

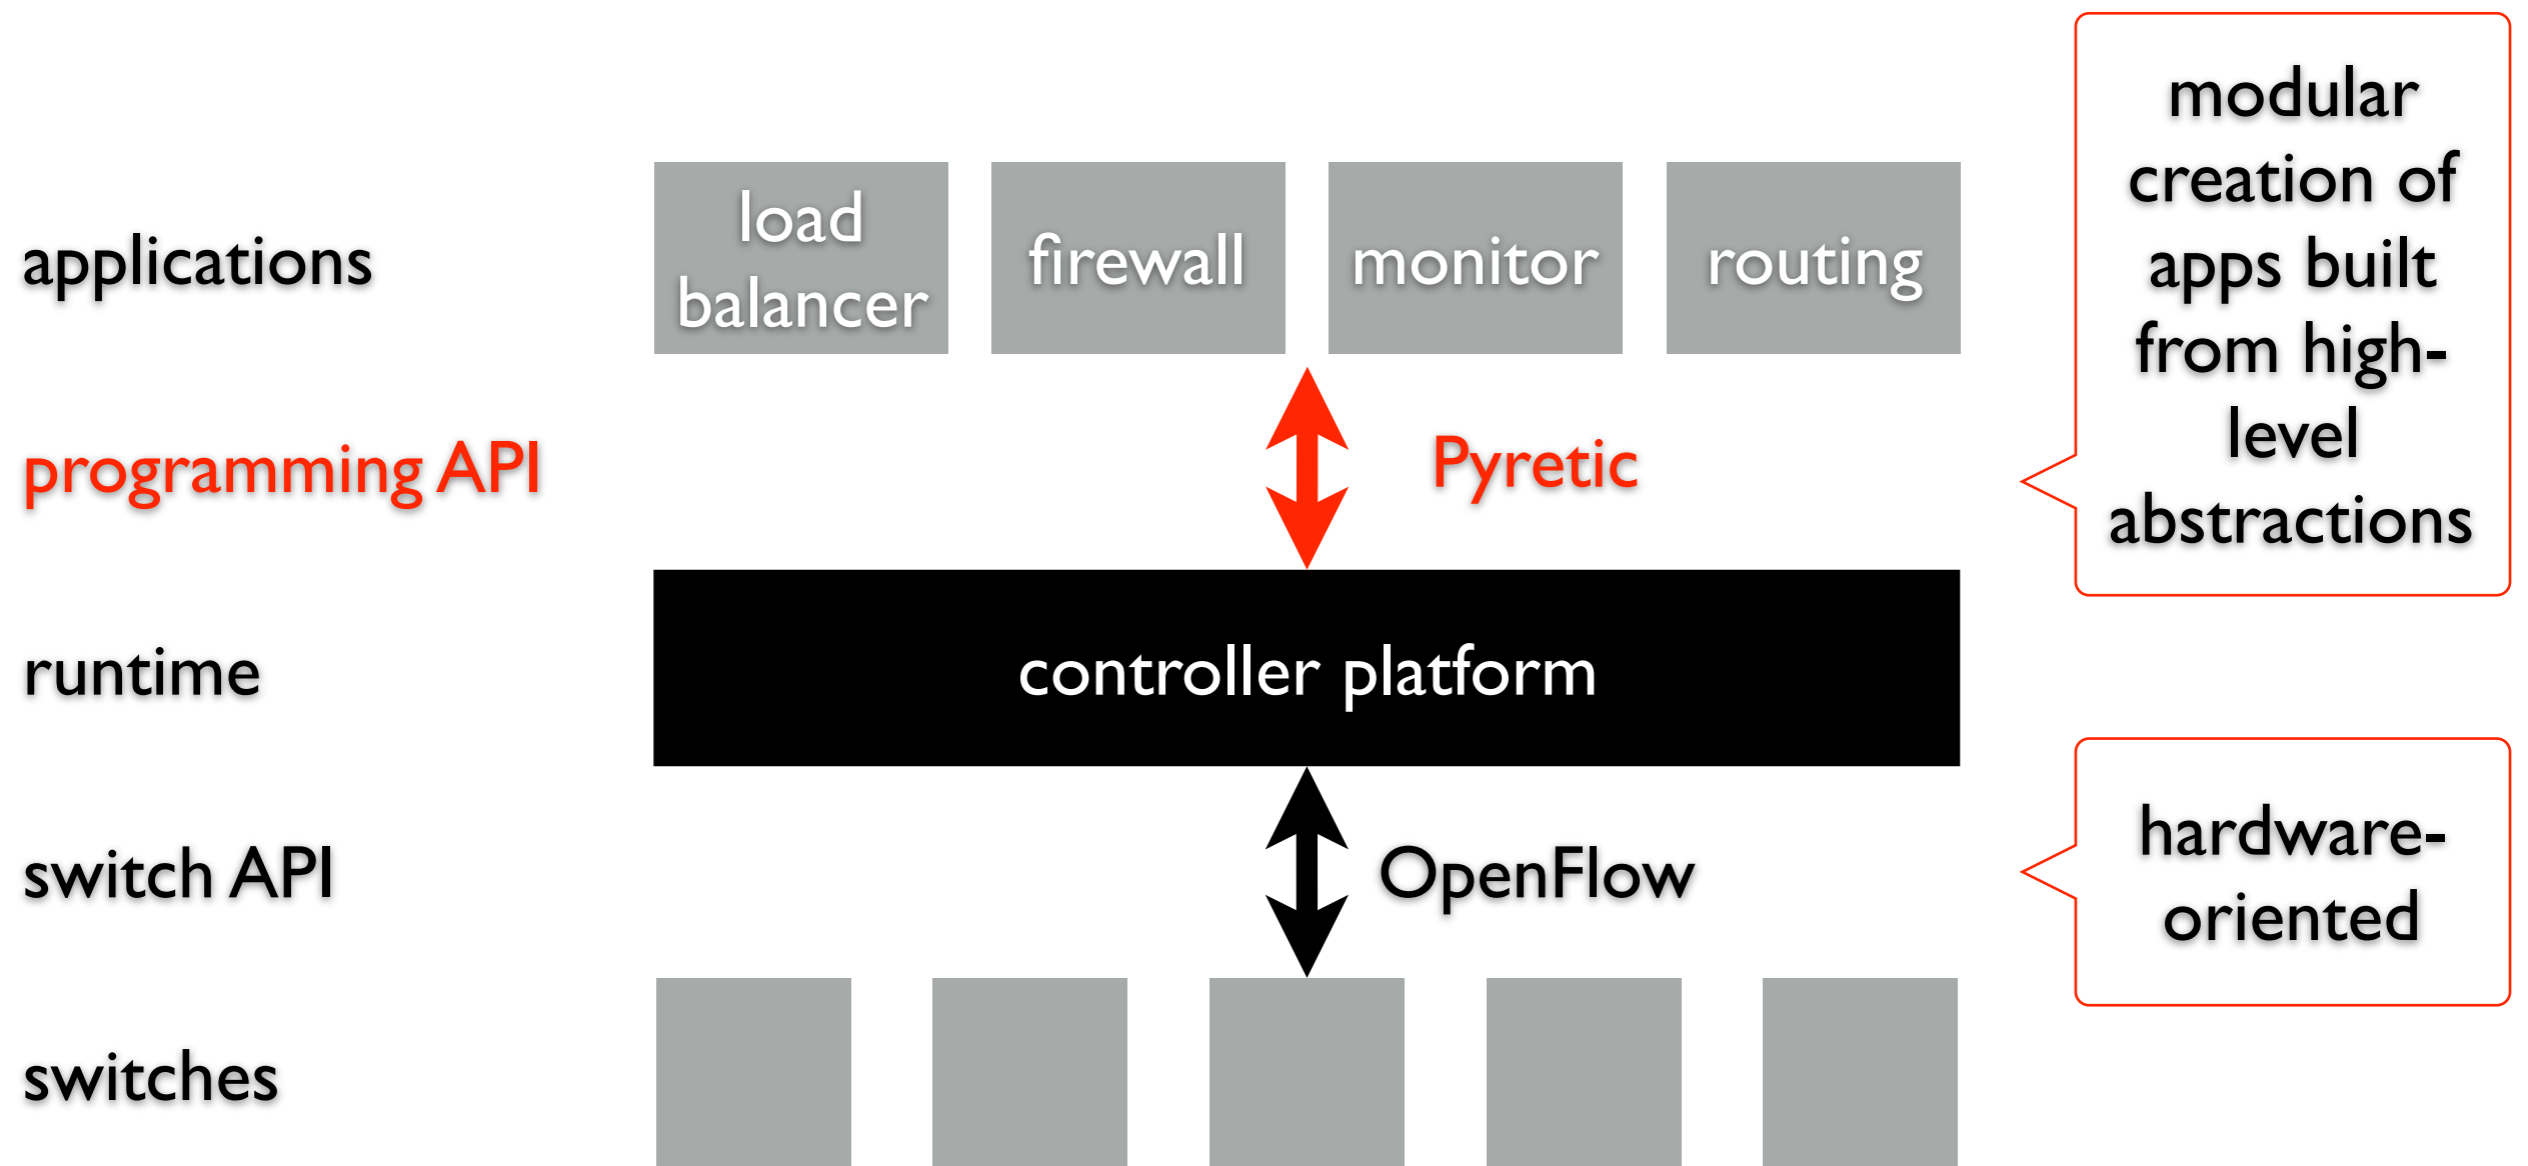
lecture 11: composing controllers

5590: software defined networking

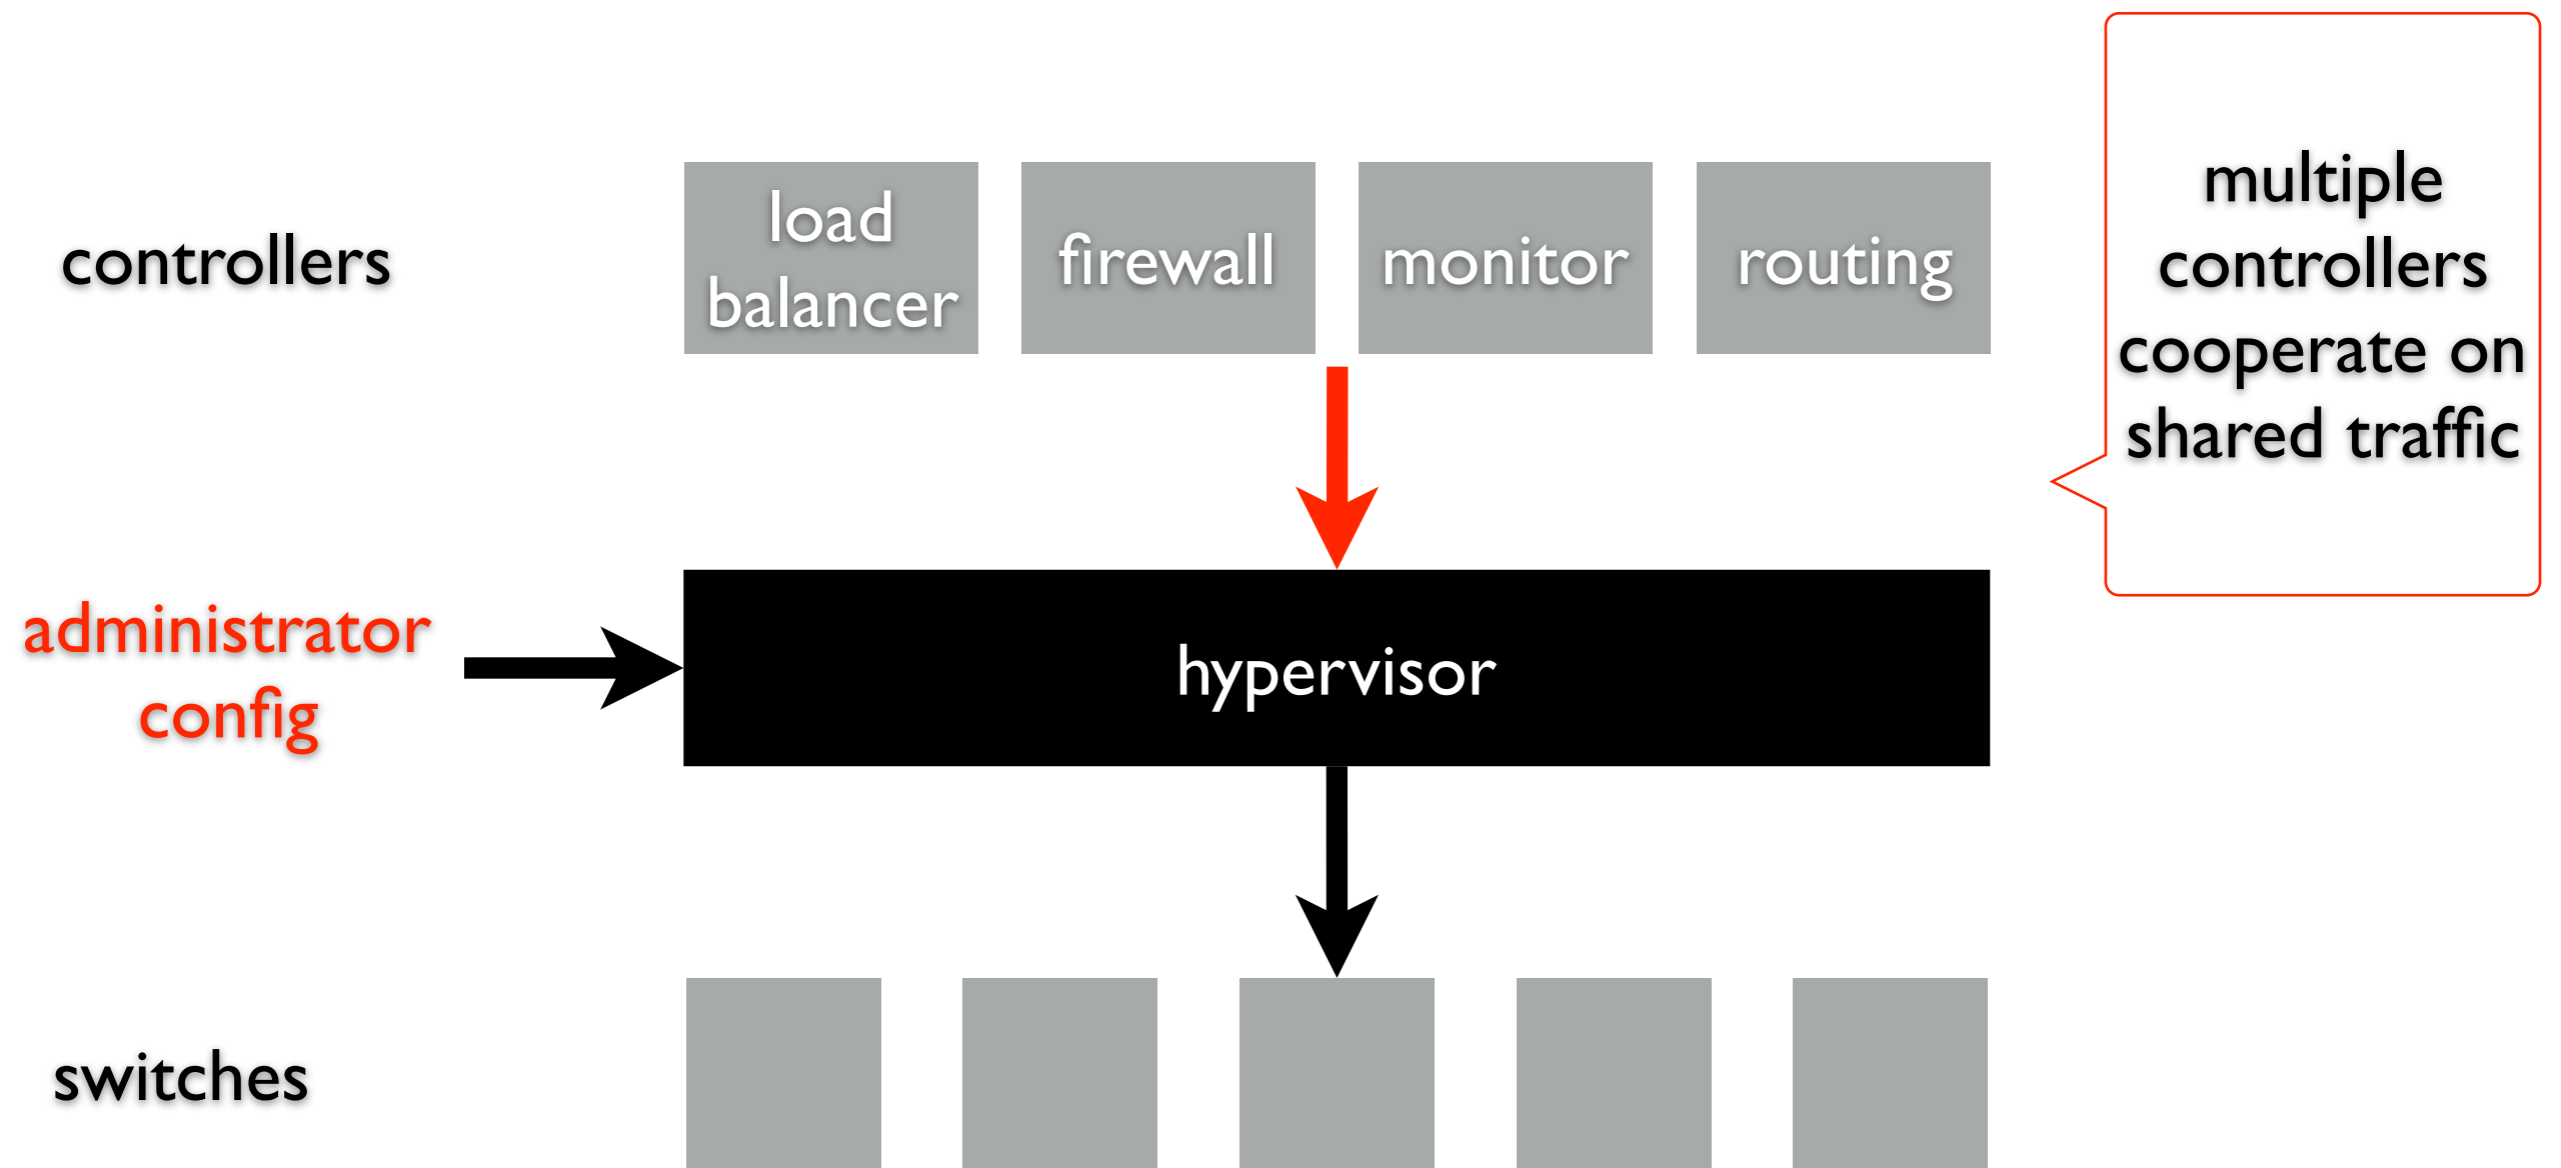
anduo wang, Temple University

TTLMAN 401B, R 17:30-20:00

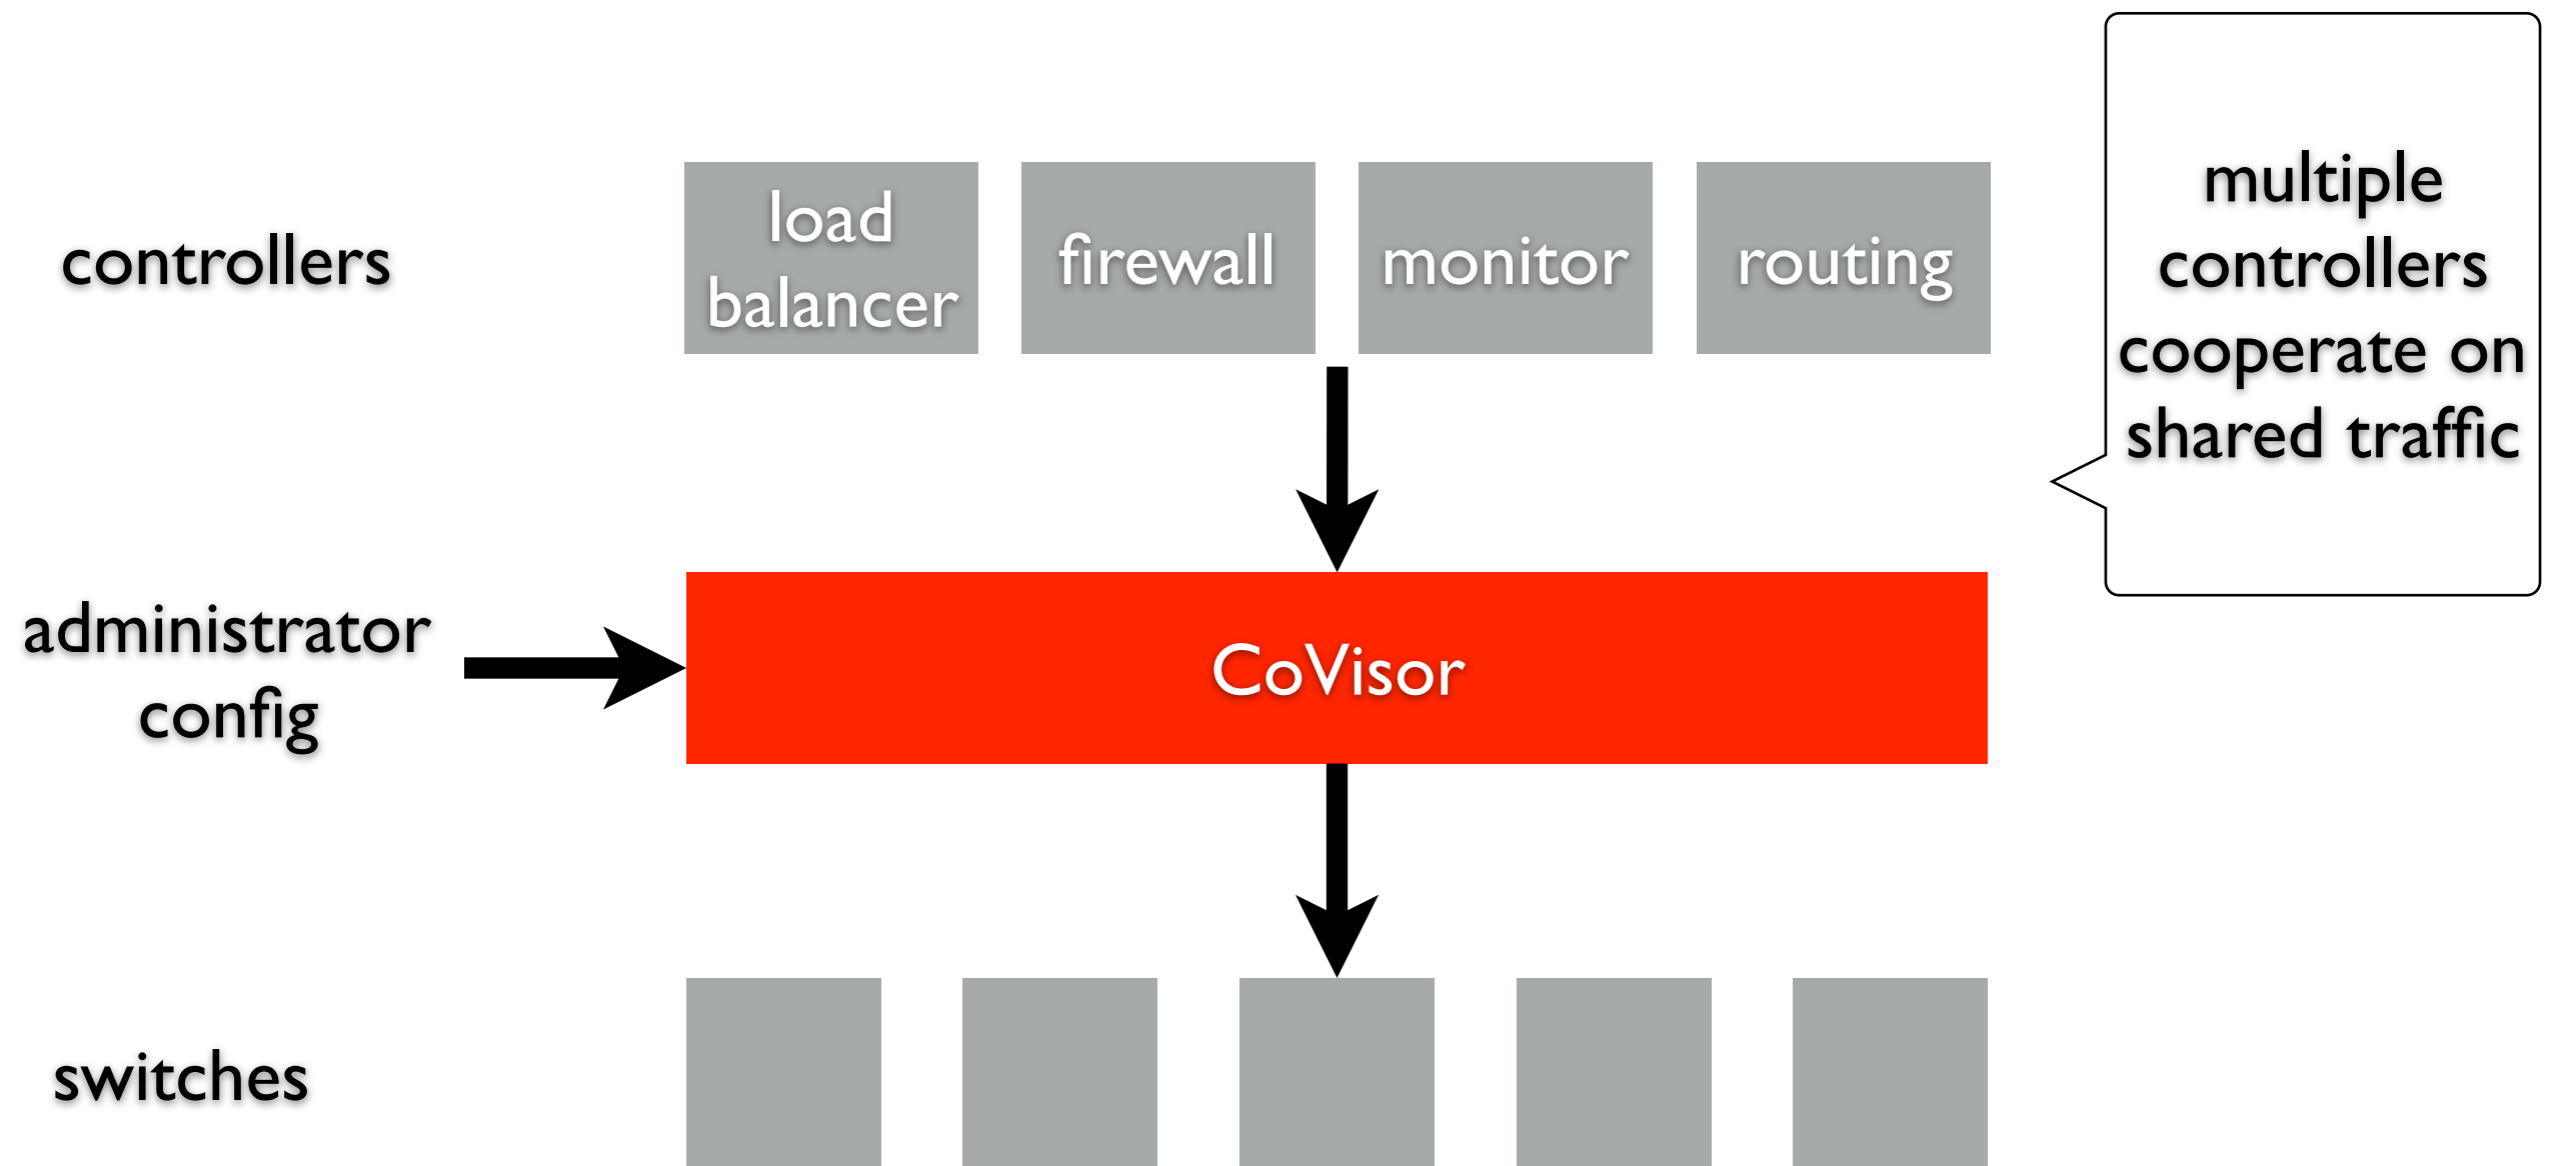
Pyretic: composing policies



composing controllers



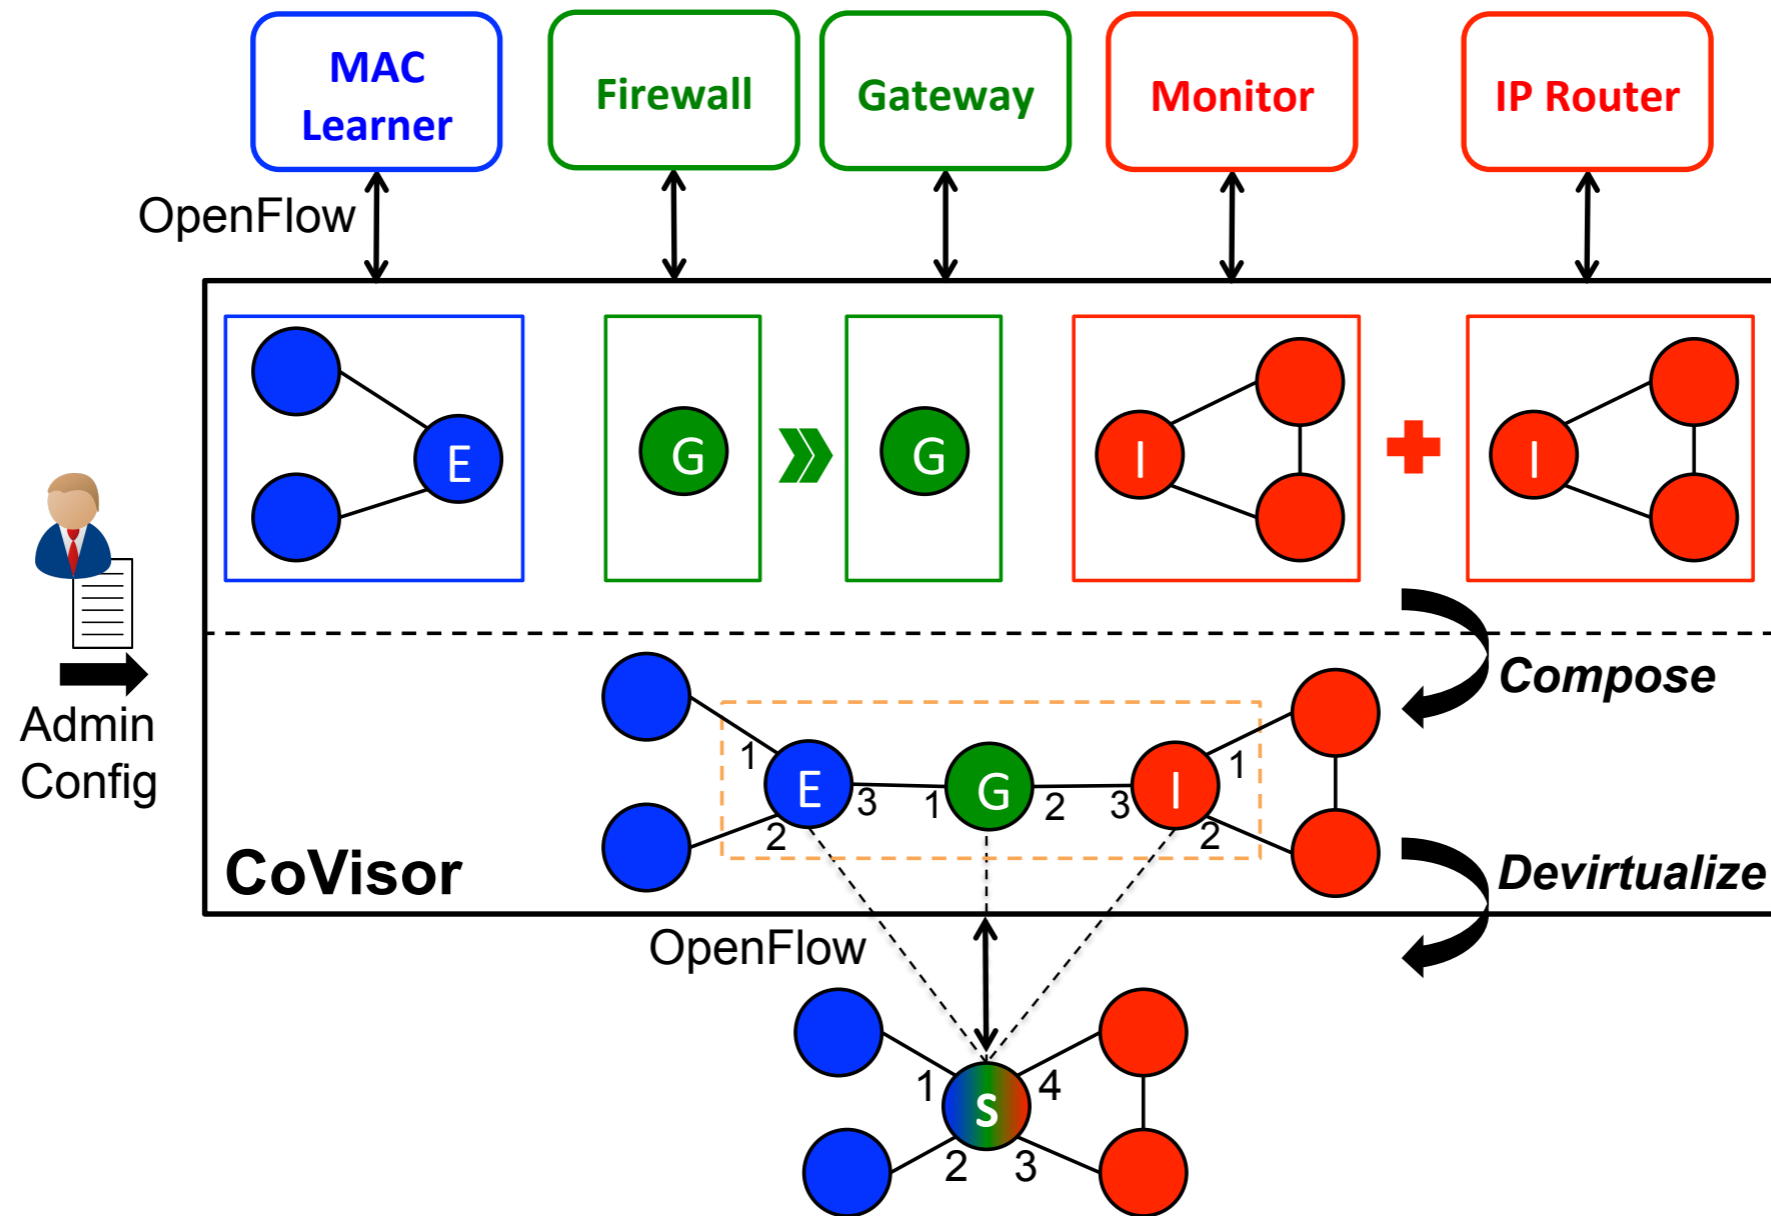
CoVisor




challenges and technical contribution

- efficient algorithms

CoVisor




CoVisor



assemble multiple controllers

- parallel, sequential, override

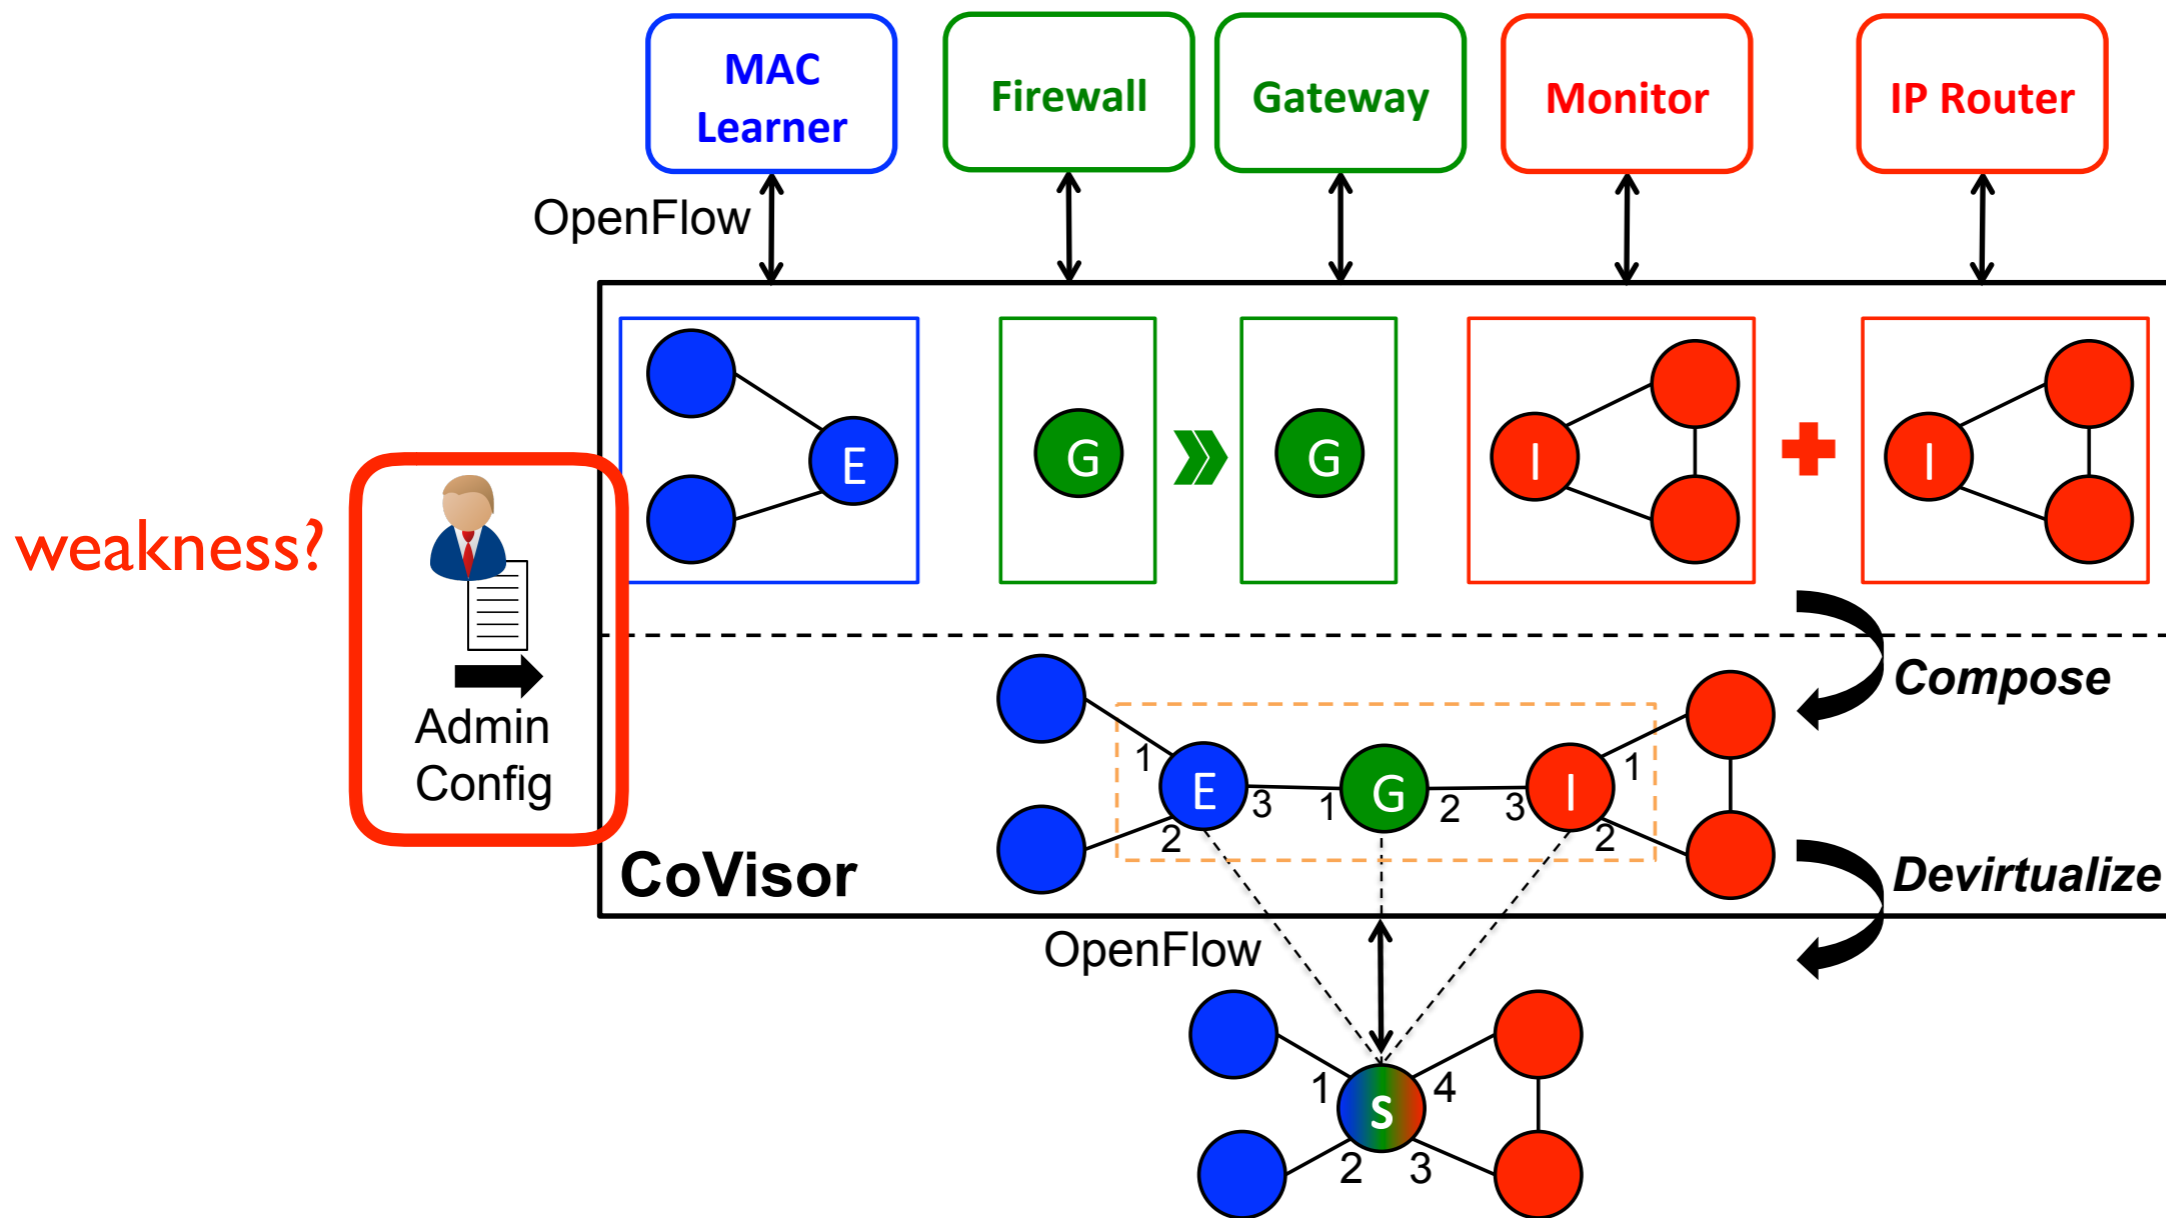
CoVisor



abstract topology

- customer virtual topology to each controller

CoVisor



protection

- fine-grained control over how a controller can operate

administrator role

administrator role

configure CoVisor to compose policies

administrator role

configure CoVisor to compose policies

- *manual spec*: $T_1 + T_2, T_1 > T_2, T_1 \triangleright T_2$

administrator role

configure CoVisor to compose policies

- *manual*

- ***proactive* incremental compilation, optimization**

administrator role

configure CoVisor to compose policies

- *manual*
- *proactive*

virtualize the network, sets packet-processing constraints

administrator role

configure CoVisor to compose policies

- *manual*
- *proactive*

virtualize the network, sets packet-processing constraints

- **virtual topo: many-to-one, one-to-many (physical-to-virtual)**

administrator role

configure CoVisor to compose policies

- *manual*
- *proactive*

virtualize the network, sets packet-processing constraints

- virtual topo: many-to-one, one-to-many (physical-to-virtual)
- **packet handling: match, action**

the “efficiency” challenge

to host tens of controllers

- each installs tens of thousands rules
- constantly updated rules

naive approach — prohibitively expensive

- time to recompile new policy
- time to install new rules on switches

efficient CoVisor algorithms

incrementally composing controller policies

- priorities form a convenient algebra, obviating recompiling from scratch


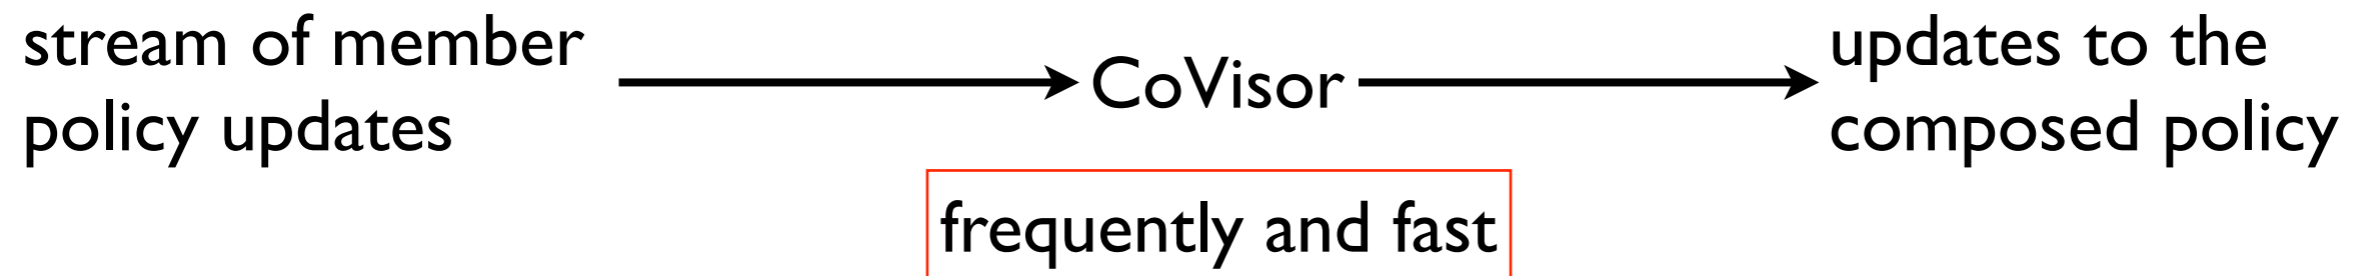
devirtualization

- one(physical)-to-many(virtual)

optimizing composition

- smart data structure accelerate compilation

incremental composition



policy composition revisit

$\text{comp}_+(R_1, R_2)$

- for every (r_1, r_2) in $(R_1 \times R_2)$
- generate new **r** if $r_1.\text{mSet}$ intersects with $r_2.\text{mSet}$
 - **r.match** = intersection of $r_1.\text{mSet}$ and $r_2.\text{mSet}$
 - **r.action** = union of $r_1.\text{action}$ and $r_2.\text{action}$

policy composition revisit

$\text{comp}_{\gg}(R_1, R_2)$

- for every (r_1, r_2) in $(R_1 \times R_2)$
- generate new r if packets produced by $r_1.\text{action}$ intersects with $r_2.\text{mSet}$
 - $r.\text{match} = ?$
 - $r.\text{action} = ?$

policy composition revisit

$\text{comp}_{\triangleright}(R_1, R_2)$

- stacking R_1 on top of R_2 with higher priority

role of priority

ideally (goal)

- single rule addition in a member policy will NOT
 - recomputing entire composed policy
 - cleaning the physical switch's flow tables

i.e., reduce update overhead

- computation
 - # of rule pairs *comp* needs to iterate
- rule update
 - # of flowmods to update a switch

strawman priority assignment

Monitoring M_R
$(1; srcip = 1.0.0.0/24; count)$
$(0; *; drop)$

Routing Q_R
$(1; dstip = 2.0.0.1; fwd(1))$
$(1; dstip = 2.0.0.2; fwd(2))$
$(0; *; drop)$

$(1; dstip=2.0.0.3; fwd(3))$

+ ↑

Parallel composition: $comp_+(M_R, Q_R)$
$(7; srcip=1.0.0.0/24, dstip=2.0.0.1; fwd(1), count)$
$(6; srcip=1.0.0.0/24, dstip=2.0.0.2; fwd(2), count)$
$(5; srcip=1.0.0.0/24, dstip=2.0.0.3; fwd(3), count)$
$(4; srcip=1.0.0.0/24; count)$
$(3; dstip=2.0.0.1; fwd(1))$
$(2; dstip=2.0.0.2; fwd(2))$
$(1; dstip=2.0.0.3; fwd(3))$
$(0; *; drop)$

position of the rule indicates
relative priority

strawman priority assignment

Monitoring M_R
$(1; srcip = 1.0.0.0/24; count)$
$(0; *; drop)$

Routing Q_R
$(1; dstip = 2.0.0.1; fwd(1))$
$(1; dstip = 2.0.0.2; fwd(2))$
$(0; *; drop)$

$(1; dstip=2.0.0.3; fwd(3))$
+ ↑

Parallel composition: $comp_+(M_R, Q_R)$
$(7; srcip=1.0.0.0/24, dstip=2.0.0.1; fwd(1), count)$
$(6; srcip=1.0.0.0/24, dstip=2.0.0.2; fwd(2), count)$
$(5; srcip=1.0.0.0/24, dstip=2.0.0.3; fwd(3), count)$
$(4; srcip=1.0.0.0/24; count)$
$(3; dstip=2.0.0.1; fwd(1))$
$(2; dstip=2.0.0.2; fwd(2))$
$(1; dstip=2.0.0.3; fwd(3))$
$(0; *; drop)$

rules in bold count toward rule update overhead

strawman priority assignment

Monitoring M_R
$(1; \text{srcip} = 1.0.0.0/24; \text{count})$
$(0; *, \text{drop})$

Routing Q_R
$(1; \text{dstip} = 2.0.0.1; \text{fwd}(1))$
$(1; \text{dstip} = 2.0.0.2; \text{fwd}(2))$
$(0; *, \text{drop})$

$(1; \text{dstip}=2.0.0.3; \text{fwd}(3))$

↑
+

Parallel composition: $comp_+(M_R, Q_R)$
$(7; \text{srcip}=1.0.0.0/24, \text{dstip}=2.0.0.1; \text{fwd}(1), \text{count})$
$(6; \text{srcip}=1.0.0.0/24, \text{dstip}=2.0.0.2; \text{fwd}(2), \text{count})$
$(5; \text{srcip}=1.0.0.0/24, \text{dstip}=2.0.0.3; \text{fwd}(3), \text{count})$
$(4; \text{srcip}=1.0.0.0/24; \text{count})$
$(3; \text{dstip}=2.0.0.1; \text{fwd}(1))$
$(2; \text{dstip}=2.0.0.2; \text{fwd}(2))$
$(1; \text{dstip}=2.0.0.3; \text{fwd}(3))$
$(0; *, \text{drop})$

rules in bold count toward rule update overhead

smartly set priority

- to make updates incremental

incremental update and priority algebra

r is computed from r_1 and r_2

- $r.\text{priority} \leftarrow r_1.\text{priority}, r_2.\text{priority}$
- incremental update without modifying existing priorities

incremental update and priority algebra

r is computed from r_1 and r_2

- $r.\text{priority} \leftarrow r_1.\text{priority}, r_2.\text{priority}$

comp_+

- $r.\text{priority} = r_1.\text{priority} + r_2.\text{priority}$

$\text{comp}_{<<}$

- $r.\text{priority} = r_1.\text{priority} \times \text{MAX}_2 + r_2.\text{priority}$

incremental update and priority algebra

r is computed ($\text{comp}_{\triangleright}$) from R_1 and R_2

- $r.\text{priority} = r.\text{priority} + \text{MAX}_2$ if r in R_1
- $r.\text{priority} = r.\text{priority}$ if r in R_2

algebra properties


identify and prove properties

- the assignment schema ensures newly generated priority
 - leaves existing priority unchanged
 - together, the new and existing priorities are compliant with the straw man scheme

devirtualization

topology transformation for one-to-many

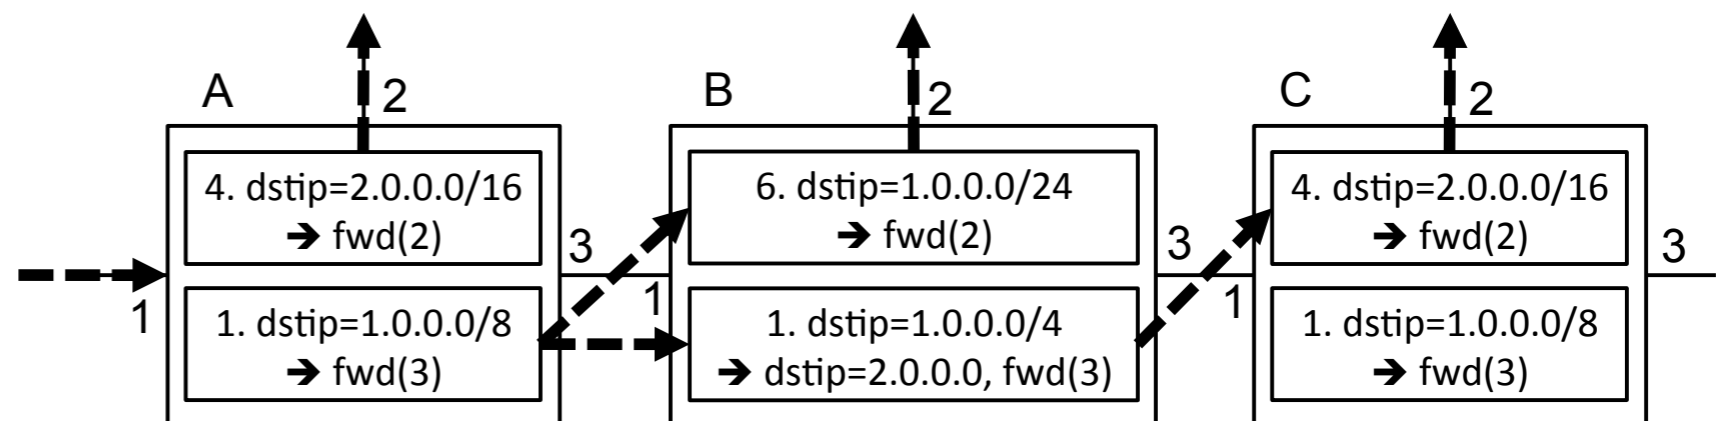

- generate symbolic path (from the virtual ingress to egress)
- on each virtual path, sequentially compose virtual policies to into a single (physical) rule



devirtualization

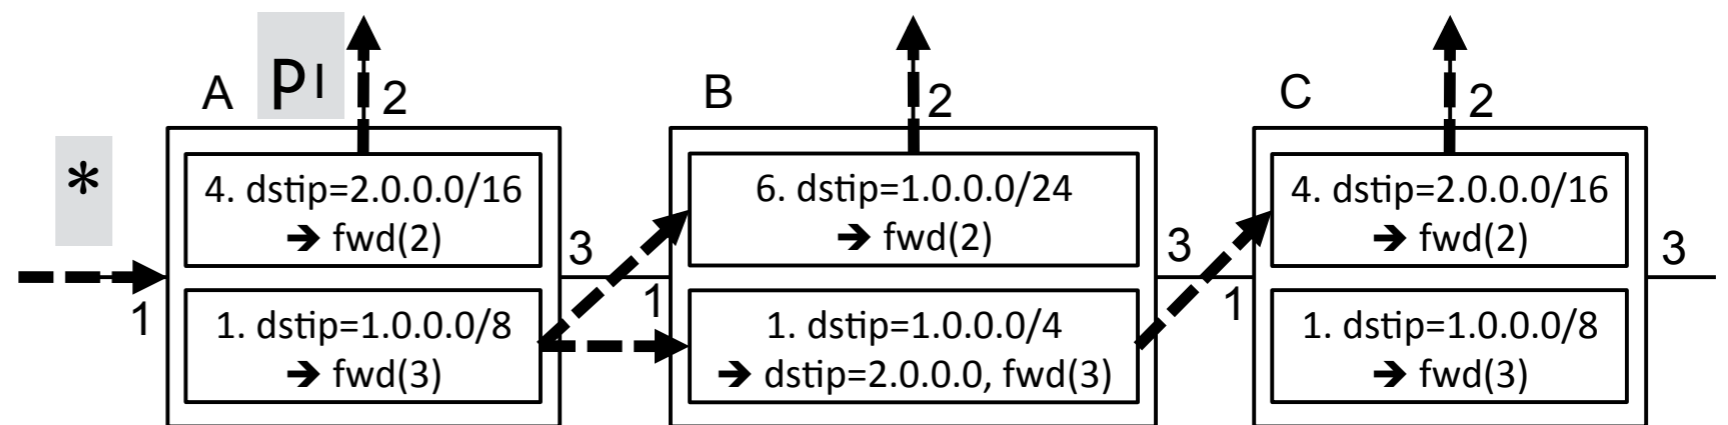


on the virtual topology, find symbolic paths

- inject wildcard packet * at ingress
- at each hop
 - evaluate the virtual policy, resulting in new packets
- until all packets reach egress



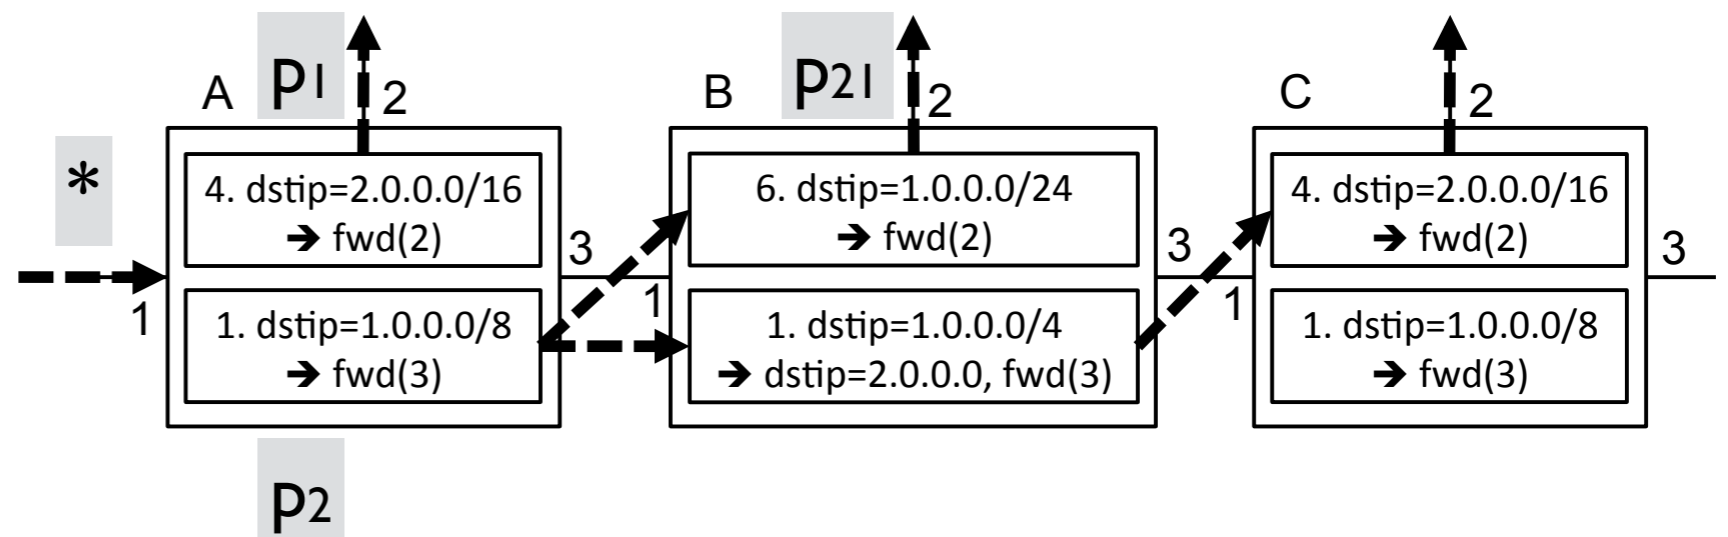


devirtualization

on the virtual topology, find symbolic paths



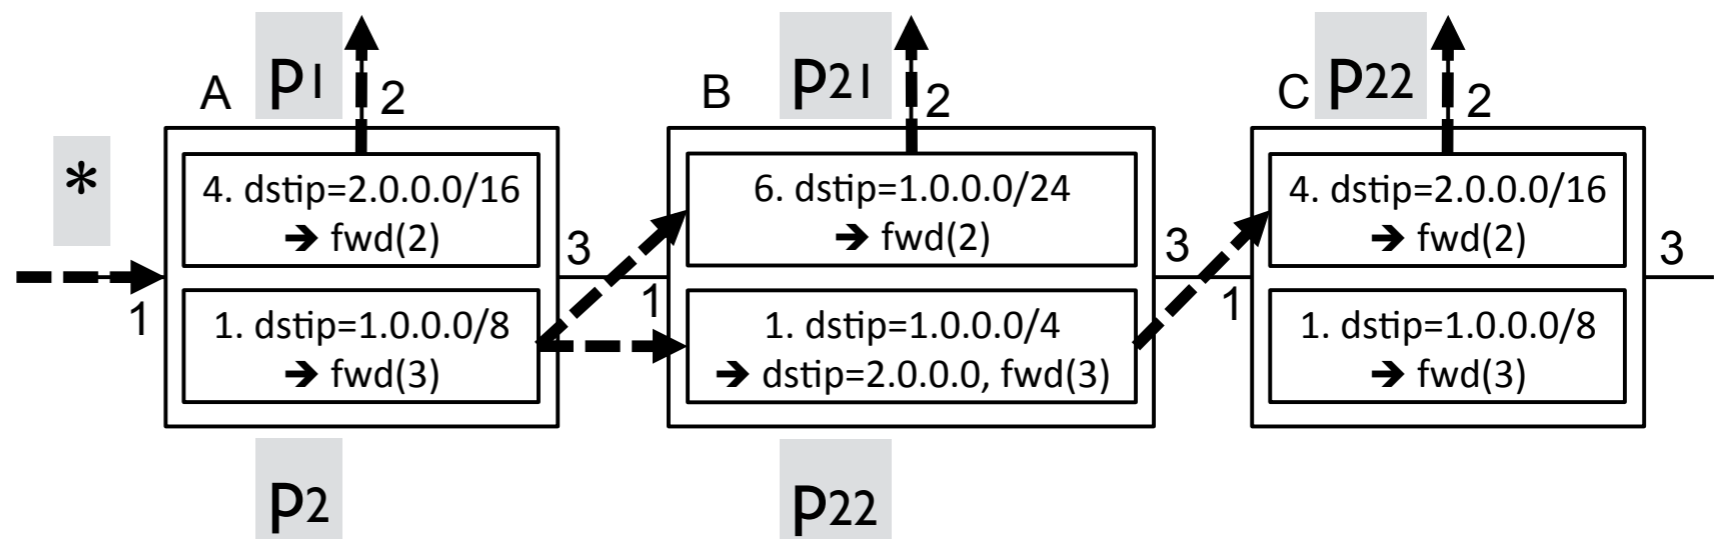


devirtualization

on the virtual topology, find symbolic paths



devirtualization

on the virtual topology, find symbolic paths

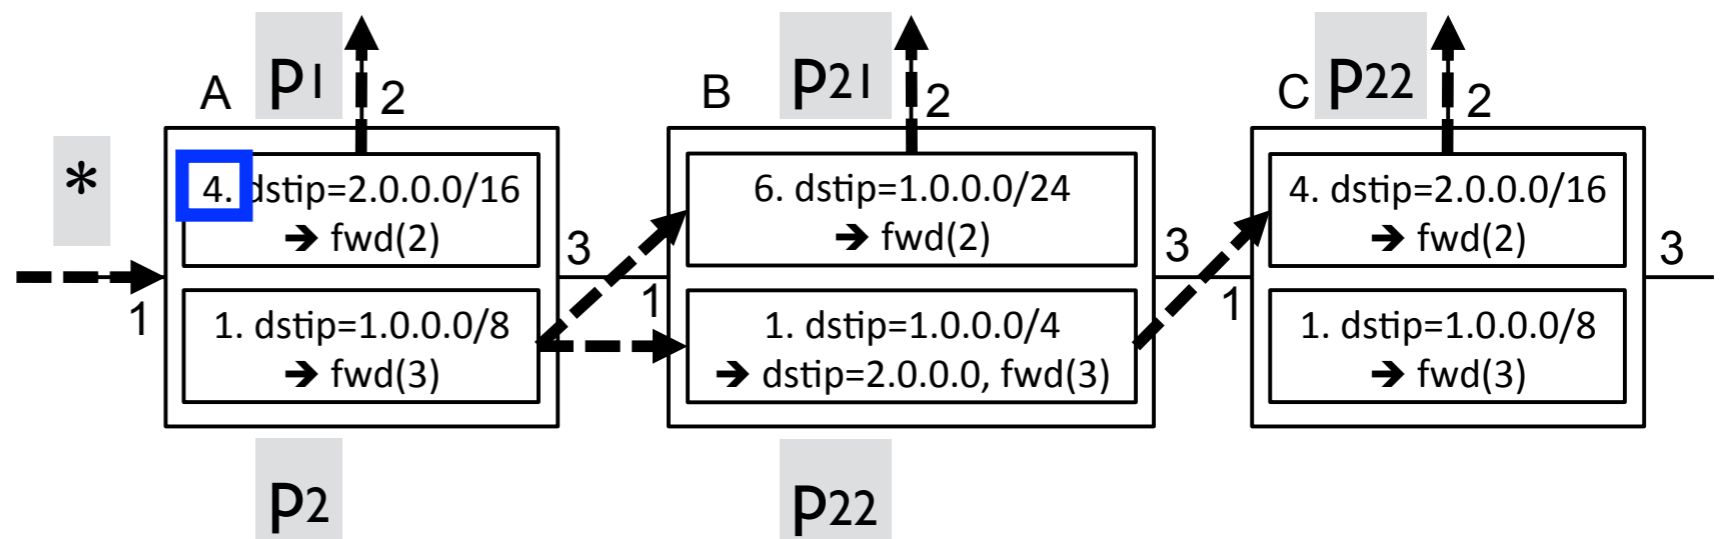




devirtualization

sequentially compose policies on each path

priority

4



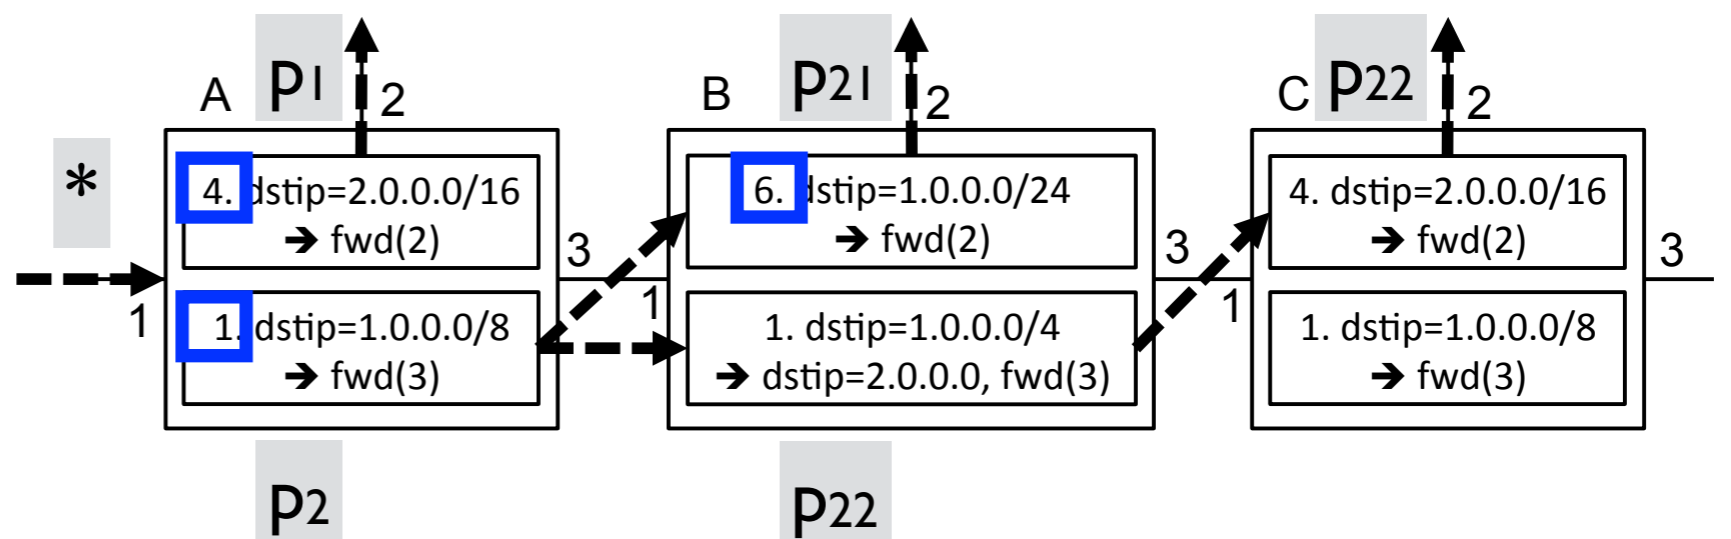




devirtualization

sequentially compose policies on each path

priority

4

1 ◦ 6 (=14)



devirtualization

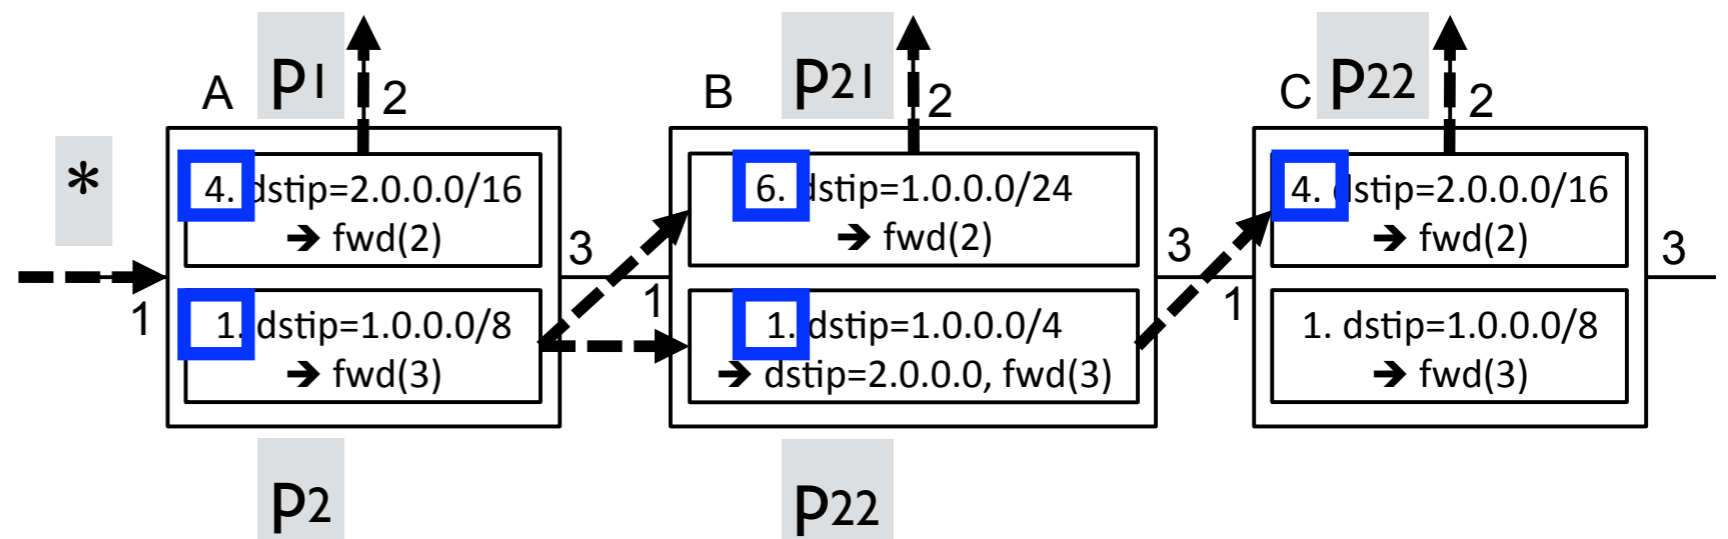



sequentially compose policies on each path

priority

4

1 ◦ 6 (=14)

1 ◦ 1 ◦ 4 (=76)



devirtualization

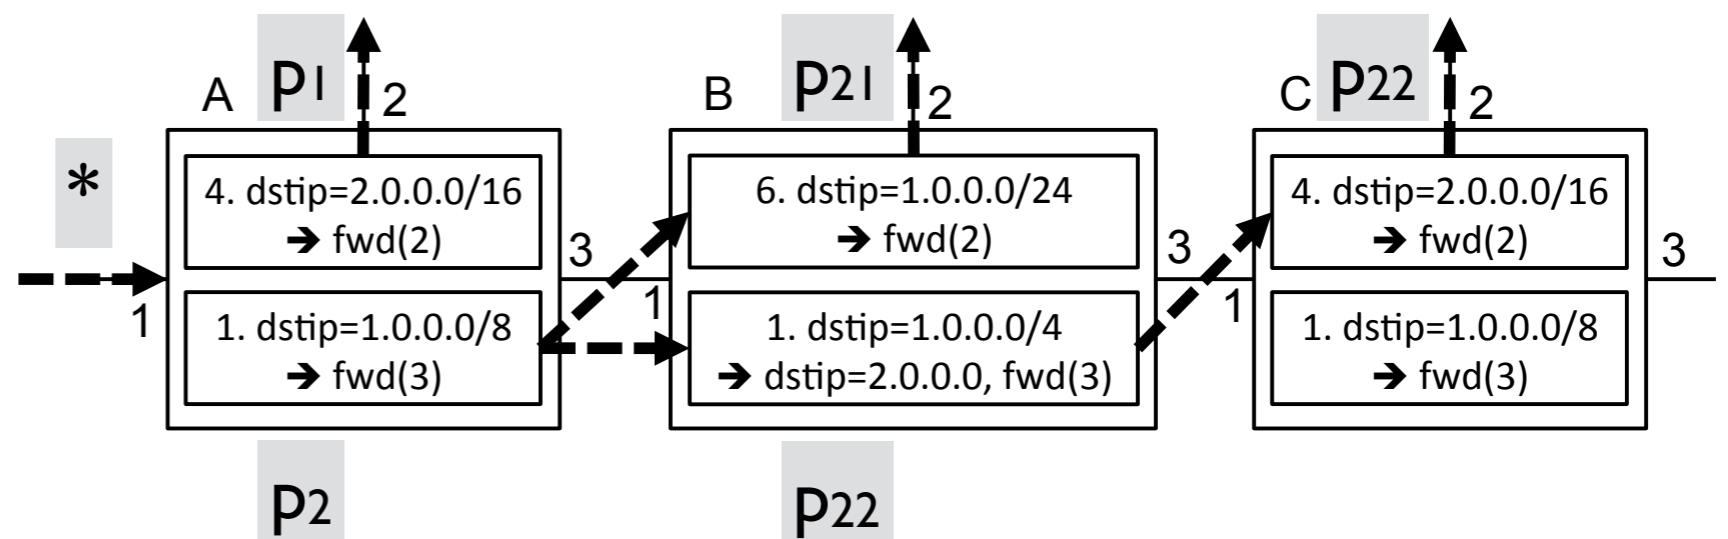




sequentially compose policies on each path

priority

4 0 0 0 (=256)


1 0 6 0 (=112)

1 0 1 4 (=76)



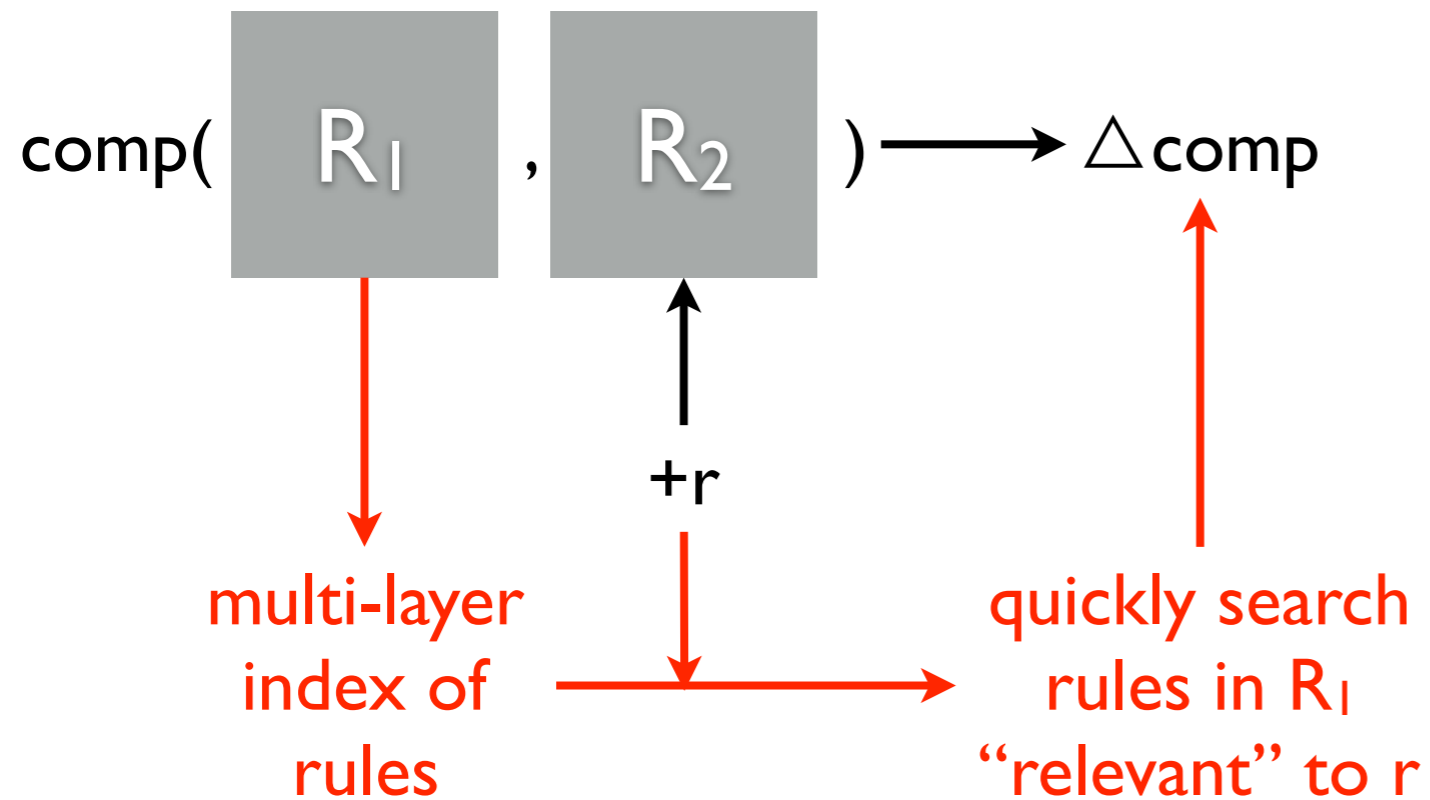
optimization: indexing rules

accelerate compilation with smart data structure




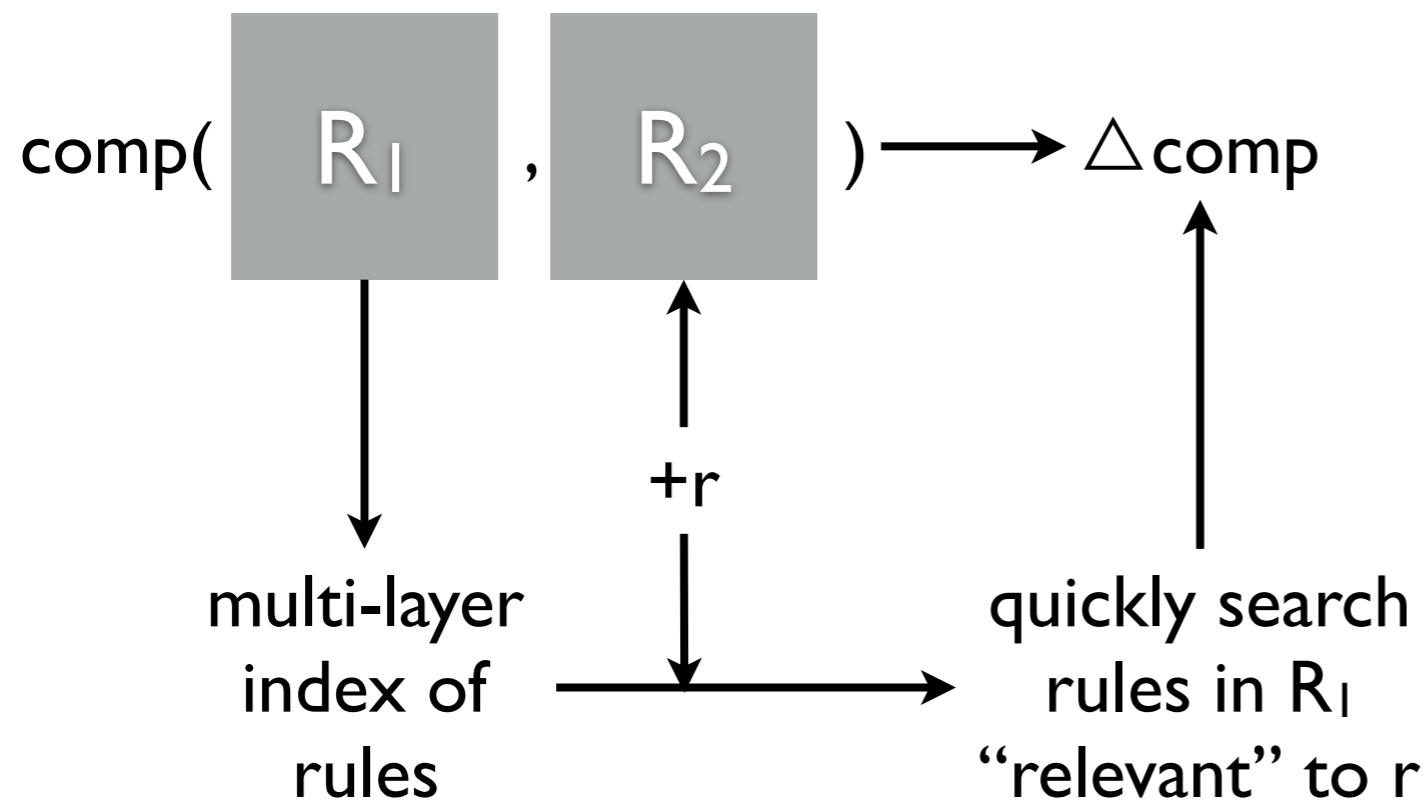
optimization: indexing rules

accelerate compilation with smart data structure



optimization: indexing rules

accelerate compilation with smart data structure

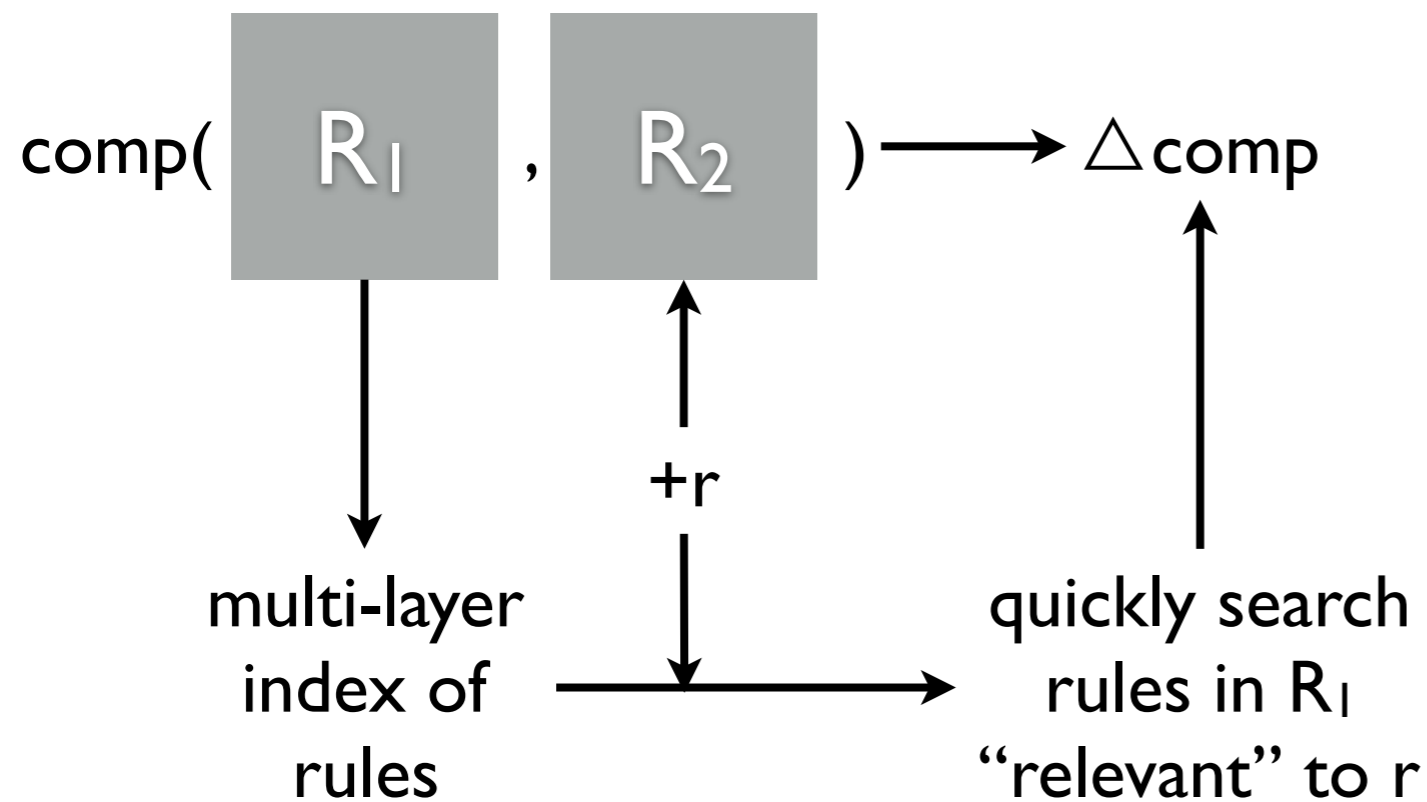


example multi-layer index

- hash table for exact match
- trie for prefix match
- list for arbitrary wildcard-match

optimization: indexing rules

reduce index size by policy correlation



only index the "correlated" info
 $R_1.index = R_2.index$
 $= R_1.fields \cap R_2.fields$