

lecture 02: review of “how the Internet works”

5590: software defined networking

anduo wang, Temple University

TTLMAN 402, R 17:30-20:00


some materials in this slide are based on lectures by
Jennifer Rexford <https://www.cs.princeton.edu/courses/archive/fall13/cos597E/>



THE GOOD



THE BAD



THE UGLY

why review

SDN interacts with “legacy” networks

- unmodified end-host computers
- hybrid deployment of SDN
- connecting to non-SDN domains

SDN is a reaction to legacy networks

- retain the “good”
- improve on the “bad” and the “ugly”

outline

brief review

- defining characteristics

“the good, the bad and the ugly” by examples

- traffic engineering in IP networks
- Ethernet
- VLAN usage in campus networks

defining characteristics

- packet switching
- layering

packet switching

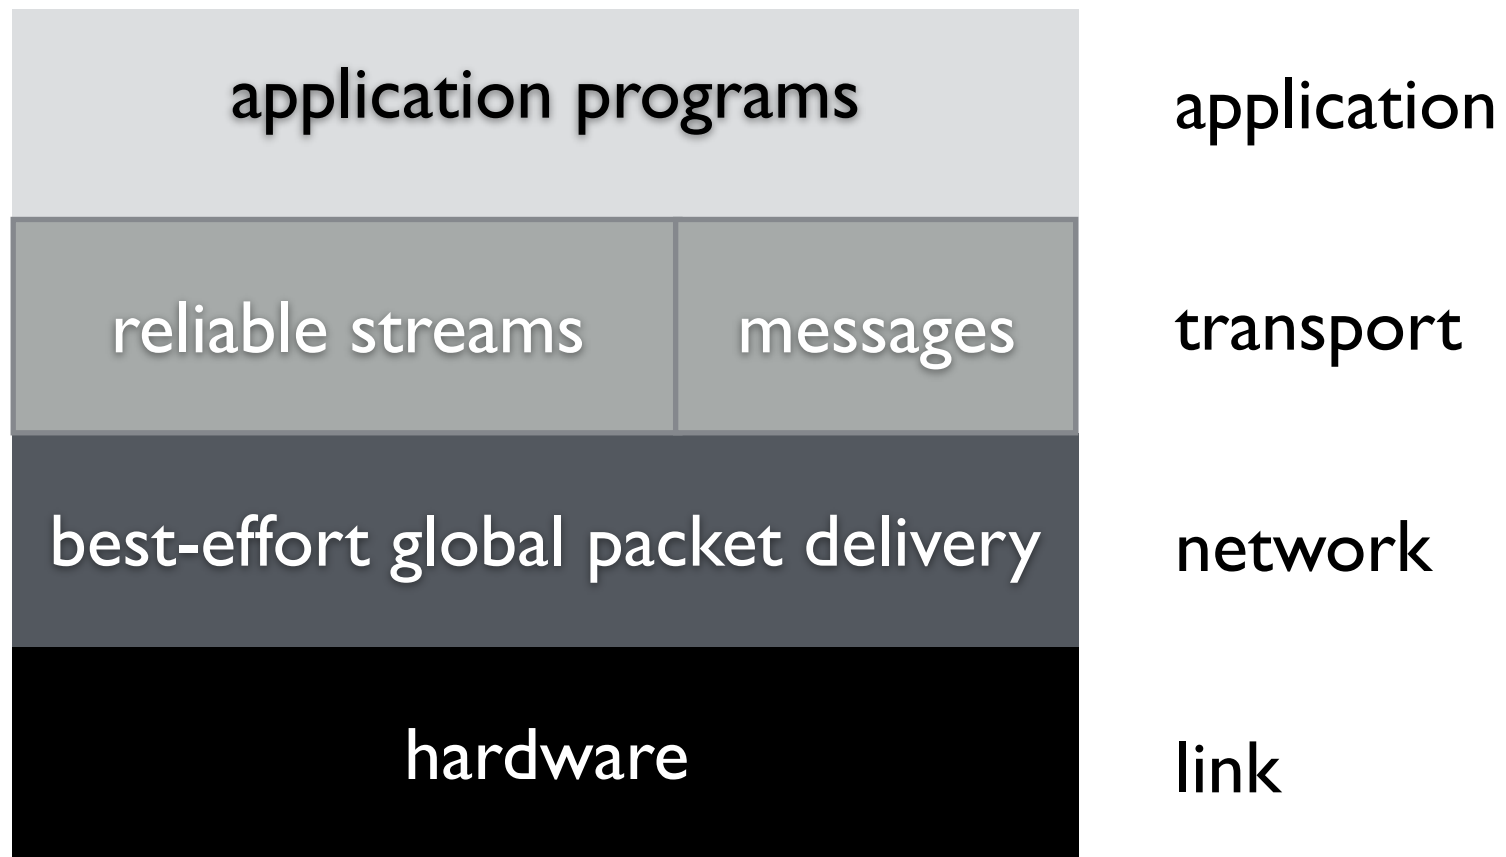
the simple and transparent core network

- the datagram, connectionless service
 - carries data without knowing what data it is
- effective for multiplexed utilization of shared interconnected networks
- open to new applications, hardwares, and new protocols


intelligence at the edges

- end hosts can run arbitrary applications



protocol layering for modularity



protocol layering for modularity





protocol layering for modularity




protocol layering for modularity

the **hourglass**



packet switching + layering



tension

- high-level network-wide objectives understood by the edges
- low-level network management of the core

the Internet is increasingly complex and notoriously hard to operate

outline

brief review

- significant ideas

“the good, the bad and the ugly” by examples

- traffic engineering in IP networks
- Ethernet
- VLAN usage in campus networks

traffic engineering
with
traditional IP routing protocols

traffic engineering

IP network manages itself

- end hosts running TCP adapt their sending rates to network congestion
- but, a particular link might be congested despite the presence of under-utilized links


TCP/IP does not adapt the routing of traffic to the prevailing demand

- *a network-wide objective: improving user performance and making more efficient use of network resources*
- this task: traffic engineering

intradomain routing

shortest path routing

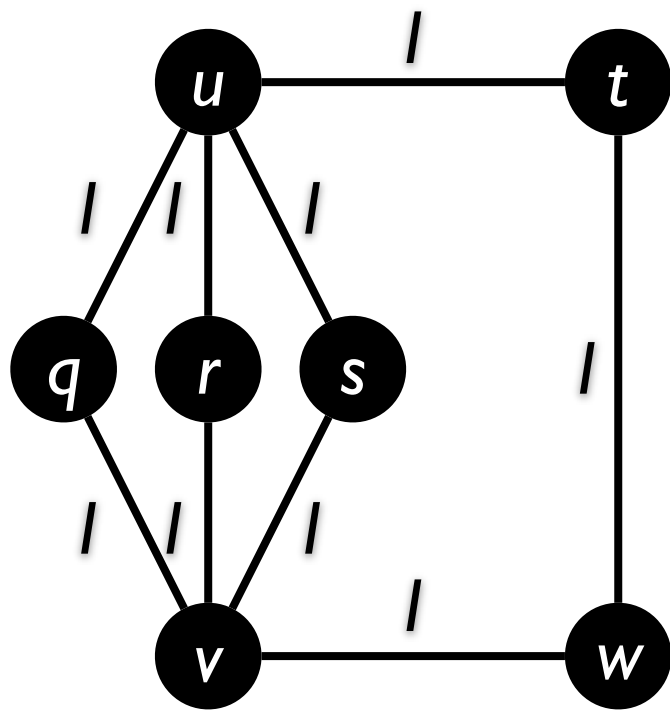
- route traffic through the shortest path within an Autonomous system based on OSPF weights



intradomain traffic engineering

routing the same demand with differing weights


- demand: q, r, s, w each has one unit of traffic to send to t




intradomain traffic engineering

routing the same demand with differing weights


- demand: q, r, s, w each has one unit of traffic to send to t



initial unit weights



local change of the congested link




global optimization of link weights

intradomain traffic engineering

globally optimized link weights

- alleviate congestion
- attractive alternative to buying additional bandwidth



traffic engineering framework

routing model

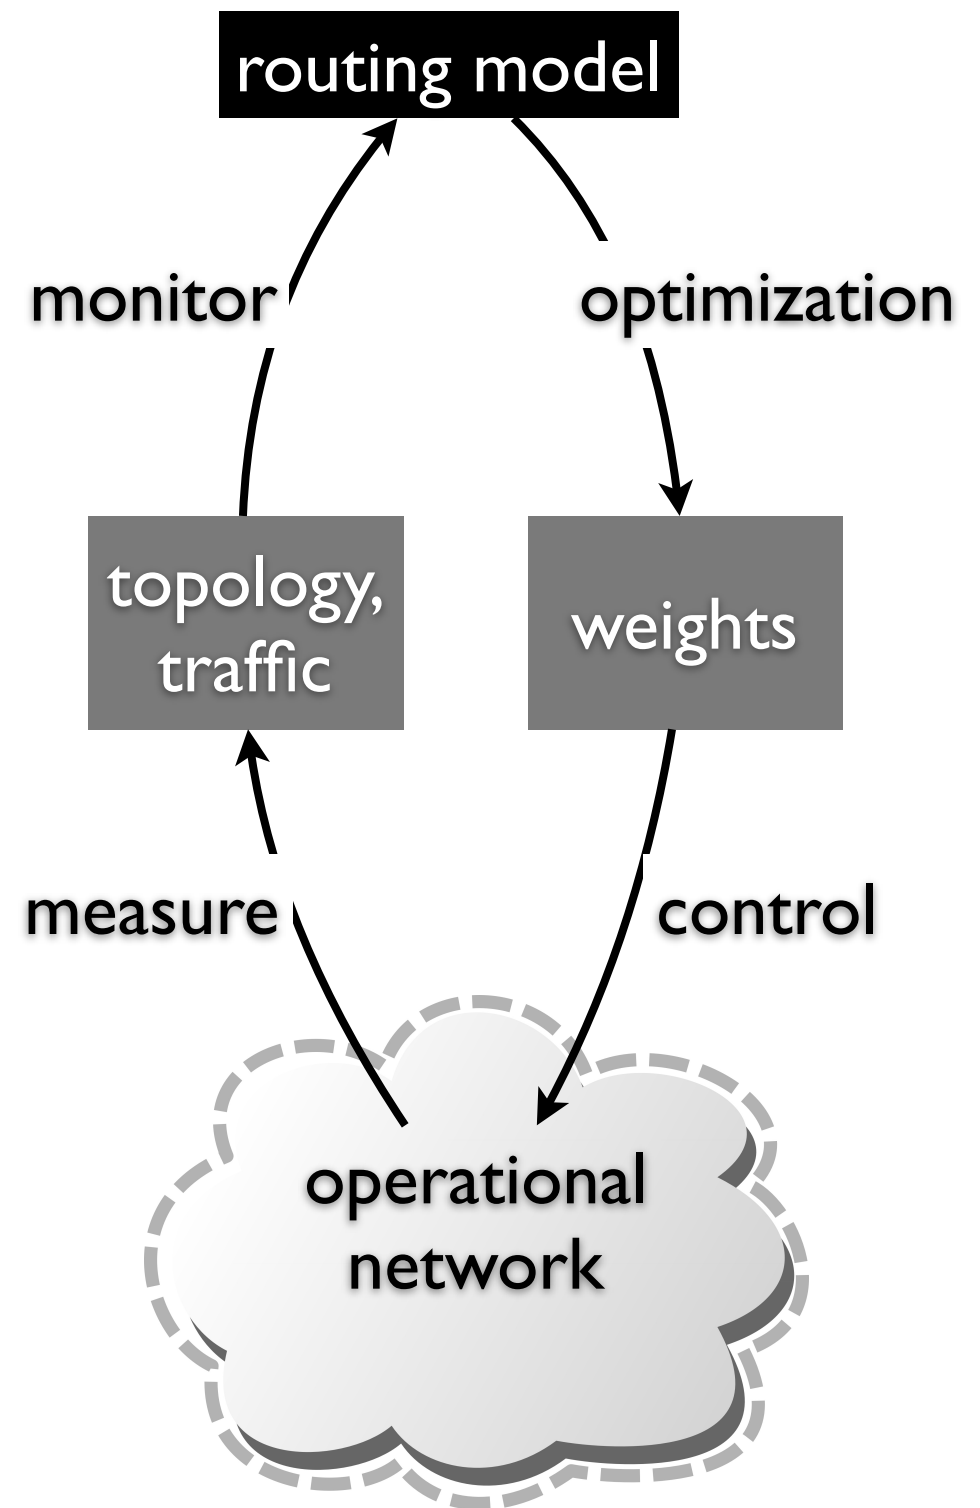
- path selection (shortest path) based on IGP weights

measurement

- lively and accurate view of the network — topology, traffic demand

reconfiguring weights

- optimize a network-wide objective
- e.g., minimize the max-utilization
- e.g., keep max-utilization under 60%



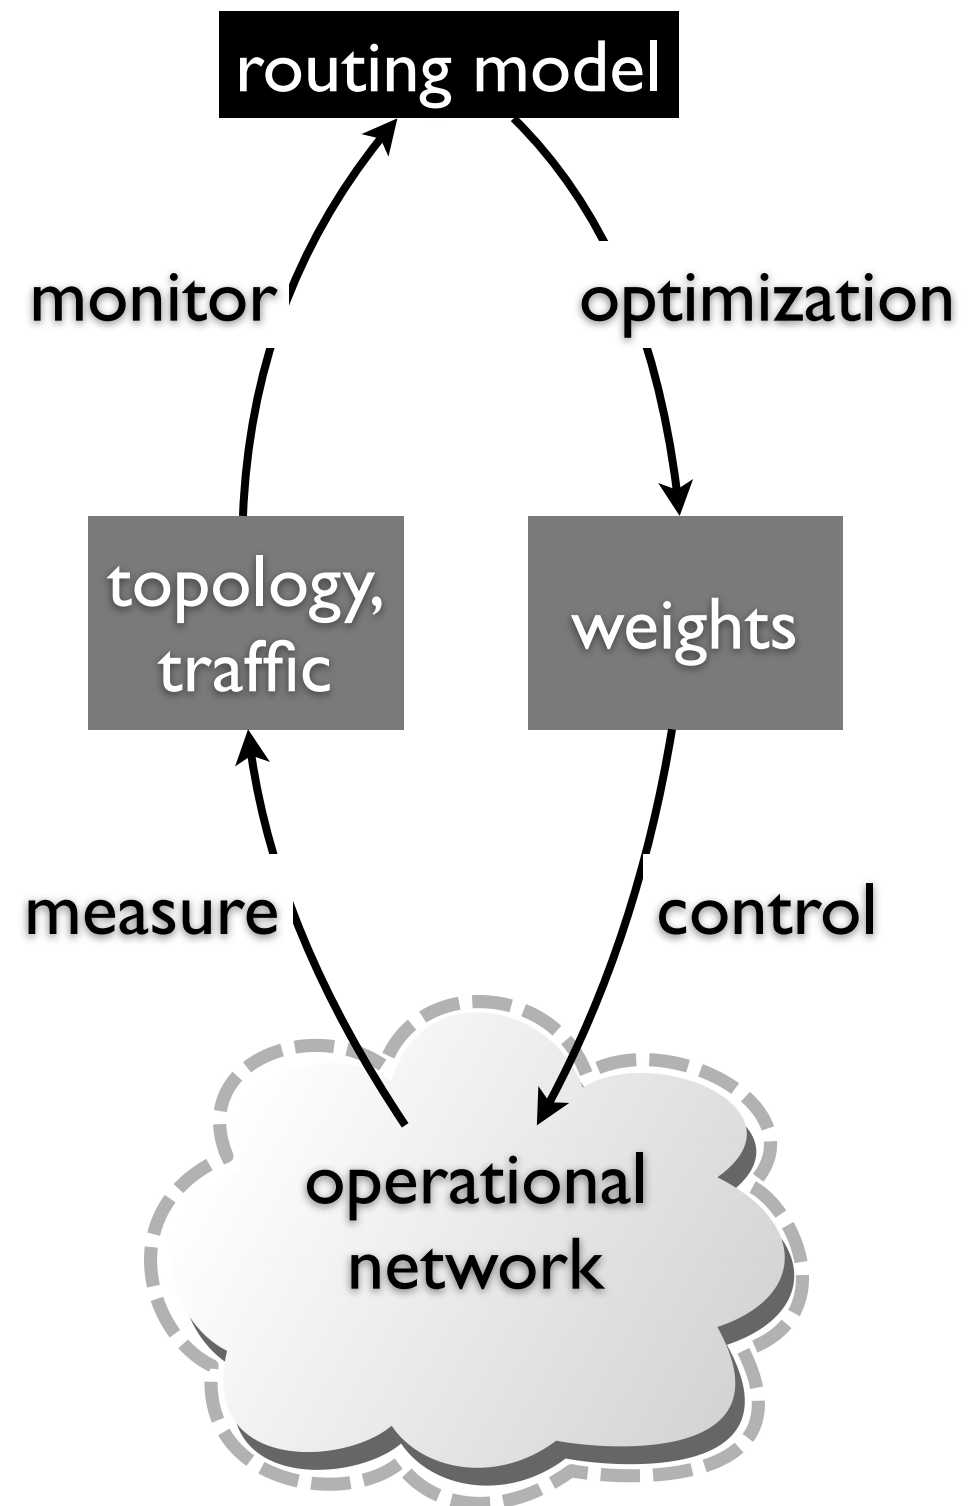
traffic engineering framework

centralized control

- stable
- lower overhead
- diverse performance objective

link weights express the routing configuration

- compatibility
- concise
- default weights and backup routes



performance

objective: link cost

- cost of using a link increases with utilization, explosive growth as utilization exceeds 100%

global optimization close to optimal


- can handle 70% more demands than Cisco or unit weights

OPT: can direct traffic along any paths in any proportions

InvCapOSPF: (Cisco) set link weight inversely proportional to its capacity

UnitOSPF: set all weights to 1

AdvancedOSPF: global optimization



results on an AT&T backbone with a projected traffic matrix

discussion

centralized control

- stable
- lower overhead
- diverse performance objective

link weights express the routing configuration

- compatibility
- concise
- default weights and backup routes

discussion

centralized control

- stable
- lower overhead
- diverse performance objective

link weights express the routing configuration

- compatibility
- concise
- default weights and backup routes

the good

- centralized control, shared with SDN
- can express diverse network-wide objective

discussion

centralized control

- stable
- lower overhead
- diverse performance objective

link weights express the routing configuration

- compatibility
- concise
- default weights and backup routes

the good

- centralized control, shared with SDN
- can express diverse network-wide objective

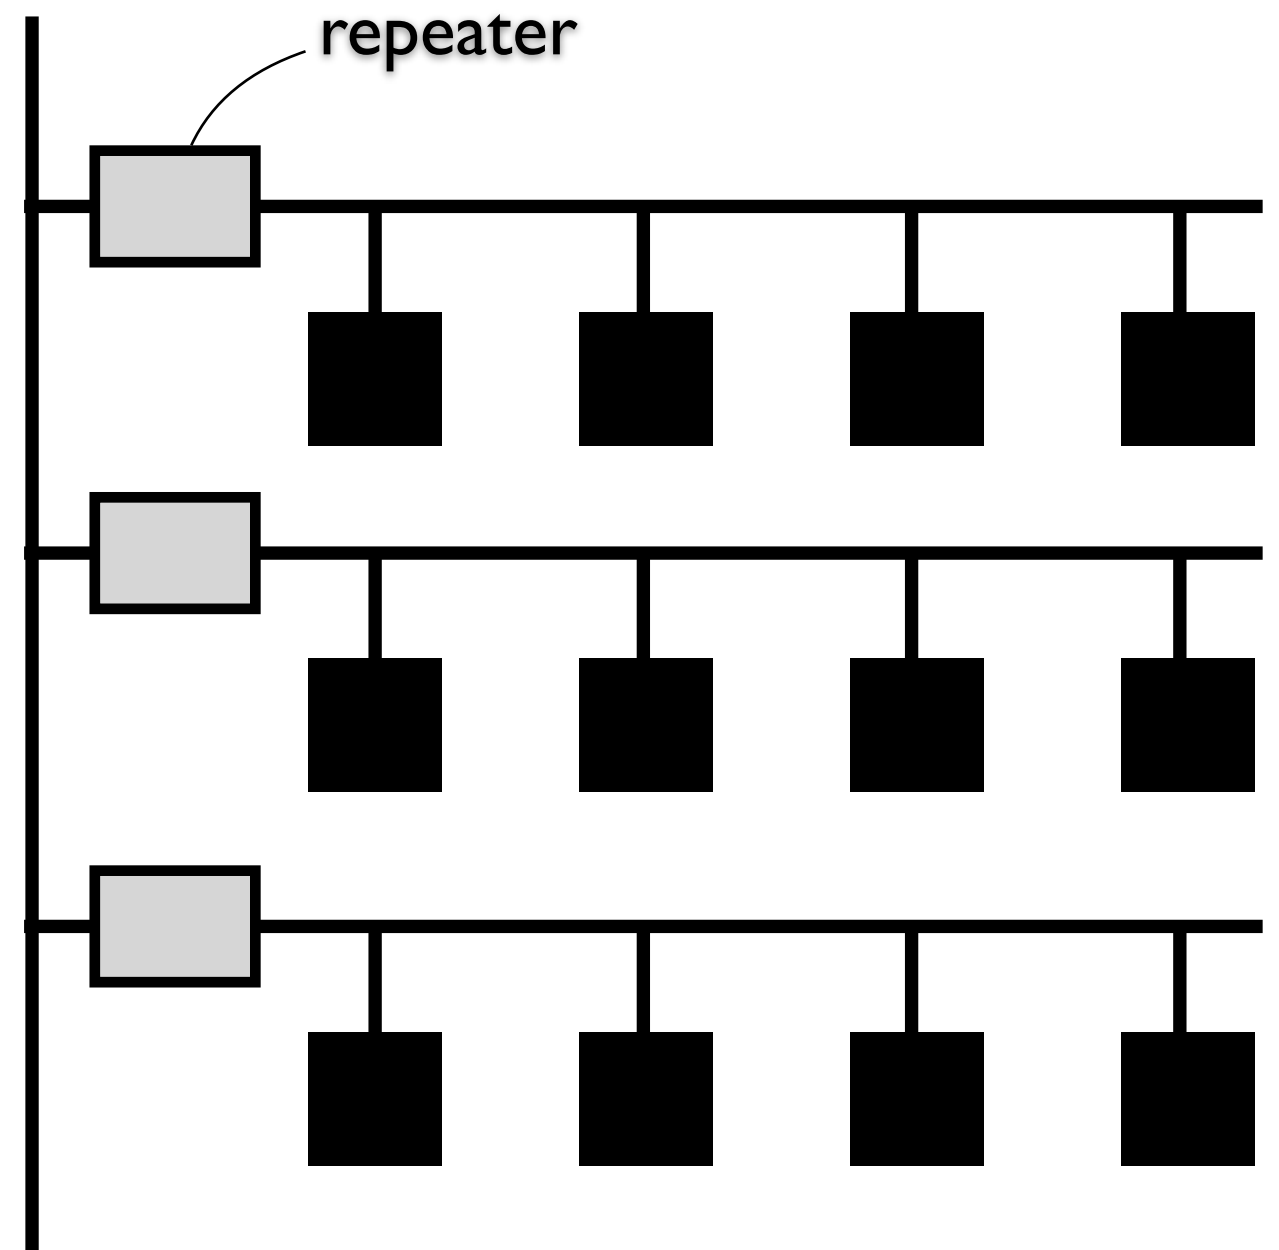

the bad and the ugly?

- inflexible: limited expressiveness
- indirect: link weights do not embed any semantics of higher-level network-wide goals

Ethernet

Ethernet

a local area network (LAN)



Ethernet

broadcasting communication

- message placed on the Ethernet is broadcast over
- ## media access control (MAC) algorithm
- **1-persistent**
 - adaptor with a frame to send transmits with probability 1 whenever busy line goes idle
 - **exponential backoff**
 - upon detection of collision, adaptor stops transmission, waits a certain amount of time (and doubles before trying again)

Ethernet — “the” LAN technology

“zero” configuration

- extremely simple to configure and maintain: no switch, no routing, no configuration tables

inexpensive

- cable is cheap
- only cost: the adaptor

switched Ethernet ...

discussion

distributed control

- coordination of access is distributed among contending senders
- colliding senders: random retransmission intervals
- switching is distributed among the recipients

no central controller

- eliminate the reliability problem

zero configuration

discussion

distributed control

- coordination of access is distributed among contending senders
- colliding senders: random retransmission intervals
- switching is distributed among the recipients

no central controller

- eliminate the reliability problem

zero configuration

SDN abandons distributed control for simplicity

discussion

distributed control

- coordination of access is distributed among contending senders
- colliding senders: random retransmission intervals
- switching is distributed among the recipients

no central controller

- eliminate the reliability problem

zero configuration

SDN abandons distributed control for simplicity

reliability, a challenge for SDN

discussion

distributed control

- coordination of access is distributed among contending senders
- colliding senders: random retransmission intervals
- switching is distributed among the recipients

no central controller

- eliminate the reliability problem

zero configuration

SDN abandons distributed control for simplicity

reliability, a challenge for SDN

a goal shared with SDN

discussion

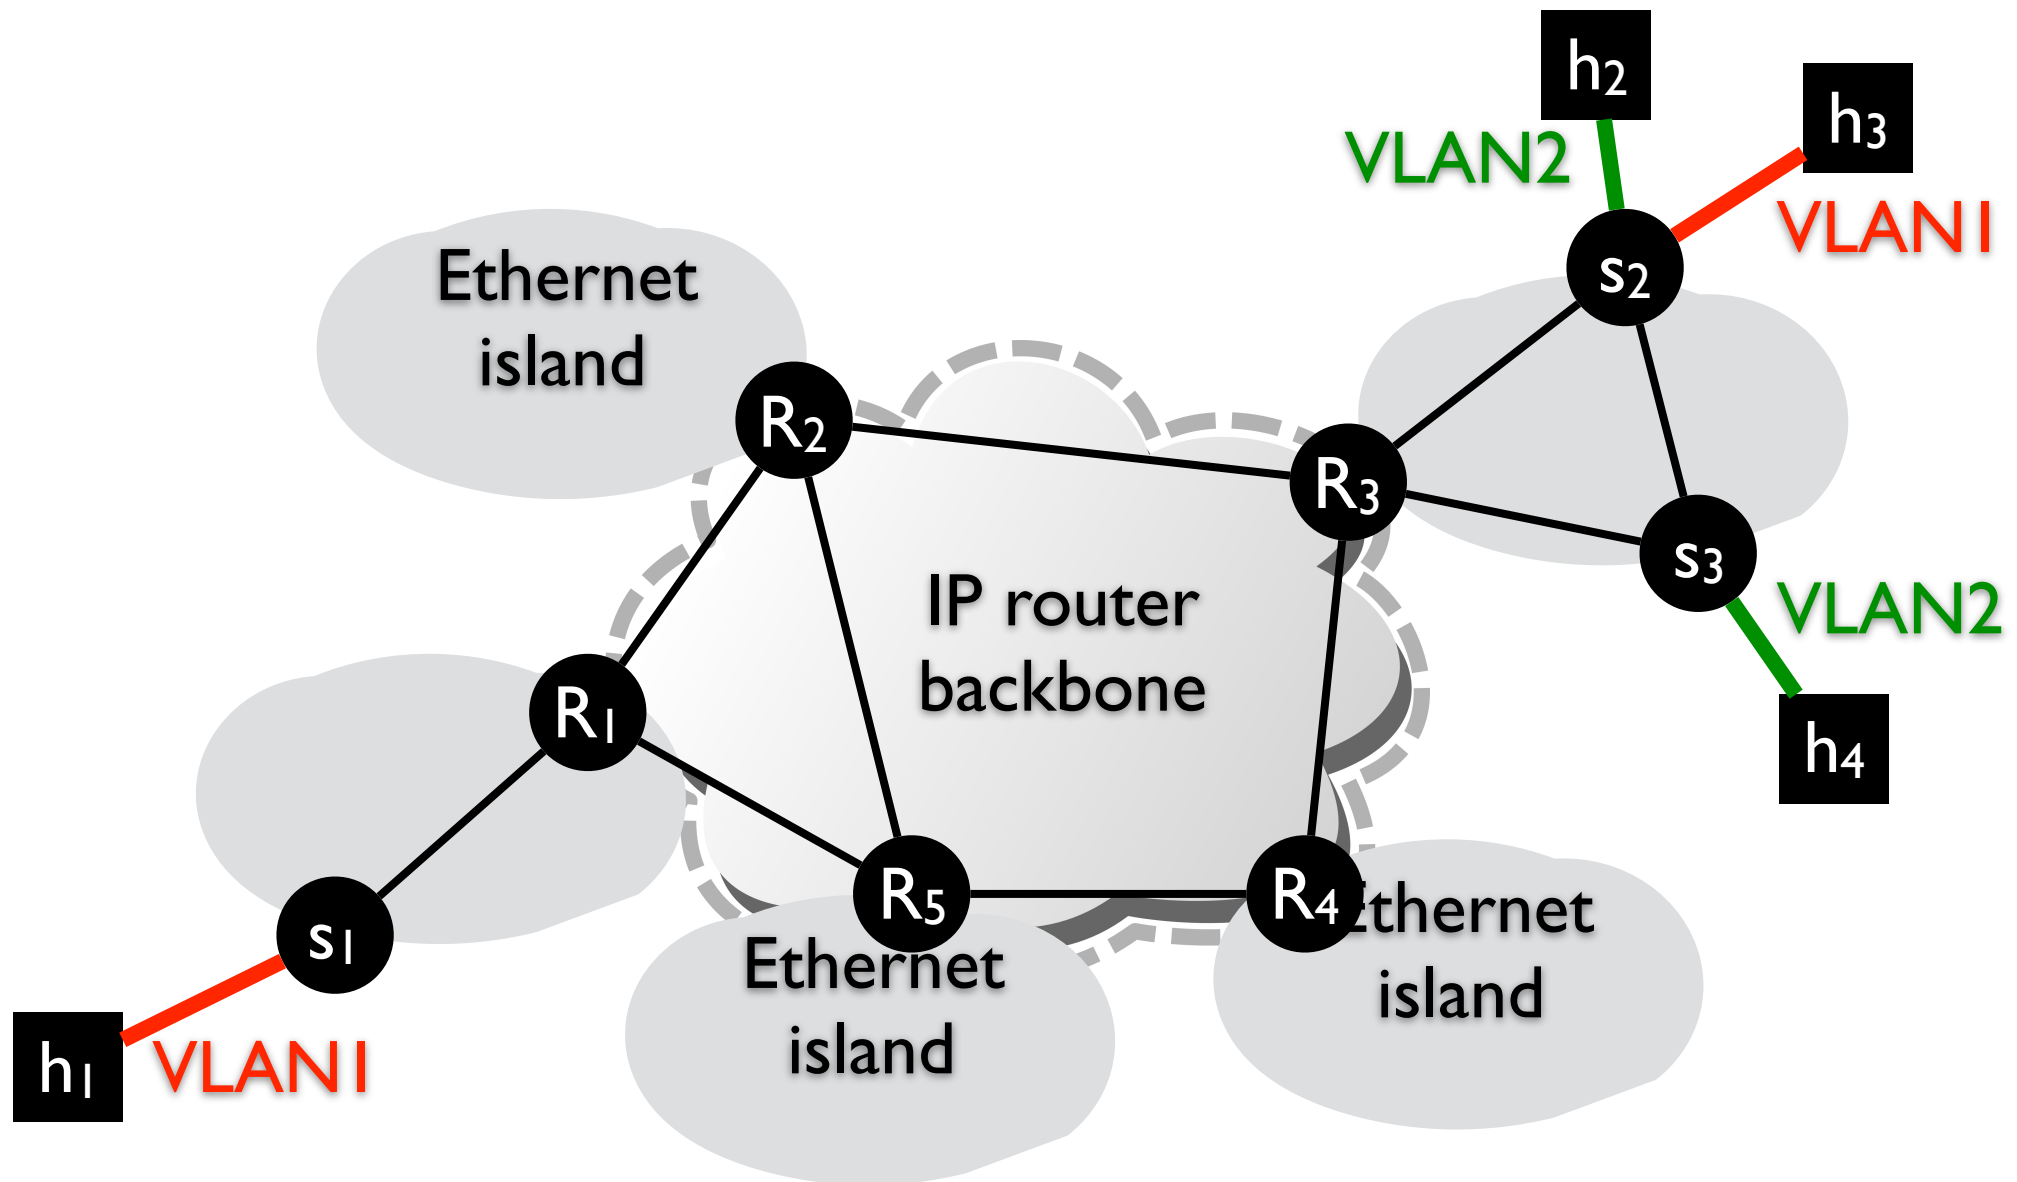
Ethernet is a real gem

- despite limitations — scalability, best effort delivery
- a rare combination of distributed control and simplicity
- arbitration of conflicting transmission demands is both distributed and statistical

VLAN for campus networks

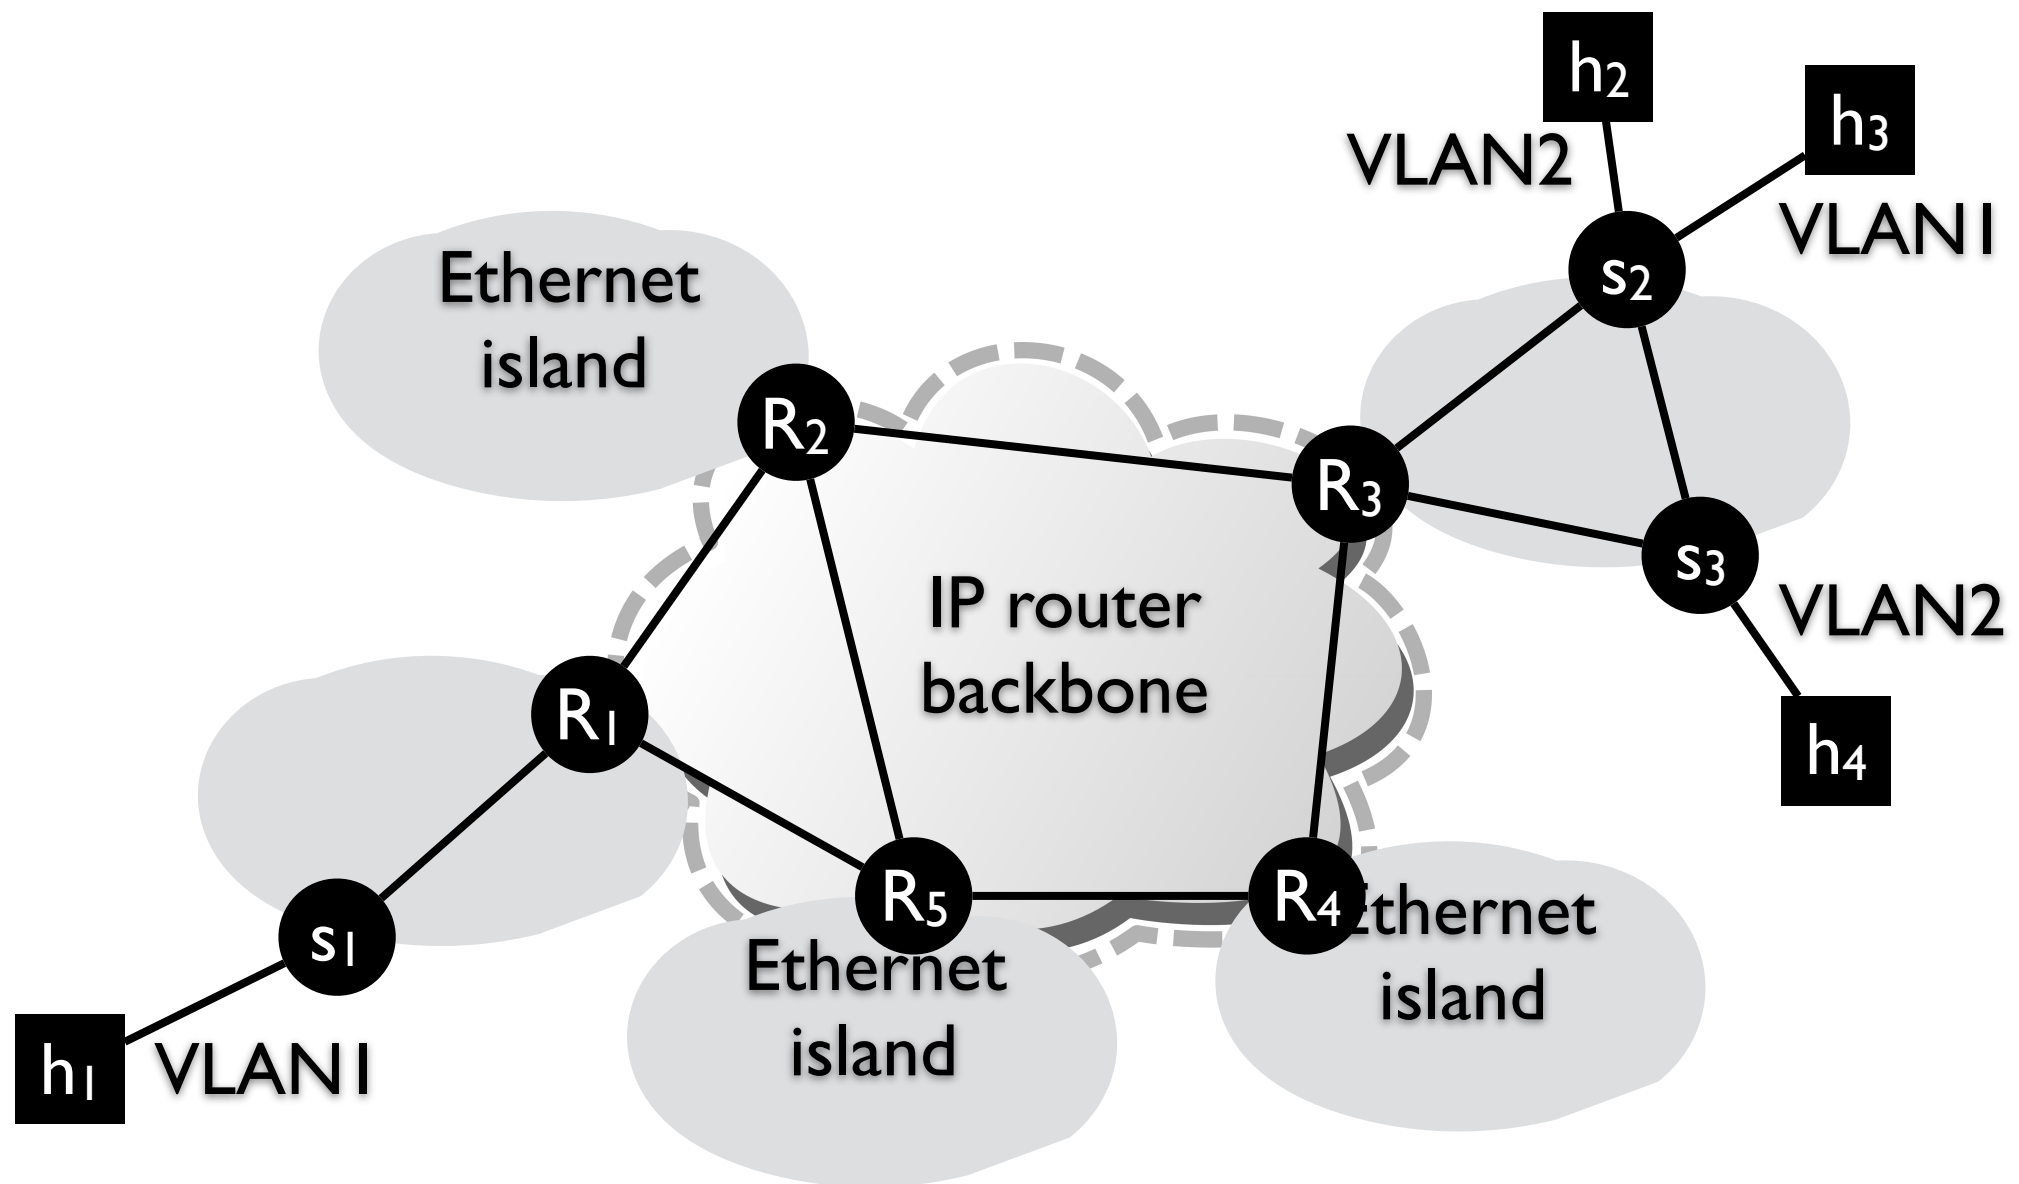
VLAN

connect hosts in the same broadcast domain, independent of their physical location



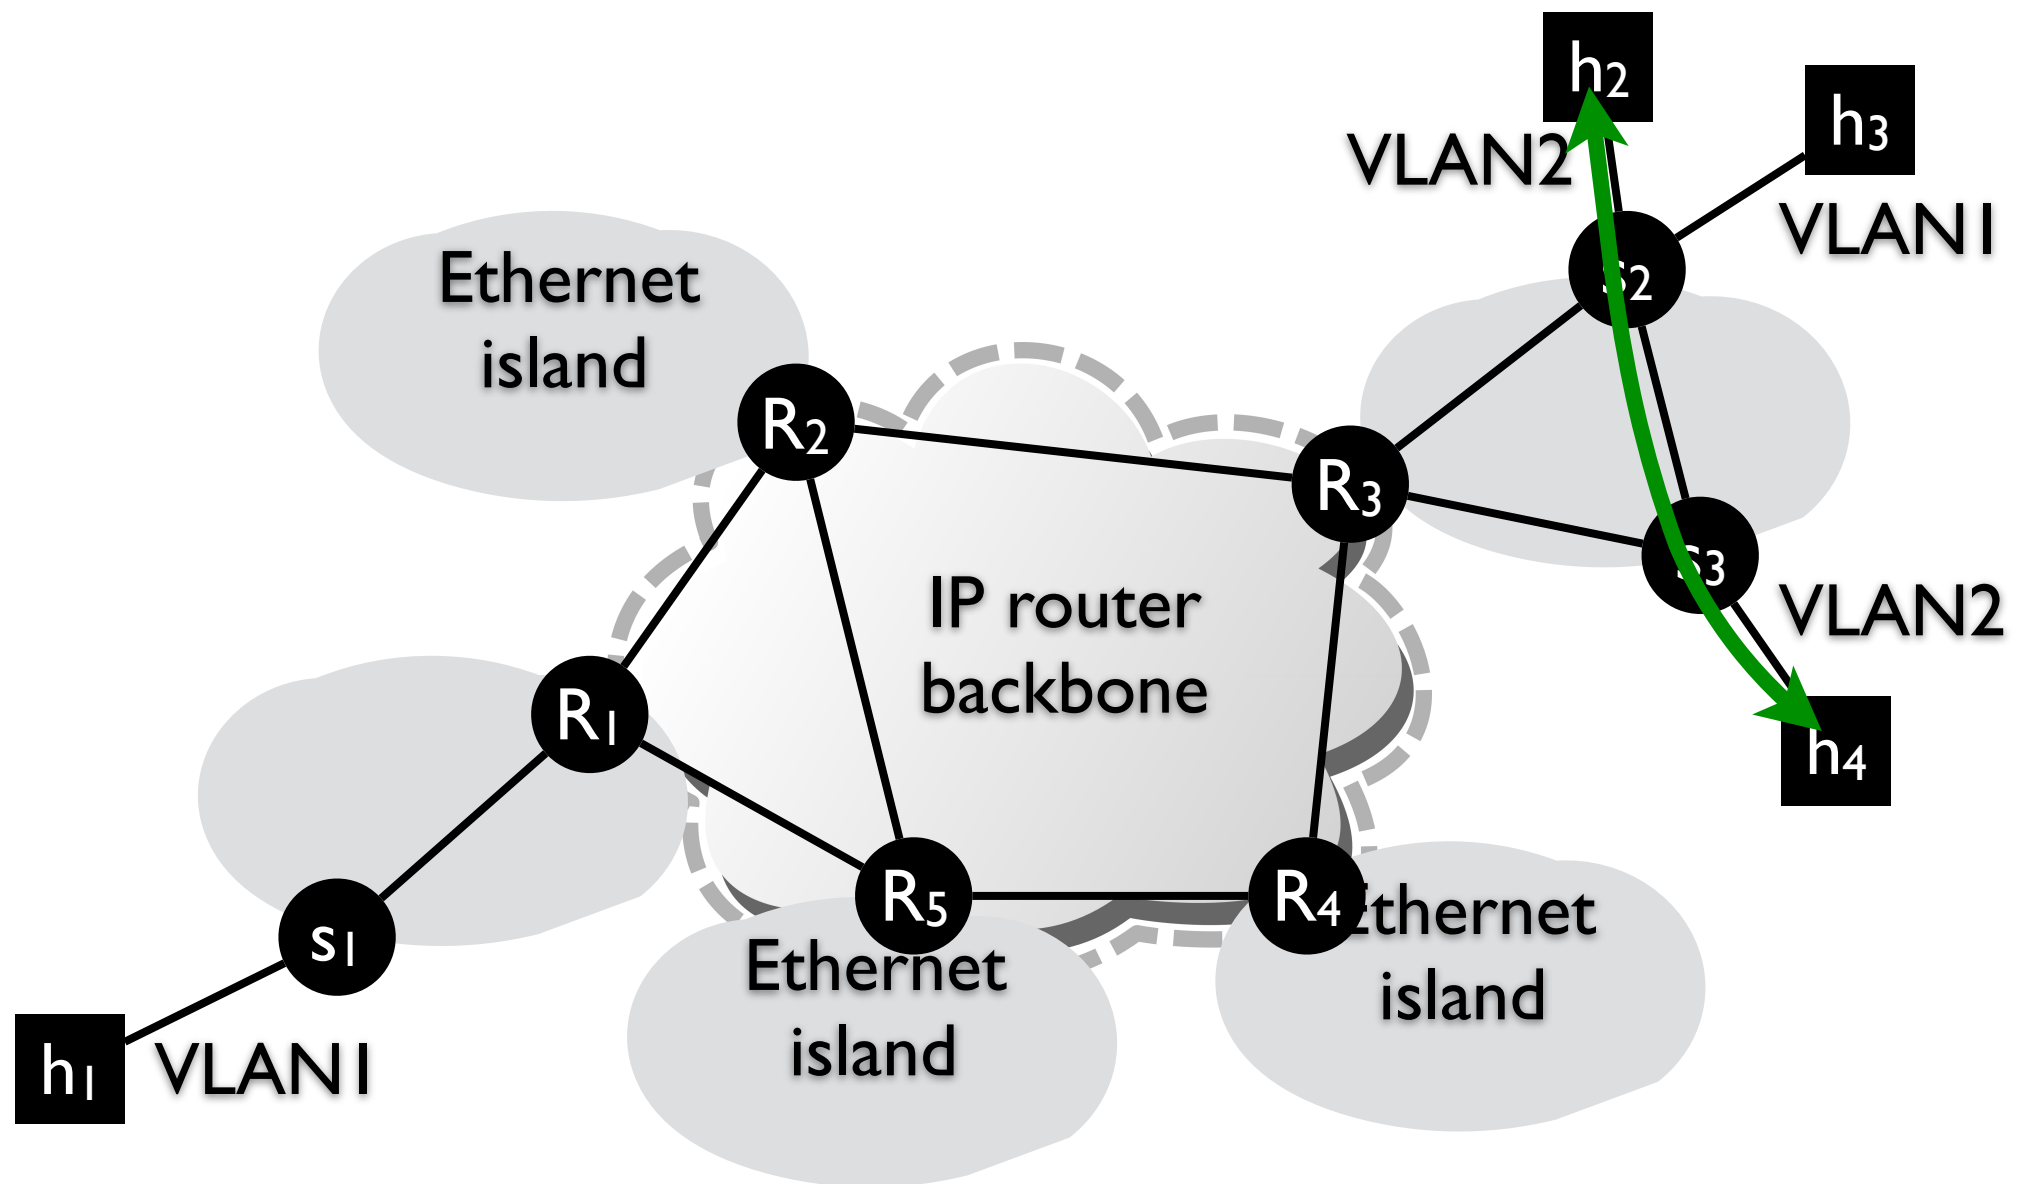
communication within a VLAN

h2 and h4 communicate over the spanning tree in VLAN2



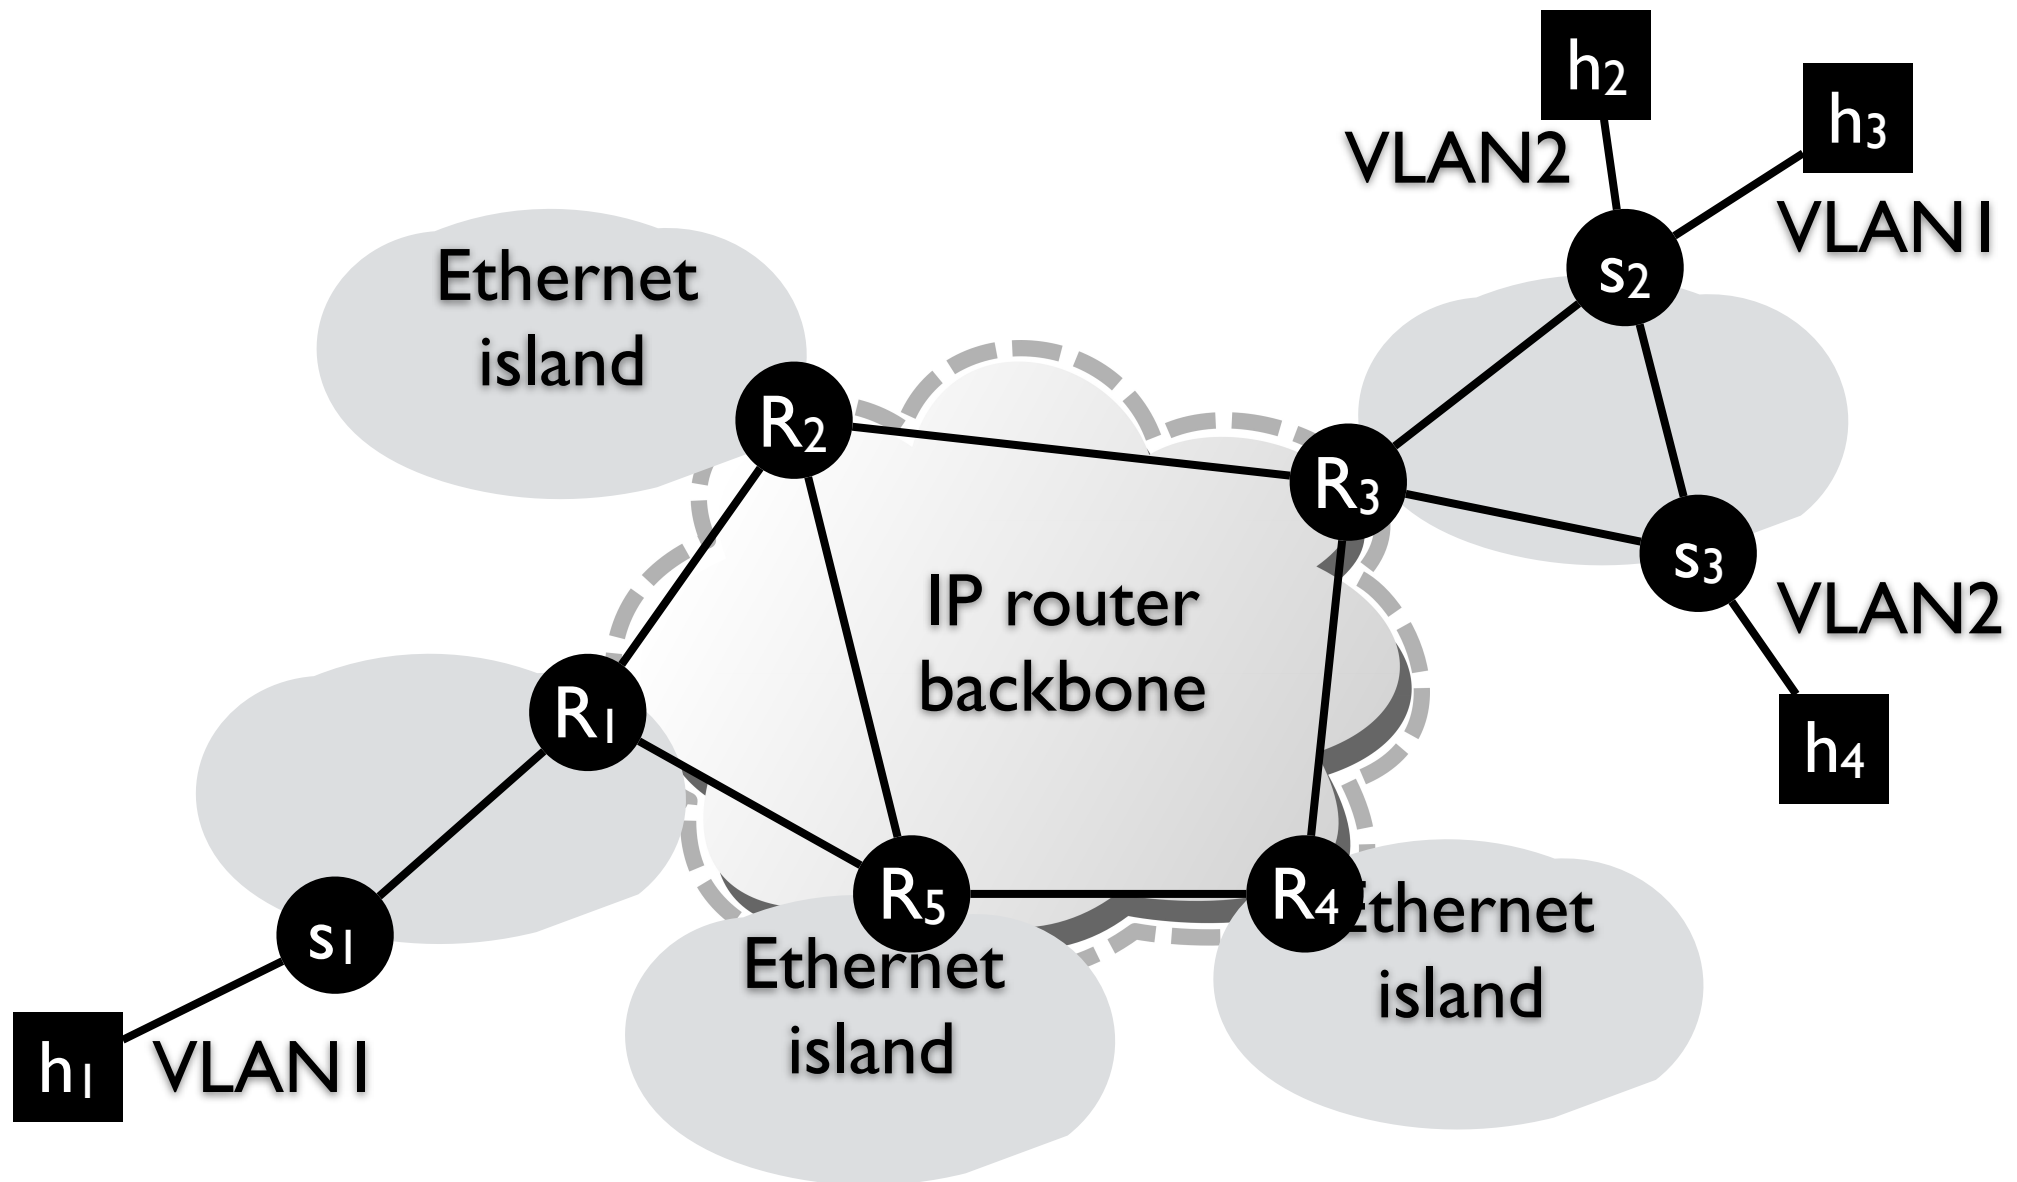
communication within a VLAN

h2 and h4 communicate over the spanning tree in VLAN2



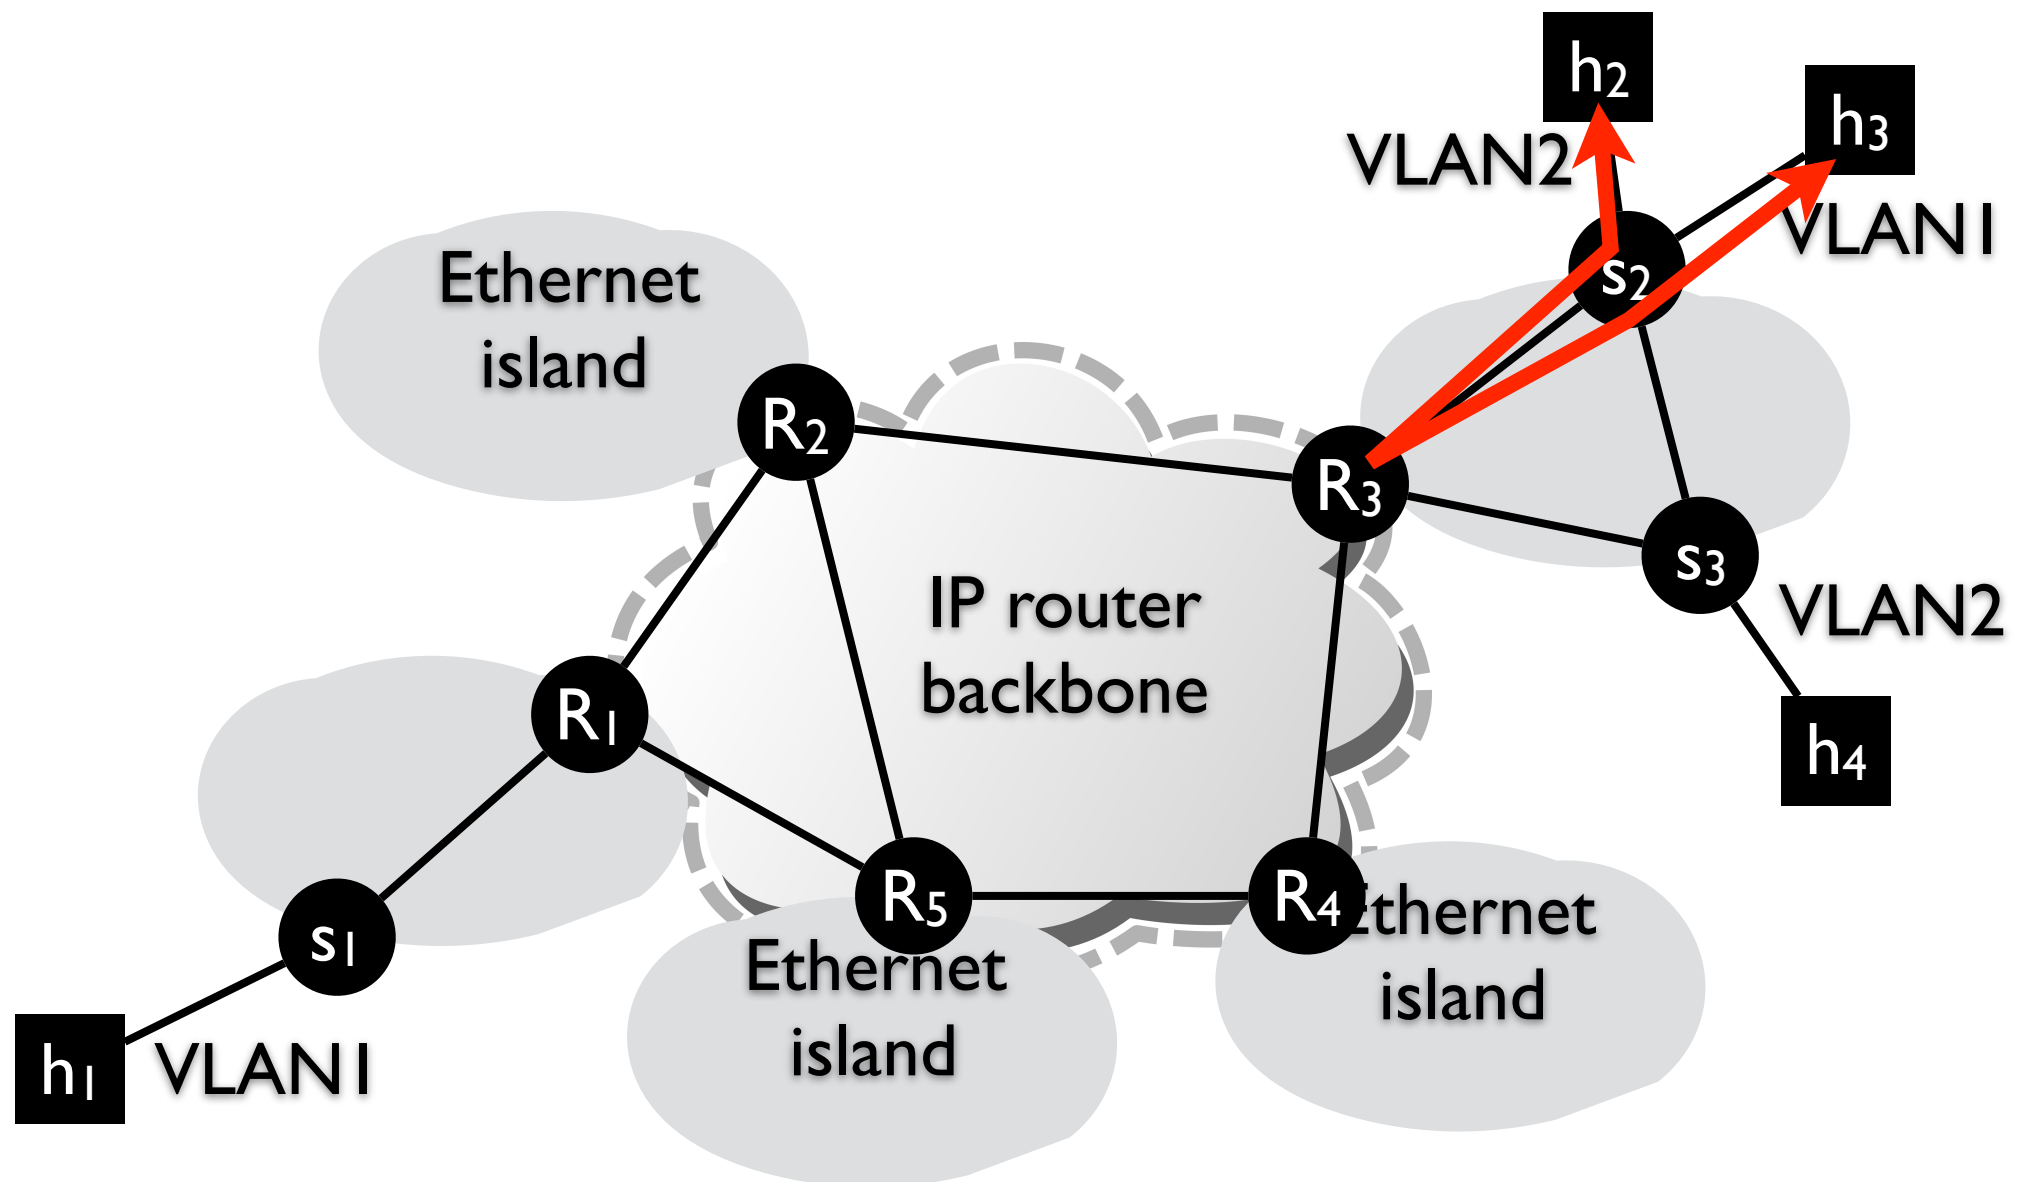
communication between VLANs

- each VLAN has a IP prefix
- IP routers forward packets based on these prefixes



communication between VLANs

- each VLAN has a IP prefix
- IP routers forward packets based on these prefixes



VLAN usage in campus networks

VLAN usage in campus networks

VLAN widely used for various policy objectives

VLAN usage in campus networks

VLAN widely used for various policy objectives

- scoping broadcast traffic

VLAN usage in campus networks

VLAN widely used for various policy objectives

- scoping broadcast traffic
 - limiting flood overhead

VLAN usage in campus networks

VLAN widely used for various policy objectives

- scoping broadcast traffic
 - limiting flood overhead
 - e.g., divide large networks into multiple VLANs

VLAN usage in campus networks

VLAN widely used for various policy objectives

- scoping broadcast traffic
 - limiting flood overhead
 - e.g., divide large networks into multiple VLANs
 - e.g., assign each building a different IP subnet, each grouped into a VLAN

VLAN usage in campus networks

VLAN widely used for various policy objectives

- scoping broadcast traffic
 - limiting flood overhead
 - e.g., divide large networks into multiple VLANs
 - e.g., assign each building a different IP subnet, each grouped into a VLAN
 - protecting security and privacy

VLAN usage in campus networks

VLAN widely used for various policy objectives

- scoping broadcast traffic
 - limiting flood overhead
 - e.g., divide large networks into multiple VLANs
 - e.g., assign each building a different IP subnet, each grouped into a VLAN
 - protecting security and privacy
 - e.g., separate VLANs for faculty, students

VLAN usage in campus networks

VLAN widely used for various policy objectives

- scoping broadcast traffic
 - limiting flood overhead
 - e.g., divide large networks into multiple VLANs
 - e.g., assign each building a different IP subnet, each grouped into a VLAN
 - protecting security and privacy
 - e.g., separate VLANs for faculty, students
- simplifying access control

VLAN usage in campus networks

VLAN widely used for various policy objectives

- **scoping broadcast traffic**
 - limiting flood overhead
 - e.g., divide large networks into multiple VLANs
 - e.g., assign each building a different IP subnet, each grouped into a VLAN
 - protecting security and privacy
 - e.g., separate VLANs for faculty, students
- **simplifying access control**
 - VLANs group hosts with common access control policy

VLAN usage in campus networks

VLAN widely used for various policy objectives

- **scoping broadcast traffic**
 - limiting flood overhead
 - e.g., divide large networks into multiple VLANs
 - e.g., assign each building a different IP subnet, each grouped into a VLAN
 - protecting security and privacy
 - e.g., separate VLANs for faculty, students
- **simplifying access control**
 - VLANs group hosts with common access control policy
 - e.g., allow user machines (faculty, student VLANs) to server (infrastructure VLAN)

VLAN usage in campus networks

VLAN widely used for various policy objectives

- **scoping broadcast traffic**
 - limiting flood overhead
 - e.g., divide large networks into multiple VLANs
 - e.g., assign each building a different IP subnet, each grouped into a VLAN
 - protecting security and privacy
 - e.g., separate VLANs for faculty, students
- **simplifying access control**
 - VLANs group hosts with common access control policy
 - e.g., allow user machines (faculty, student VLANs) to server (infrastructure VLAN)
- **decentralizing network management**

VLAN usage in campus networks

VLAN widely used for various policy objectives

- **scoping broadcast traffic**
 - limiting flood overhead
 - e.g., divide large networks into multiple VLANs
 - e.g., assign each building a different IP subnet, each grouped into a VLAN
 - protecting security and privacy
 - e.g., separate VLANs for faculty, students
- **simplifying access control**
 - VLANs group hosts with common access control policy
 - e.g., allow user machines (faculty, student VLANs) to server (infrastructure VLAN)
- **decentralizing network management**
 - delegate tasks to individual VLANs

VLAN usage in campus networks

VLAN widely used for various policy objectives

- **scoping broadcast traffic**
 - limiting flood overhead
 - e.g., divide large networks into multiple VLANs
 - e.g., assign each building a different IP subnet, each grouped into a VLAN
 - protecting security and privacy
 - e.g., separate VLANs for faculty, students
- **simplifying access control**
 - VLANs group hosts with common access control policy
 - e.g., allow user machines (faculty, student VLANs) to server (infrastructure VLAN)
- **decentralizing network management**
 - delegate tasks to individual VLANs
 - e.g., one IT group manages “classroom VLAN” across 60 buildings

VLAN usage in campus networks

VLAN widely used for various policy objectives

- **scoping broadcast traffic**
 - limiting flood overhead
 - e.g., divide large networks into multiple VLANs
 - e.g., assign each building a different IP subnet, each grouped into a VLAN
 - protecting security and privacy
 - e.g., separate VLANs for faculty, students
- **simplifying access control**
 - VLANs group hosts with common access control policy
 - e.g., allow user machines (faculty, student VLANs) to server (infrastructure VLAN)
- **decentralizing network management**
 - delegate tasks to individual VLANs
 - e.g., one IT group manages “classroom VLAN” across 60 buildings
- **enabling host mobility**

problem: inexpressiveness

problem: inexpressiveness

built-in protocol limitation

- number of VLANs < 4096 (12-bit header field)
 - multiple *isolated* group in the same VLAN
 - *isolated* VLANs share VLAN ID
- number of hosts per VLAN (flooding, spanning tree)
 - artificially divide large group into multiple VLANs

problem: inexpressiveness

built-in protocol limitation

- number of VLANs < 4096 (12-bit header field)
 - multiple *isolated* group in the same VLAN
 - *isolated* VLANs share VLAN ID
- number of hosts per VLAN (flooding, spanning tree)
 - artificially divide large group into multiple VLANs

unfit for traffic grouping

- VLAN naturally groups end hosts
 - unexpected security breach: student plugs into a hub in a faculty office
 - restricted policy: a faculty on faculty VLAN cannot participate in admin

problem: complex configuration

problem: complex configuration

tight coupling between VLANs and IP

- wasting IP addresses, complex IP assignment

problem: complex configuration

tight coupling between VLANs and IP

- wasting IP addresses, complex IP assignment

spanning tree computation

- explicitly configure switches to form spanning tree
 - determining which links participate in which VLAN is difficult
 - trunk links become inconsistent after network evolves
 - over-loading root bridge: same switch selected as the root in multiple VLANs

discussion: the bad and the ugly?

VLAN mechanism

- indirect and inflexible
 - VLAN creates broadcast domain for end-hosts
 - built-in protocol limitation
- low-level realization
 - explicit access port, trunk port

diverse high-level policy

- scoping traffic
- access control
- delegate management

discussion: the bad and the ugly?

VLAN mechanism

- indirect and inflexible
 - VLAN creates broadcast domain for end-hosts
 - built-in protocol limitation
- low-level realization
 - explicit access port, trunk port

diverse high-level policy

- scoping traffic
- access control
- delegate management

SDN mechanisms

- direct, flexible
- high-level abstraction

discussion: the bad and the ugly?

VLAN mechanism

- indirect and inflexible
 - VLAN creates broadcast domain for end-hosts
 - built-in protocol limitation
- low-level realization
 - explicit access port, trunk port

diverse high-level policy

- scoping traffic
- access control
- delegate management

SDN mechanisms

- direct, flexible
- high-level abstraction

the diverse high-level policy is a goal shared with SDN

to do

submit reviews by **5pm, September 8**

- 4D and Ethane papers

next time

- centralized control
- database defined networking