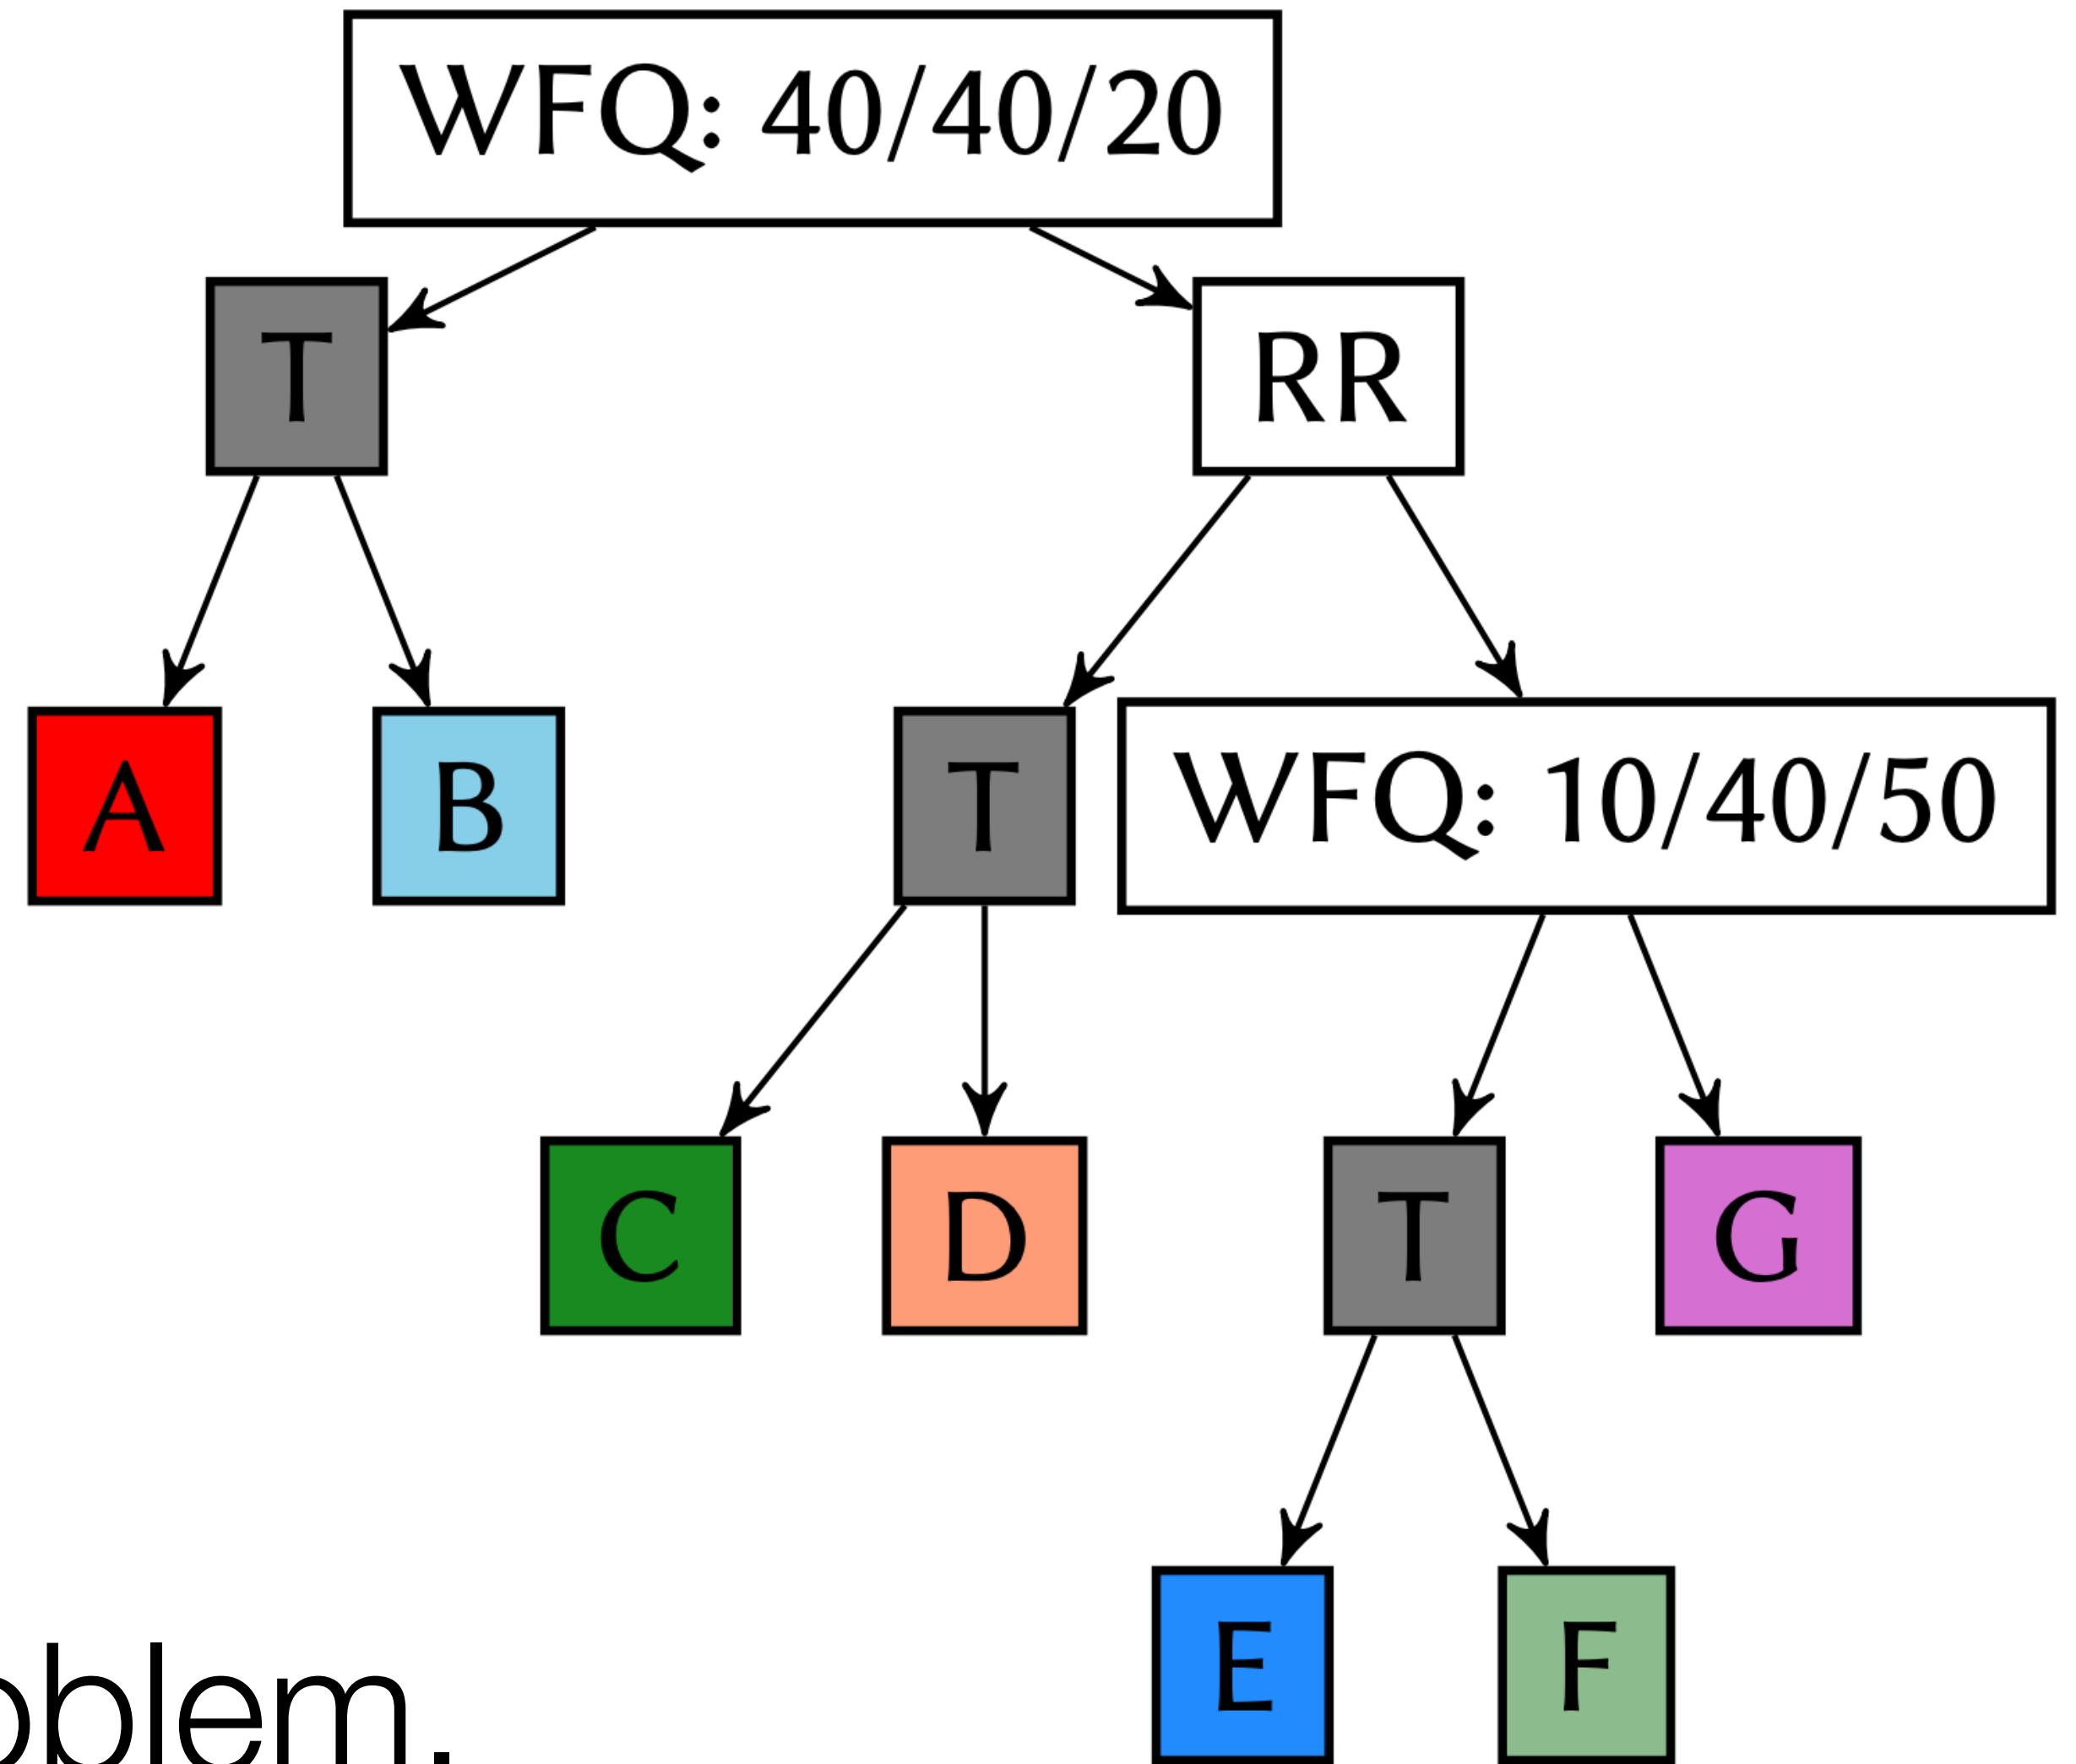

logical

But the hardware supports a regular-branching binary tree.

No problem.

Here's how I'll use that tree.

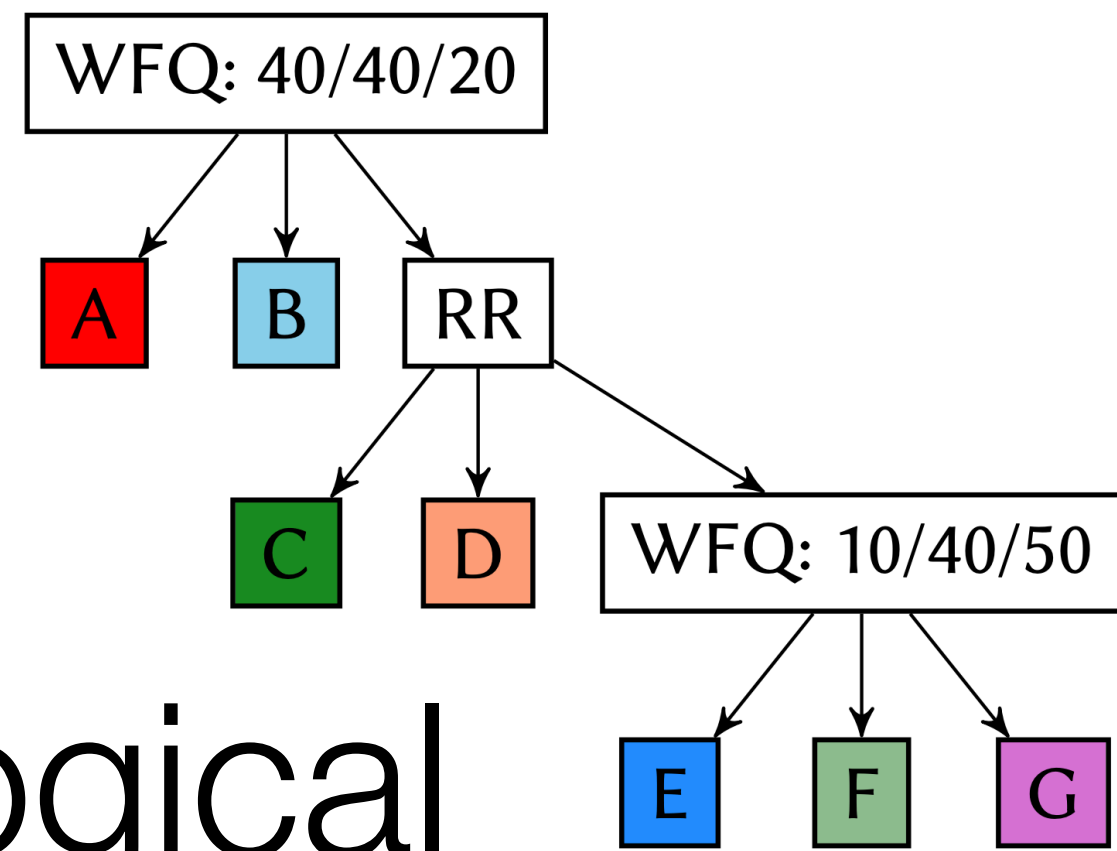
Workflow




No problem.

Here's how I'll use that tree.

Workflow




logical




actual

Simulation


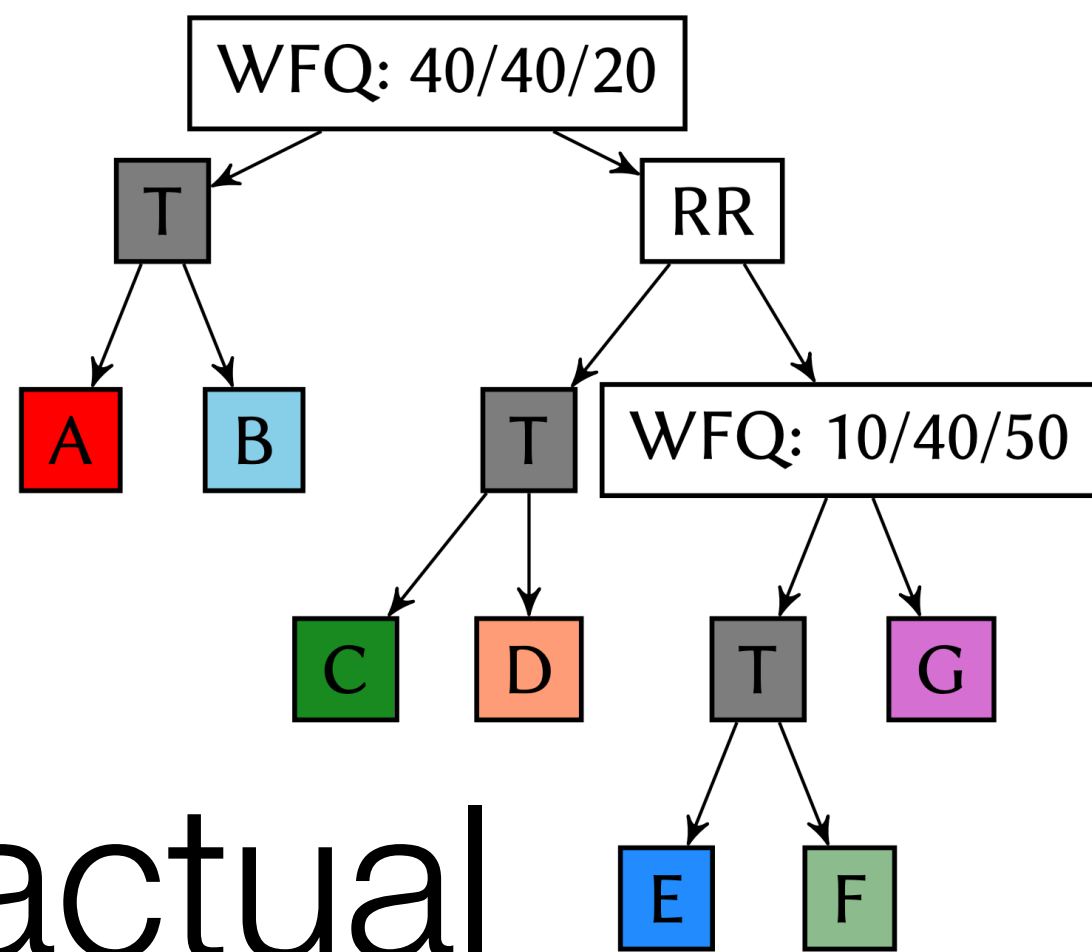
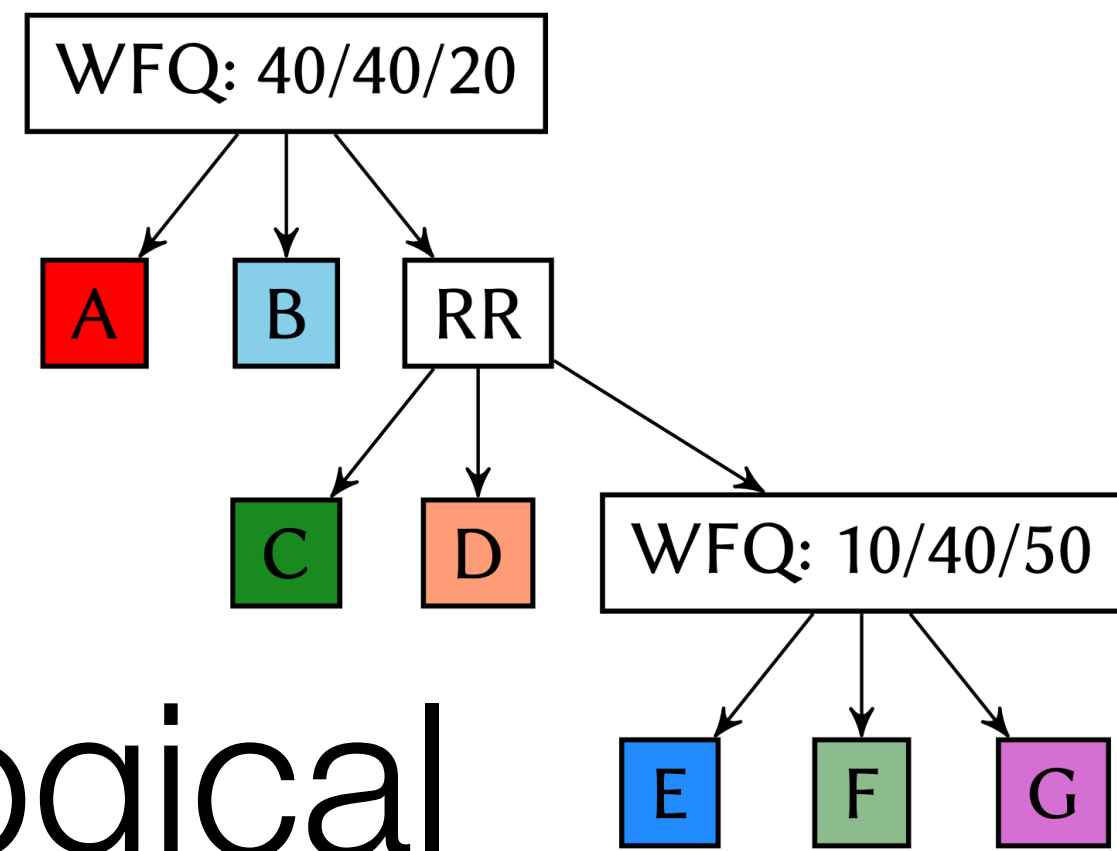


logical

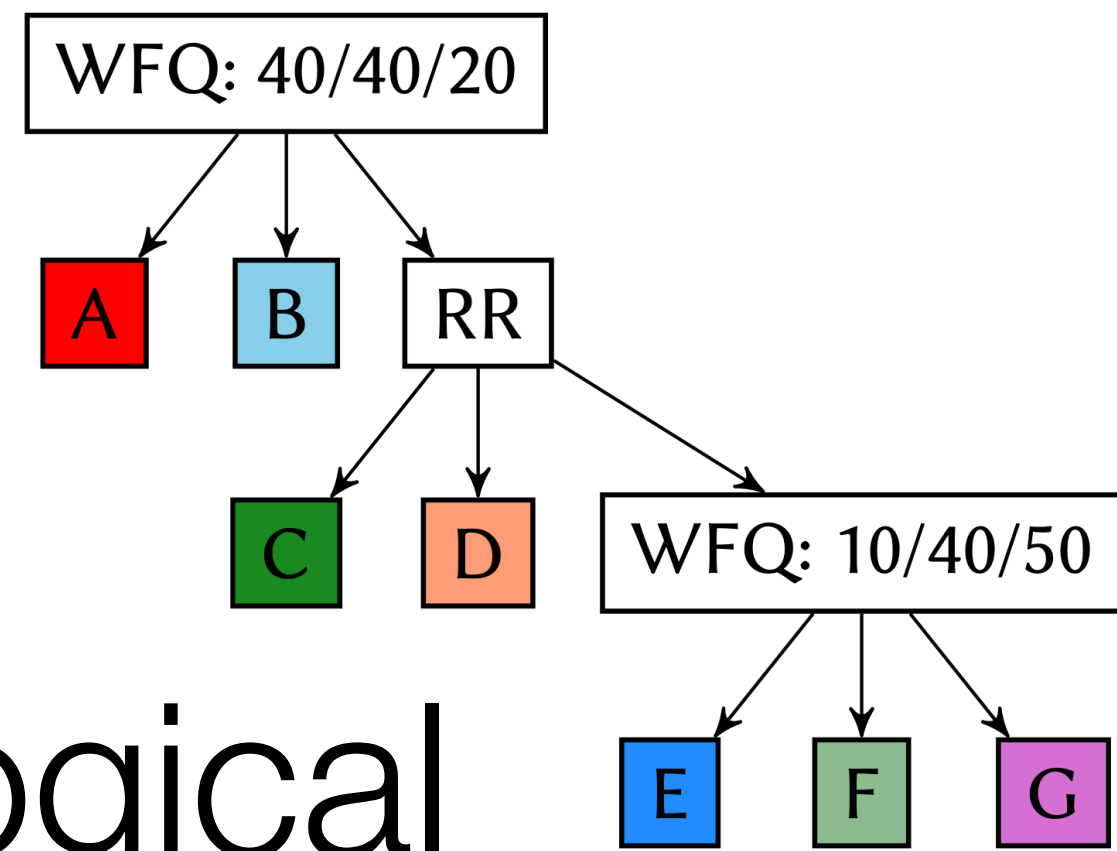


actual


Simulation




Simulation



logical



actual



Underlying formalism

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.



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.



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.



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.



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.



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.

Formal Abstractions for Packet Scheduling

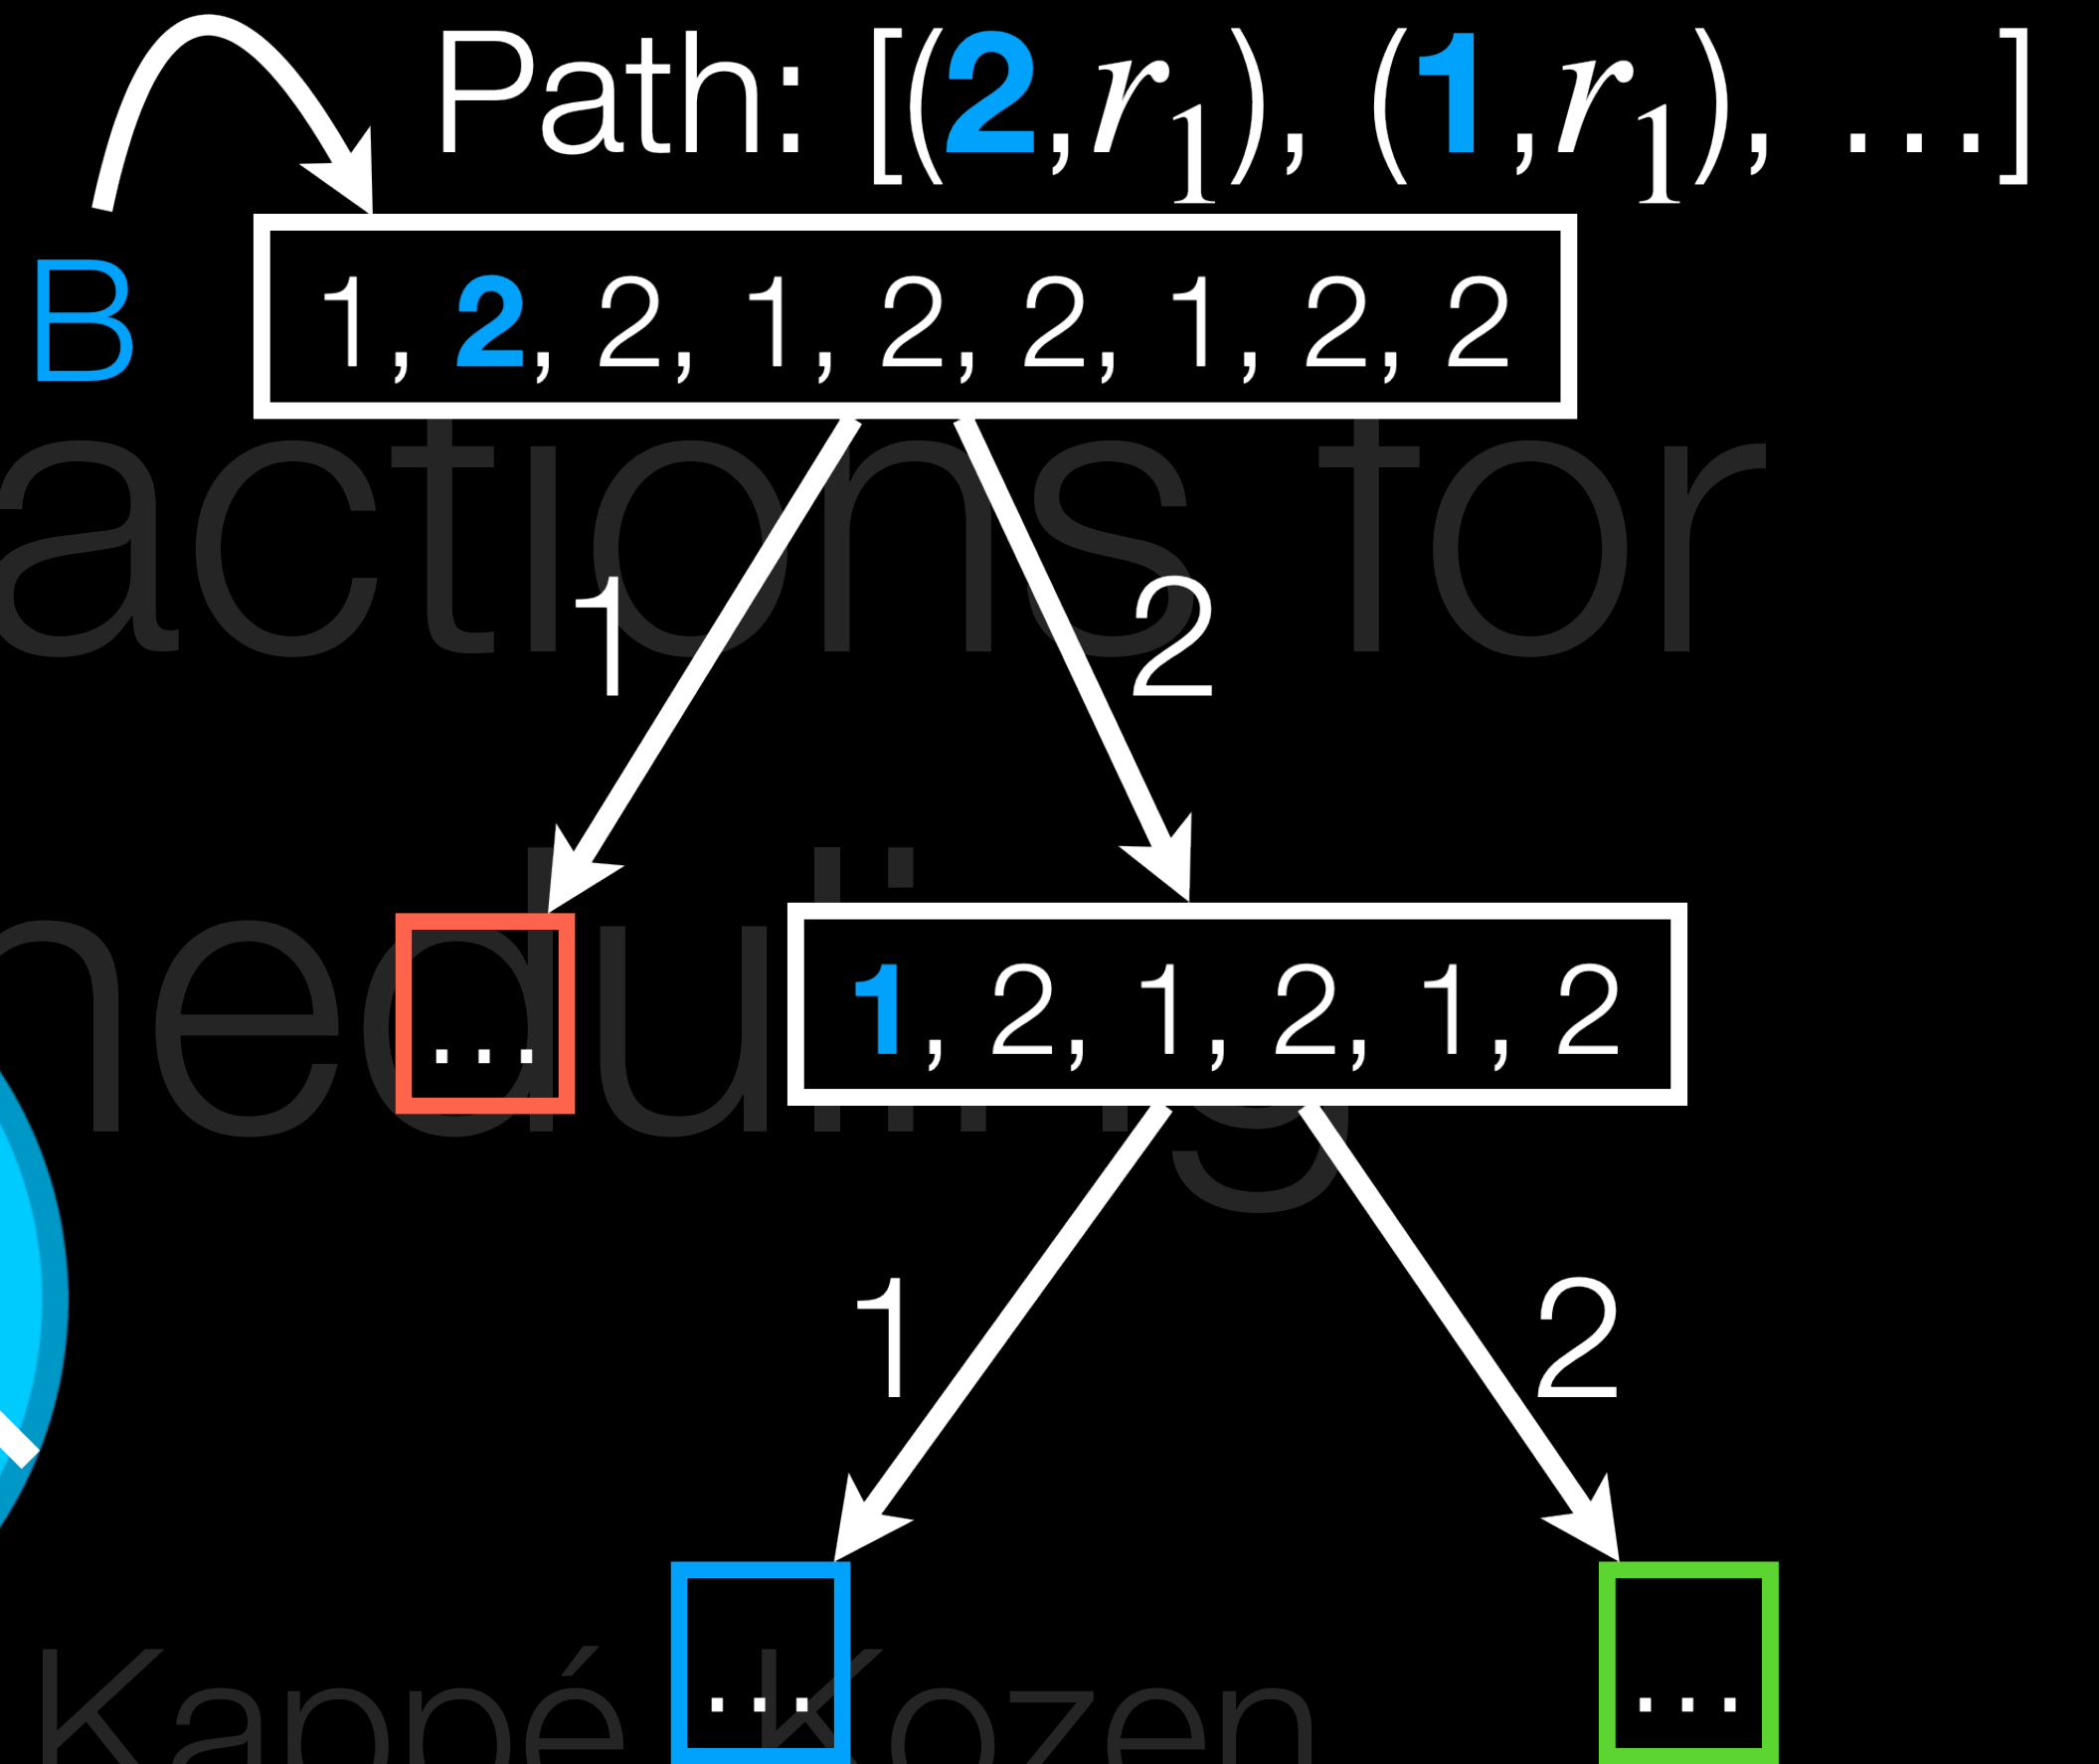
Mohan, Liu, Foster, Kappé, Kozen

cs.cornell.edu/~amohan

tree
shape



language
expressivity



Mohan, Liu, Foster, Kappé, Kozen

cs.cornell.edu/~amohan