WHEN THE MODEL REALLY MATTERS:

THE COMPOSITIONAL ARCHITECTURE OF THE INTERNET*

Pamela Zave

AT&T Labs and Princeton University

New Jersey, USA

*Joint work with Jennifer Rexford of Princeton.

IN 1992

THE EXPLOSIVE GROWTH OF THE
WORLD-WIDE WEB BEGAN

AND IN 1993

THE LAST MAJOR CHANGE WAS MADE
TO THE "CLASSIC" INTERNET ARCHITECTURE

WHAT HAS HAPPENED SINCE 1993?

- most of the world's . . .

  . . . telecommunication infrastructure

  . . . entertainment distribution . . .

  has moved to the Internet

- an explosion of security threats
- most networked devices are mobile
- cloud computing
- exhaustion of the IP address space
- the need for elastic resource allocation instead of over-provisioning

A CONUNDRUM:

The "classic" Internet architecture (how experts describe the Internet) has not changed since 1993, . . .

. . . yet the Internet has met all these new challenges, at least to some extent.

also, implementation technology has changed dramatically— networks are now software systems

THE “CLASSIC” INTERNET ARCHITECTURE

APPLICATION LAYER	applications and mnemonic names
TRANSPORT LAYER	reliable byte streams, datagrams
NETWORK LAYER	best-effort global packet delivery
LINK LAYER	best-effort local packet delivery
PHYSICAL LAYER	diverse physical media (wires, optical fibers, radio channels)

so we expect
a typical packet
to look like this

HTTP header
TCP header
IP header
Ethernet header

THE REALITY: THIS IS A TYPICAL PACKET IN THE AT&T BACKBONE

packets sampled elsewhere would look different, but might be equally complex

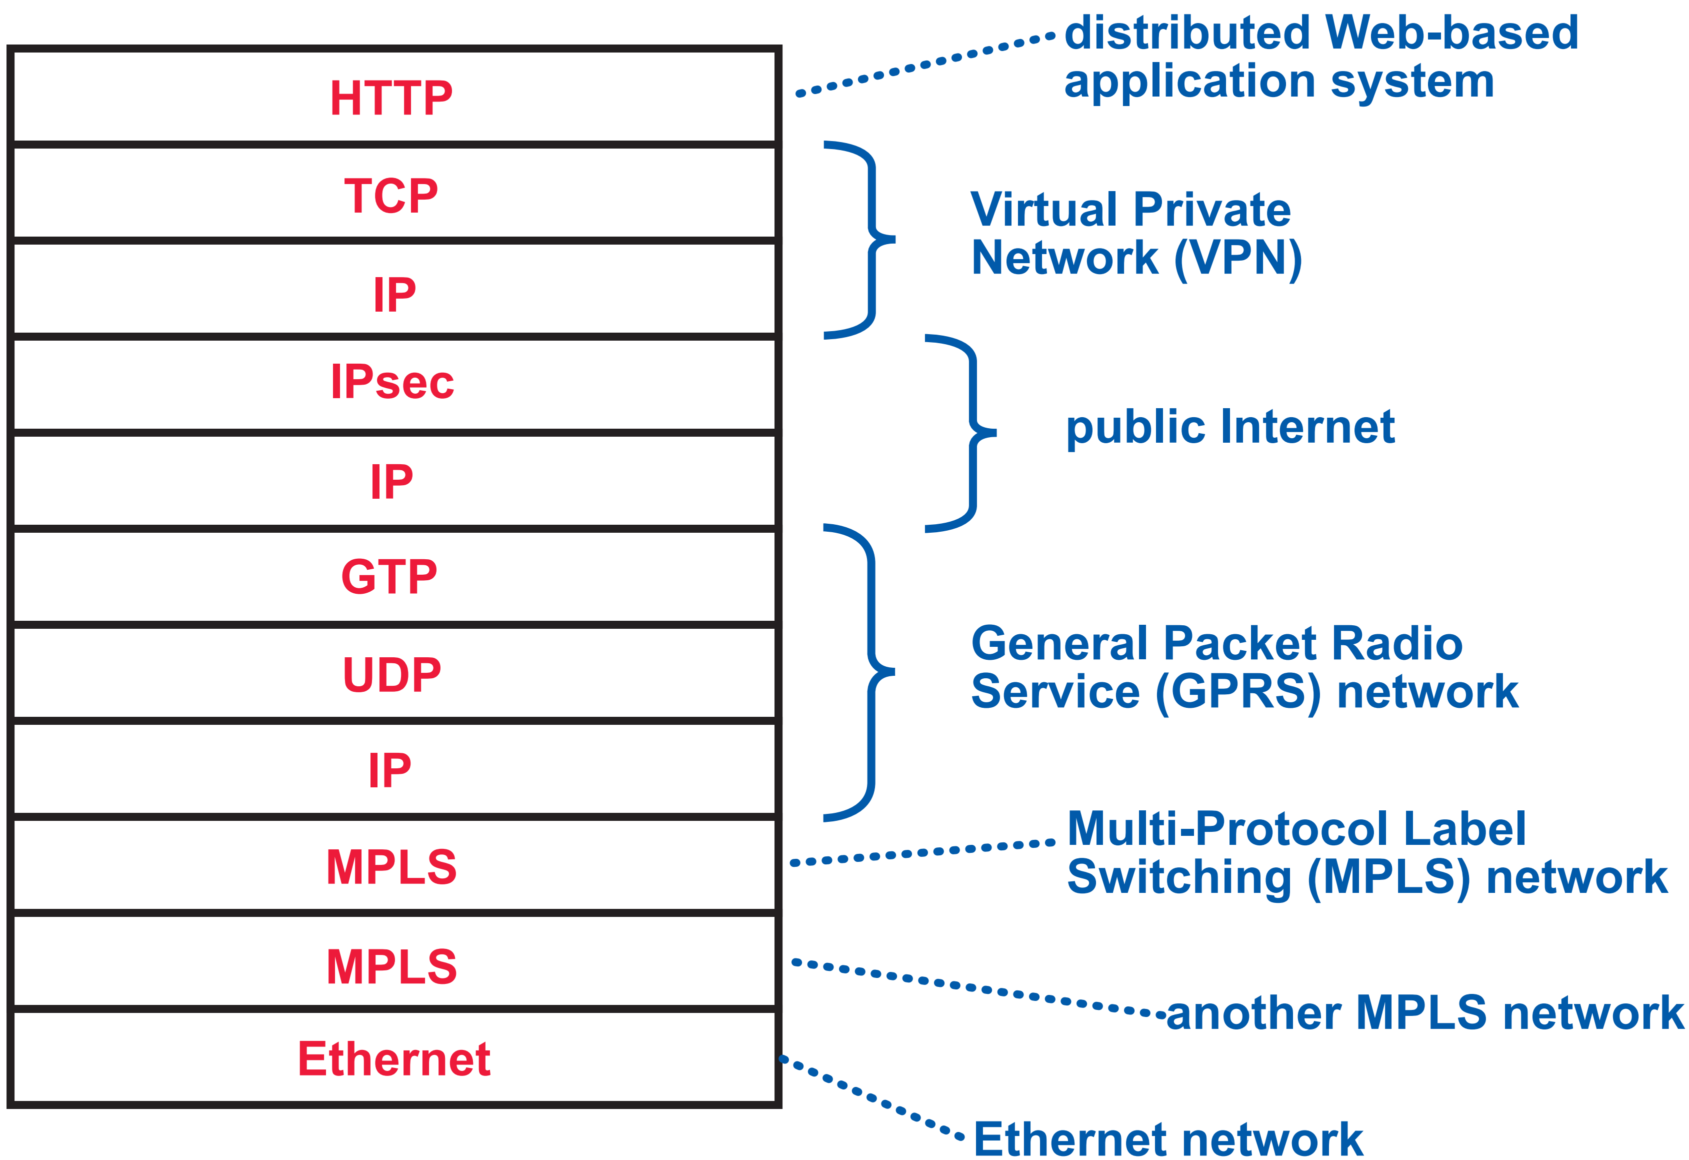

WHAT IS TAUGHT? WHAT IS RESEARCHED?

1. Teach the classic Internet architecture and how basic services (Web, email, file transfer) are implemented.

2. Note that there are exceptions everywhere, but we cannot think about them very much without being overwhelmed by complexity.

3. All other topics are studied and researched in isolation, as if they were independent rather than different aspects of the same artifacts.

   e.g., security, mobility, cloud computing, streaming, the Internet of Things— whatever is "hot" at the moment

4. Assume that solutions to narrow problems can all be composed by cramming them into the network layer of the classic Internet architecture.

   which is not the way changes are made now, is not modular or verifiable, and probably not optimal

WHY DO WE NEED A BETTER MODEL OF NETWORKING?

1. Teach the classic Internet architecture and how basic services (Web, email, file transfer) are implemented.

2. Note that there are exceptions everywhere, but we cannot think about them very much without being overwhelmed by complexity.

To talk about networking as it really is, without being overwhelmed by complexity.

To teach principles rather than just details.

3. All other topics are studied and researched in isolation, as if they were independent rather than different aspects of the same artifacts.

To understand how the aspects of networking compose.

4. Assume that solutions to narrow problems can all be composed by cramming them into the network layer of the classic Internet architecture.

A BETTER MODEL: THE INTERNET IS A FLEXIBLE COMPOSITION OF MANY NETWORKS

global networking
as we know it

many more than those acknowledged
in the classic architecture

each network has all the same
basic mechanisms, . . .

. . . but in each network they are
specialized for a particular . . .
. . . purpose,
. . . membership scope,
. . . geographical span, and
. . . level of abstraction

because all networks have
fundamental similarity, they all have
common interfaces for composition

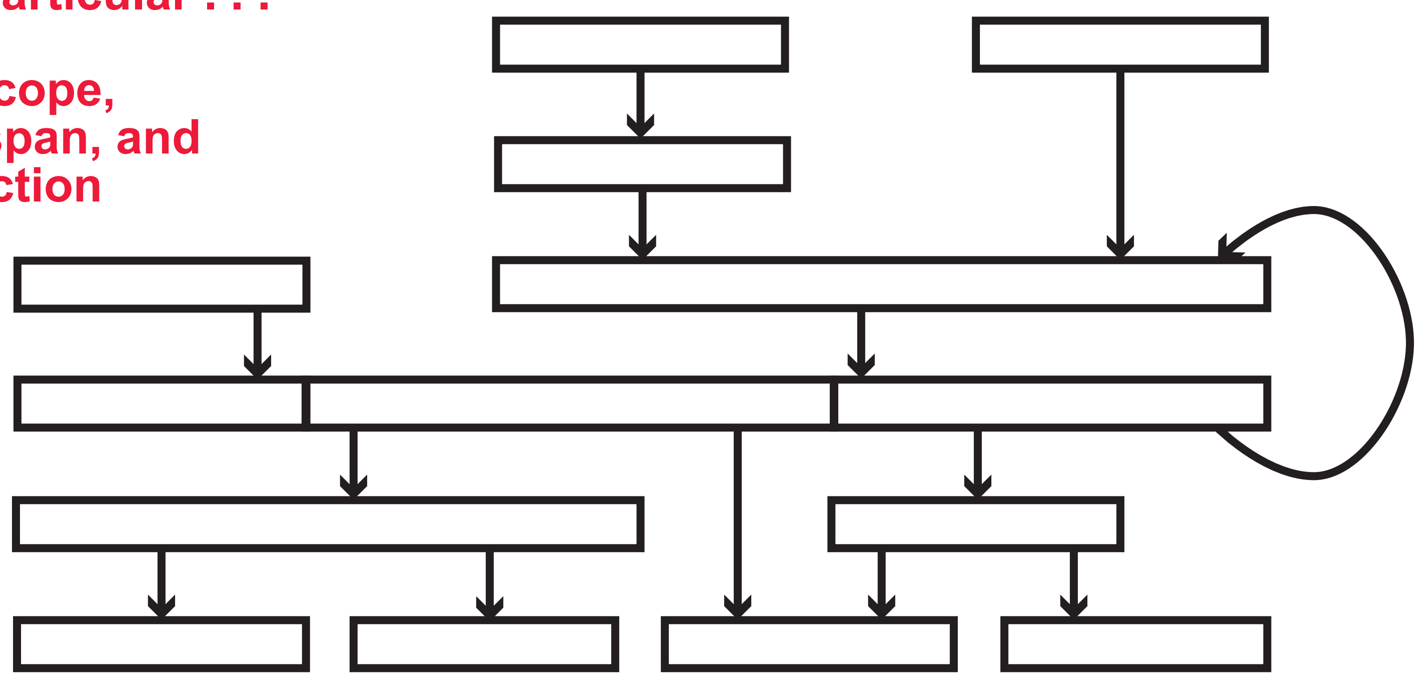

the Internet protocol suite implements a general-purpose network design and
is available on most networked devices—so it is re-used for many purposes

OLD: THE END-TO-END PRINCIPLE

The functions of a network should be minimized, so that it serves everyone efficiently, . . .

. . . and whenever possible, services should be implemented in endpoint machines.

or, "smart edge, dumb network"

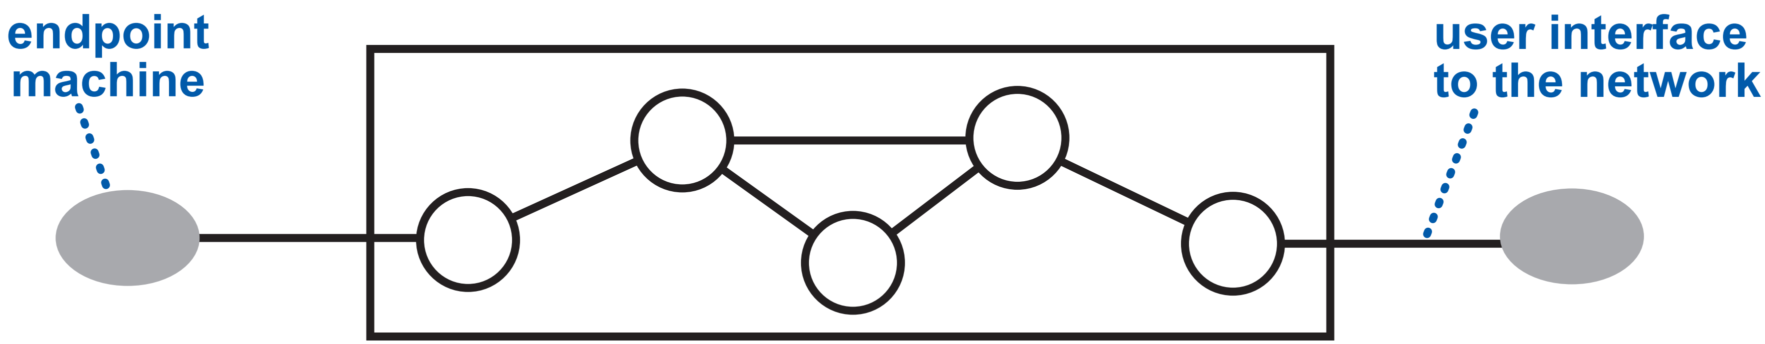

the End-to-End Principle is a design principle, but it has been so influential that it is assumed to be descriptive

today there are many exceptions:

- many service functions are implemented inside the network, . . .

  . . . by middleboxes and programmable routers

- cannot control network performance without the cooperation of endpoints

today we know . . .

. . . that if we want to verify network services . . .

. . . we must include in our model all the agents involved in providing those services

NEW: USER INTERFACES ARE INSIDE MACHINES

the user of a network is a distributed application system— its modules must communicate through network services

an instance of network service is a session; a session transmits a group of packets that the user sees as related

modules on the same machine communicate through its operating system or hardware

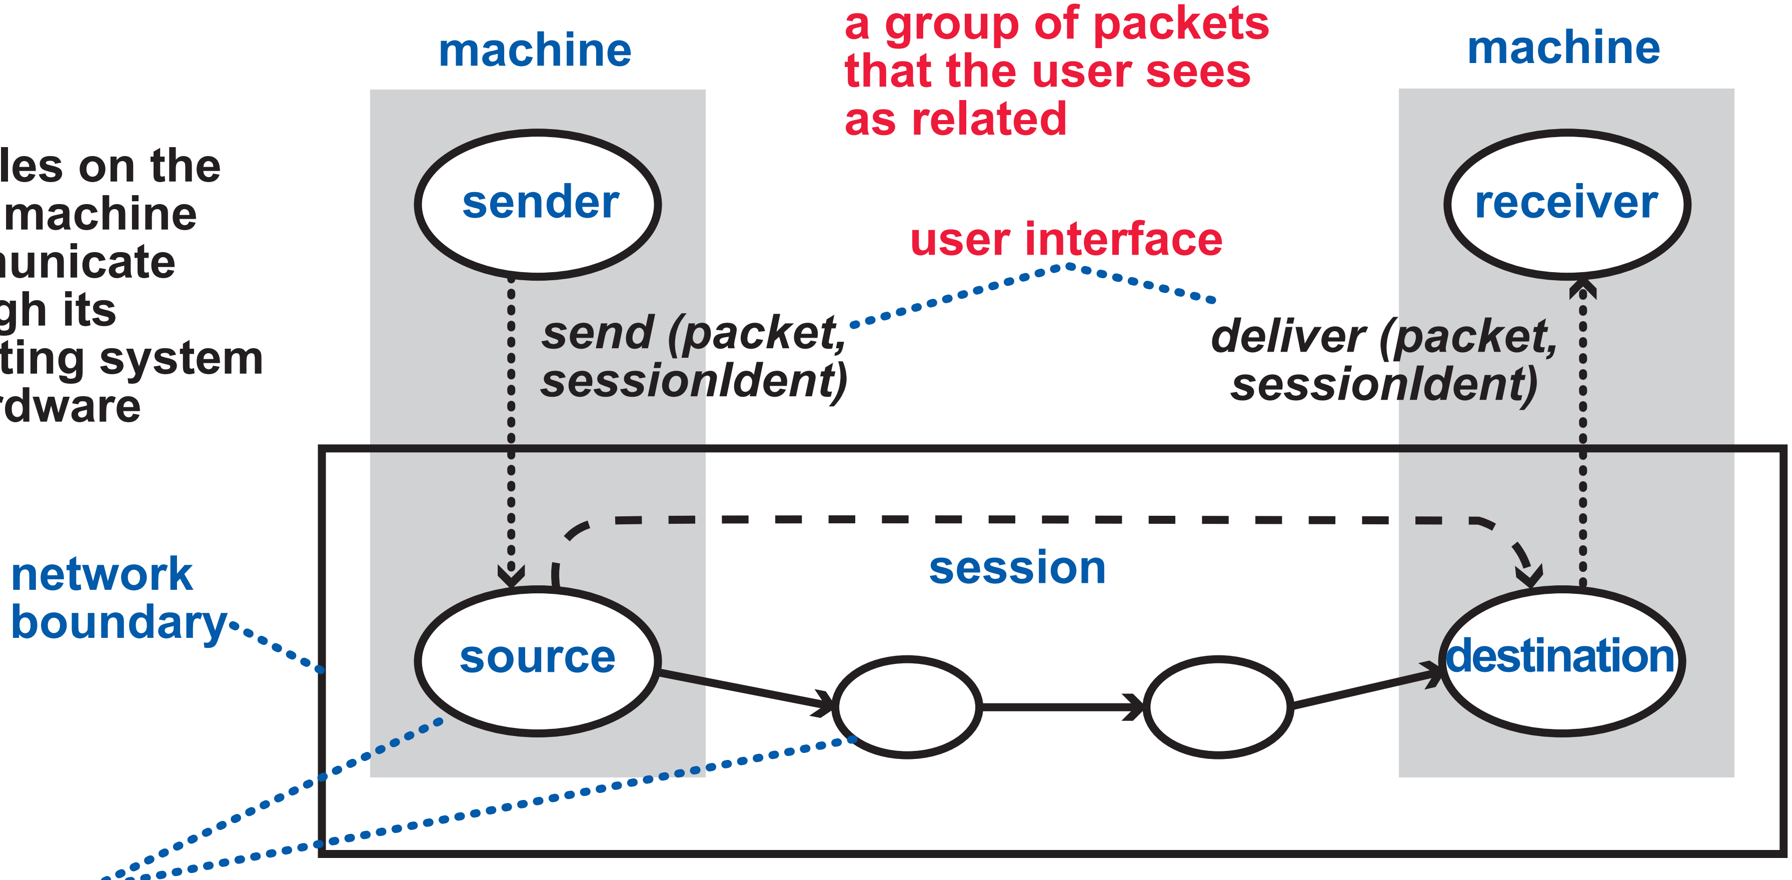

a member of a network is a software or hardware module that implements some of the network protocols

OLD: LAYERS ARE FIXED, HAVE DISTINCT FUNCTIONS

classic Internet architecture has 5 layers, OSI model has the same 5 plus 2 others

routing is the control mechanism that chooses packet paths and encodes paths in forwarding tables

forwarding is the mechanism that pushes packets along their paths

in both reference architectures, there is routing and forwarding only in the link layer (local) and network layer (global)

in this realistic example, there is routing and forwarding in each of the six networks, . . .

. . . with different purposes,

. . . over different spans,

. . . allocating different resources

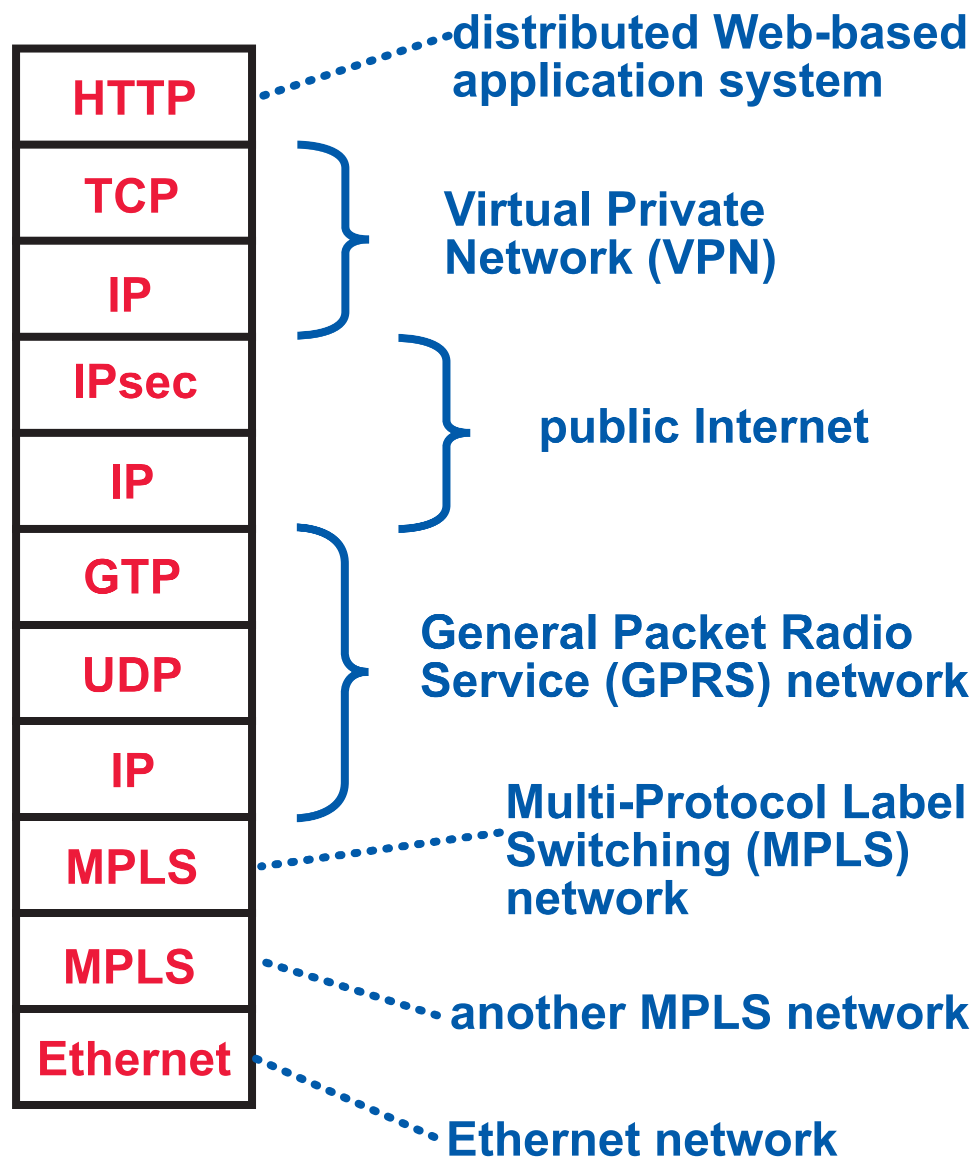

NEW: LAYERS IN A COMPOSITION HIERARCHY ARE SELF-CONTAINED NETWORKS

each network is a microcosm of networking
with all of the basic mechanisms, . . .
. . . all of which can be specialized,
. . . and some of which can be vestigial

members have names from a **namespace**

members are connected by **links** (communication channels)

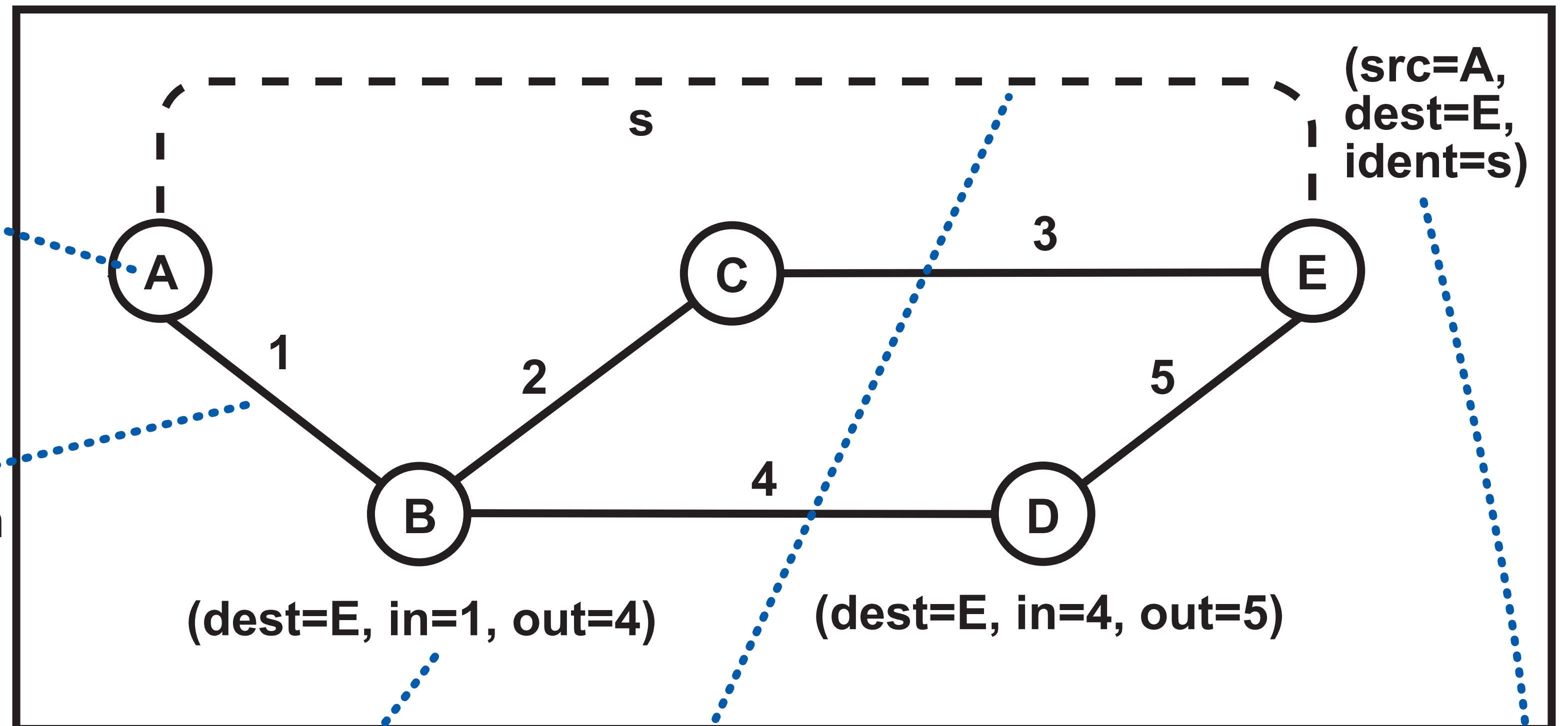

routing chooses packet paths and populates **forwarding tables**, which are used by the **forwarding protocol**

a **session** is an instance of network service

the service is implemented by a **session protocol**, with **session state** in members

REQUIREMENTS ON NETWORKS

The users put a load (of sessions and packets) on the network.

The network delivers communication services with desirable properties.

A network has a single administrative authority . . .

. . . that is responsible for satisfying properties of sessions and session aggregates.

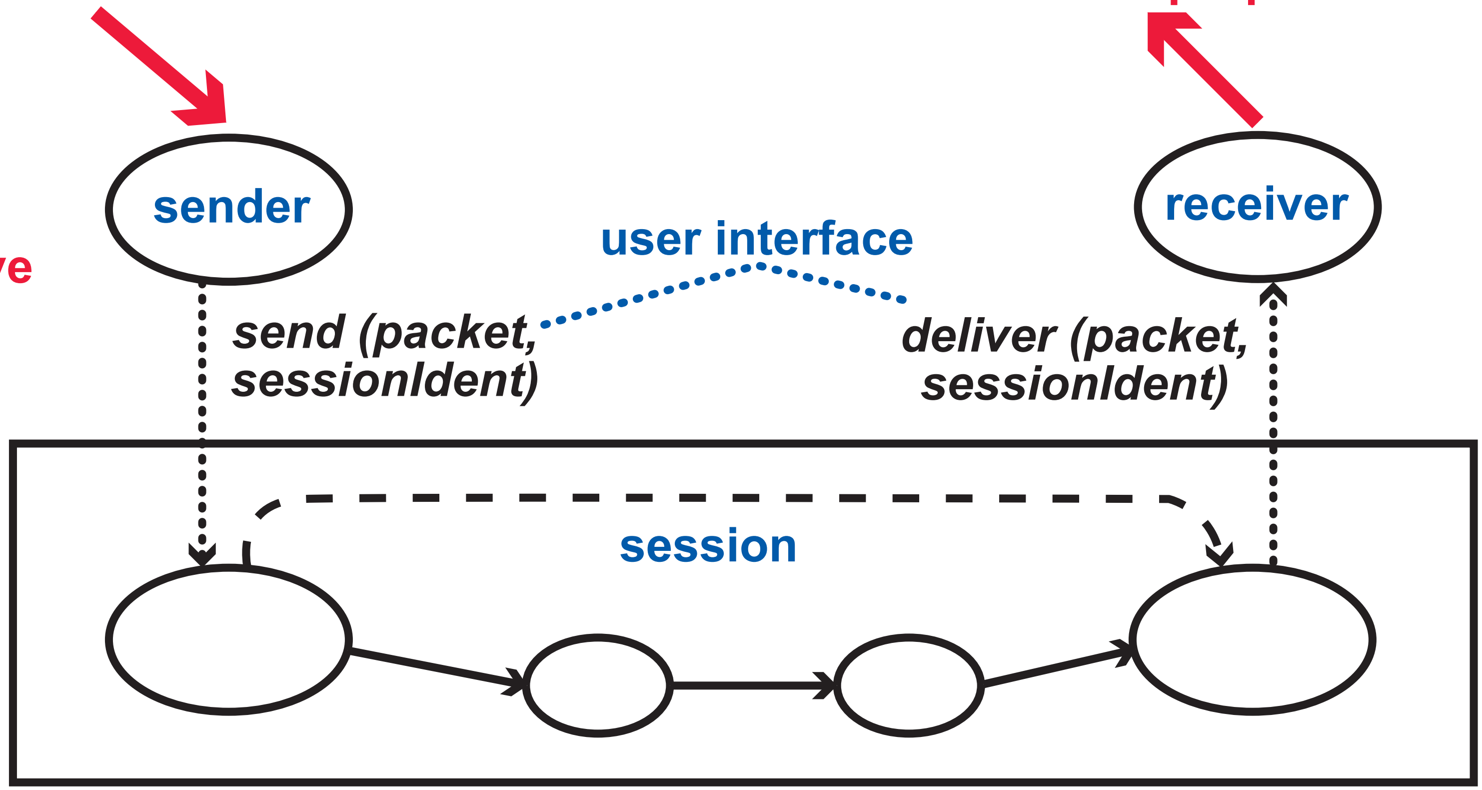

REACHABILITY
- what are the possible destinations?

PERFORMANCE
- maximum latency
- minimum bandwidth
- packet loss rate
- faults tolerated

SERVICE-SPECIFIC BEHAVIOR
- synchronization
- ordered delivery
- guaranteed delivery
- load-balancing
- session persists despite mobility of endpoints

SECURITY
- access control
- DoS protection
- authentication
- privacy
- data integrity
- lawful intercept
- availability

SELF-CONTAINED REASONING ABOUT A NETWORK

reasoning often requires assumptions about the behavior of links

SESSION PERFORMANCE

minimum bandwidth = min $(S_k(B_k))$

links in path

session's share of bandwidth

maximum latency = sum (L_k)

links in path

REACHABILITY

reachability from A is the transitive closure of the forwarding relation

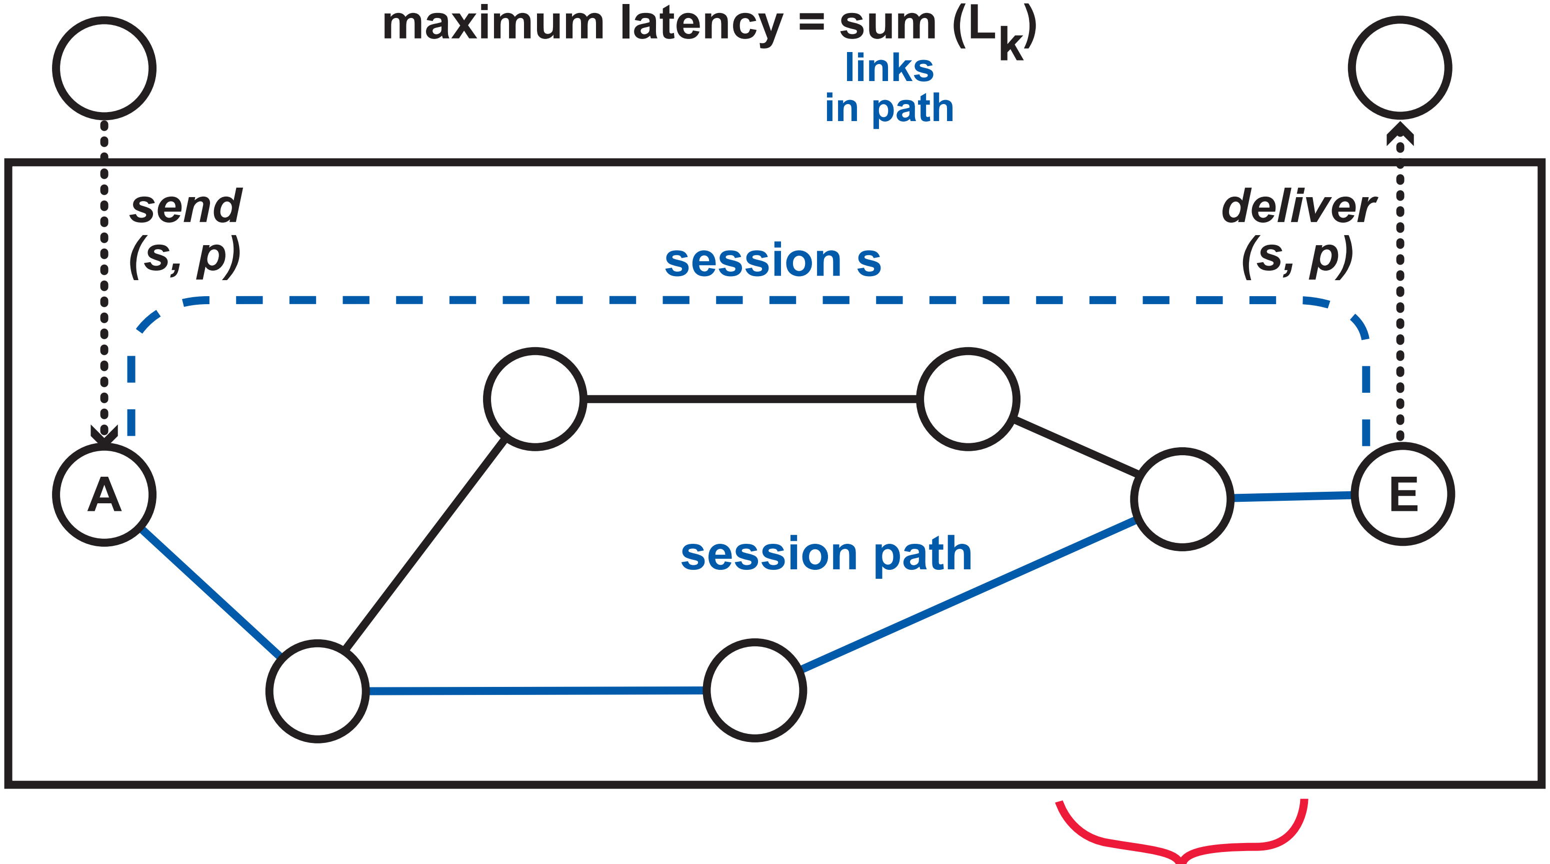

PROTOCOLS

reason about routing protocols, session protocols, etc.

SECURITY

all paths to E go through middleboxes that protect it from DoS attacks and malware

A COMPOSITION OPERATOR: LAYERING

A link in an "overlay" network . . .

. . . is implemented by a session in an "underlay" network.

now a user of a network can be a network instead of a distributed application system

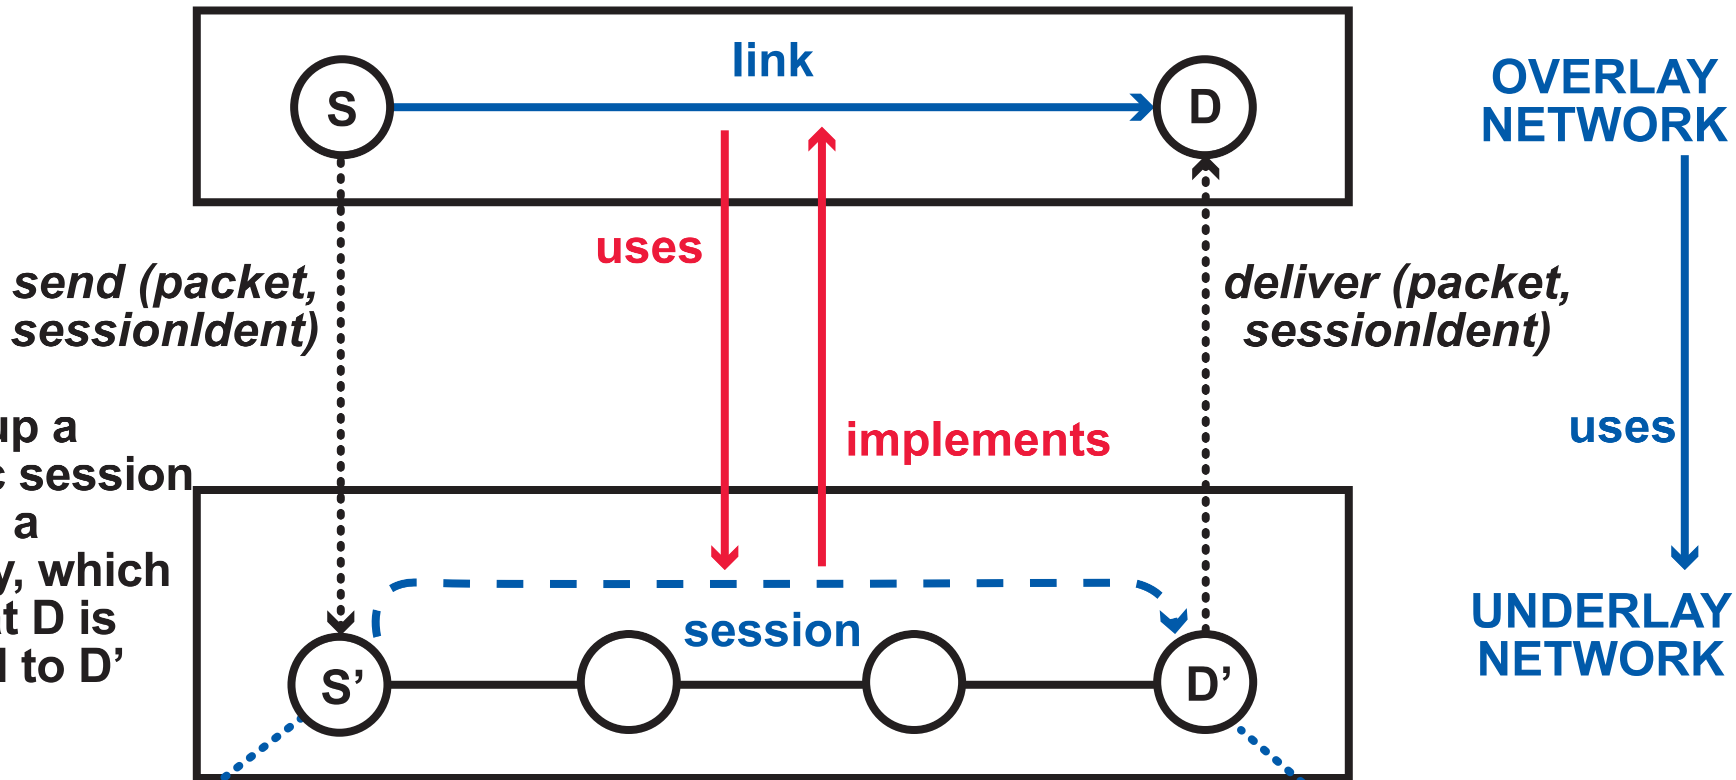

setting up a dynamic session requires a directory, which says that D is attached to D'

dis-assemble packet into smaller packets; encapsulate each in session header

compositional reasoning requires nothing new—the specified properties of the underlay session are simply the assumed properties of the overlay link

decapsulate (remove session headers), assemble into sent packet

LAYERING HAS MANY USES

to build a network with a larger span out of smaller, heterogeneous networks

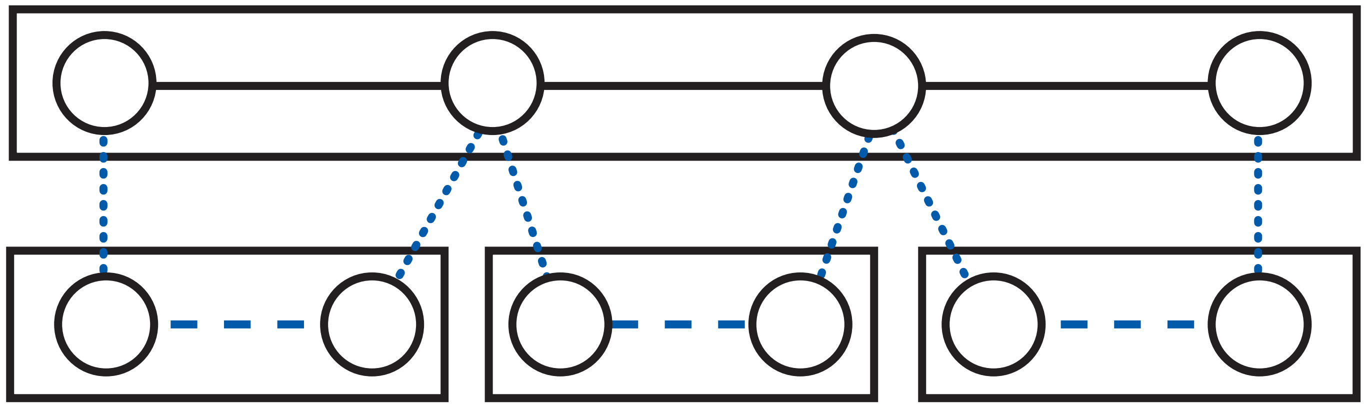

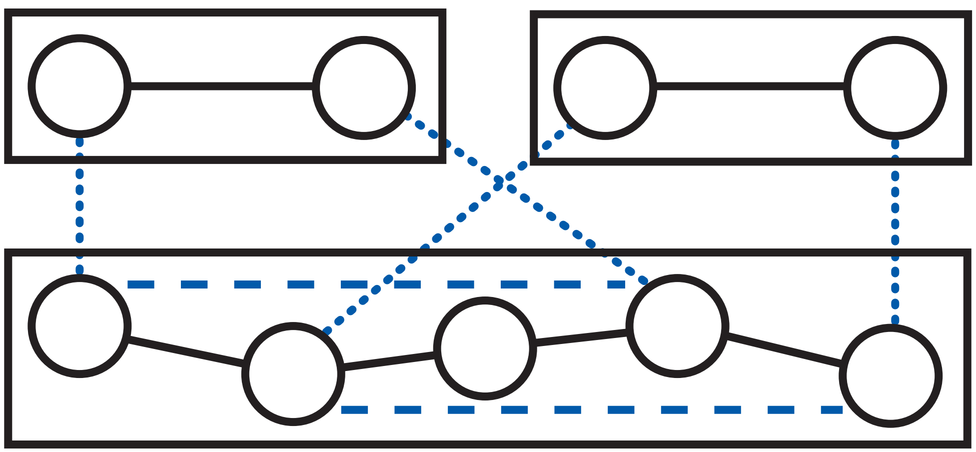

to share the resources of a network in a disciplined way

to build improved communication services on top of an existing network

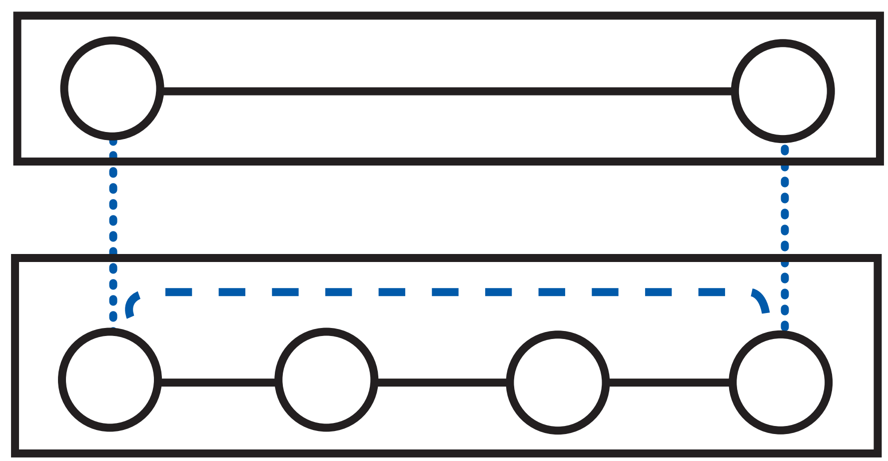

CAMPUS NETWORK WITH VLANS FOR SECURITY

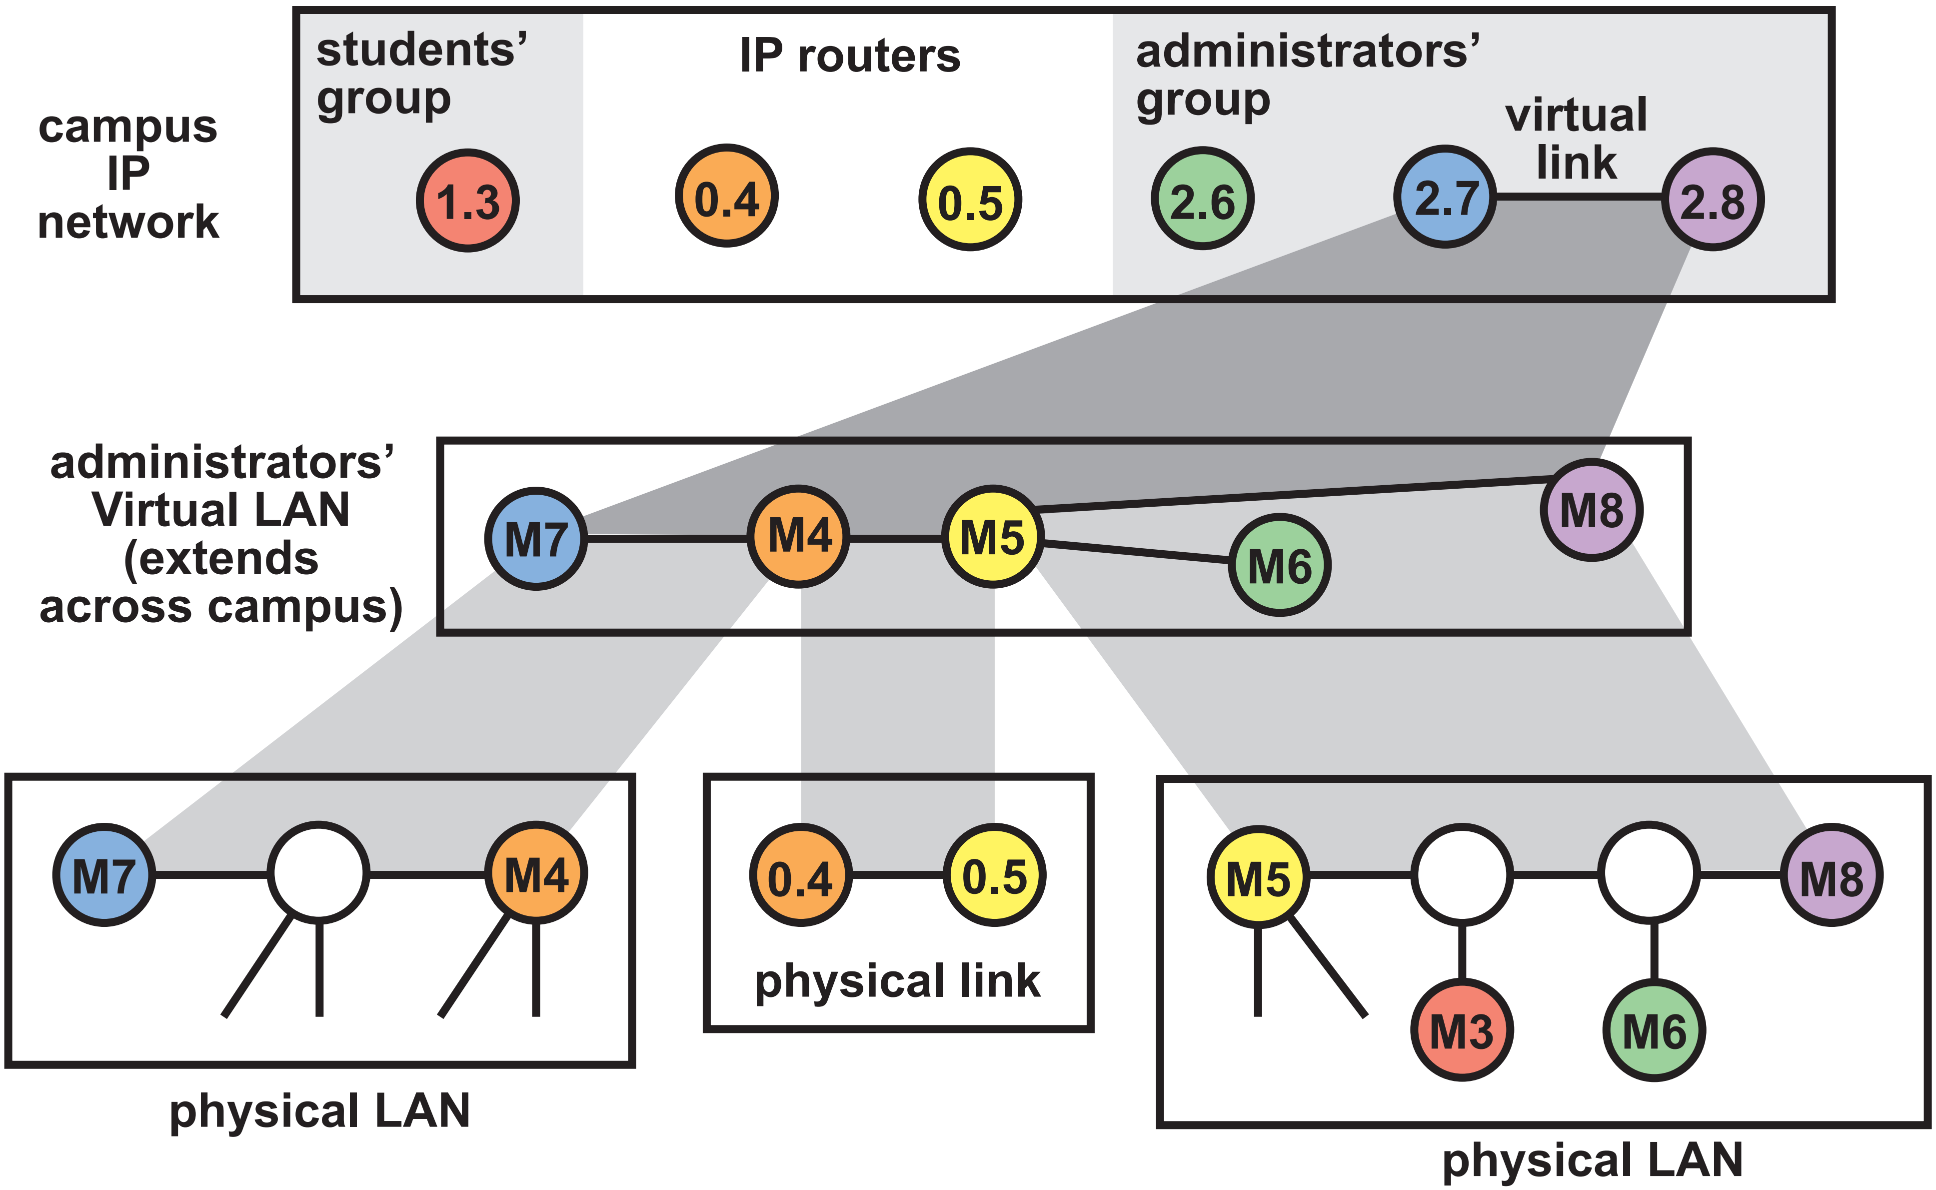

both VLANs and LANs use the Ethernet protocols

for easy configuration, security, and efficiency, each VLAN is isolated

VERIFICATION OF INTER-GROUP SECURITY

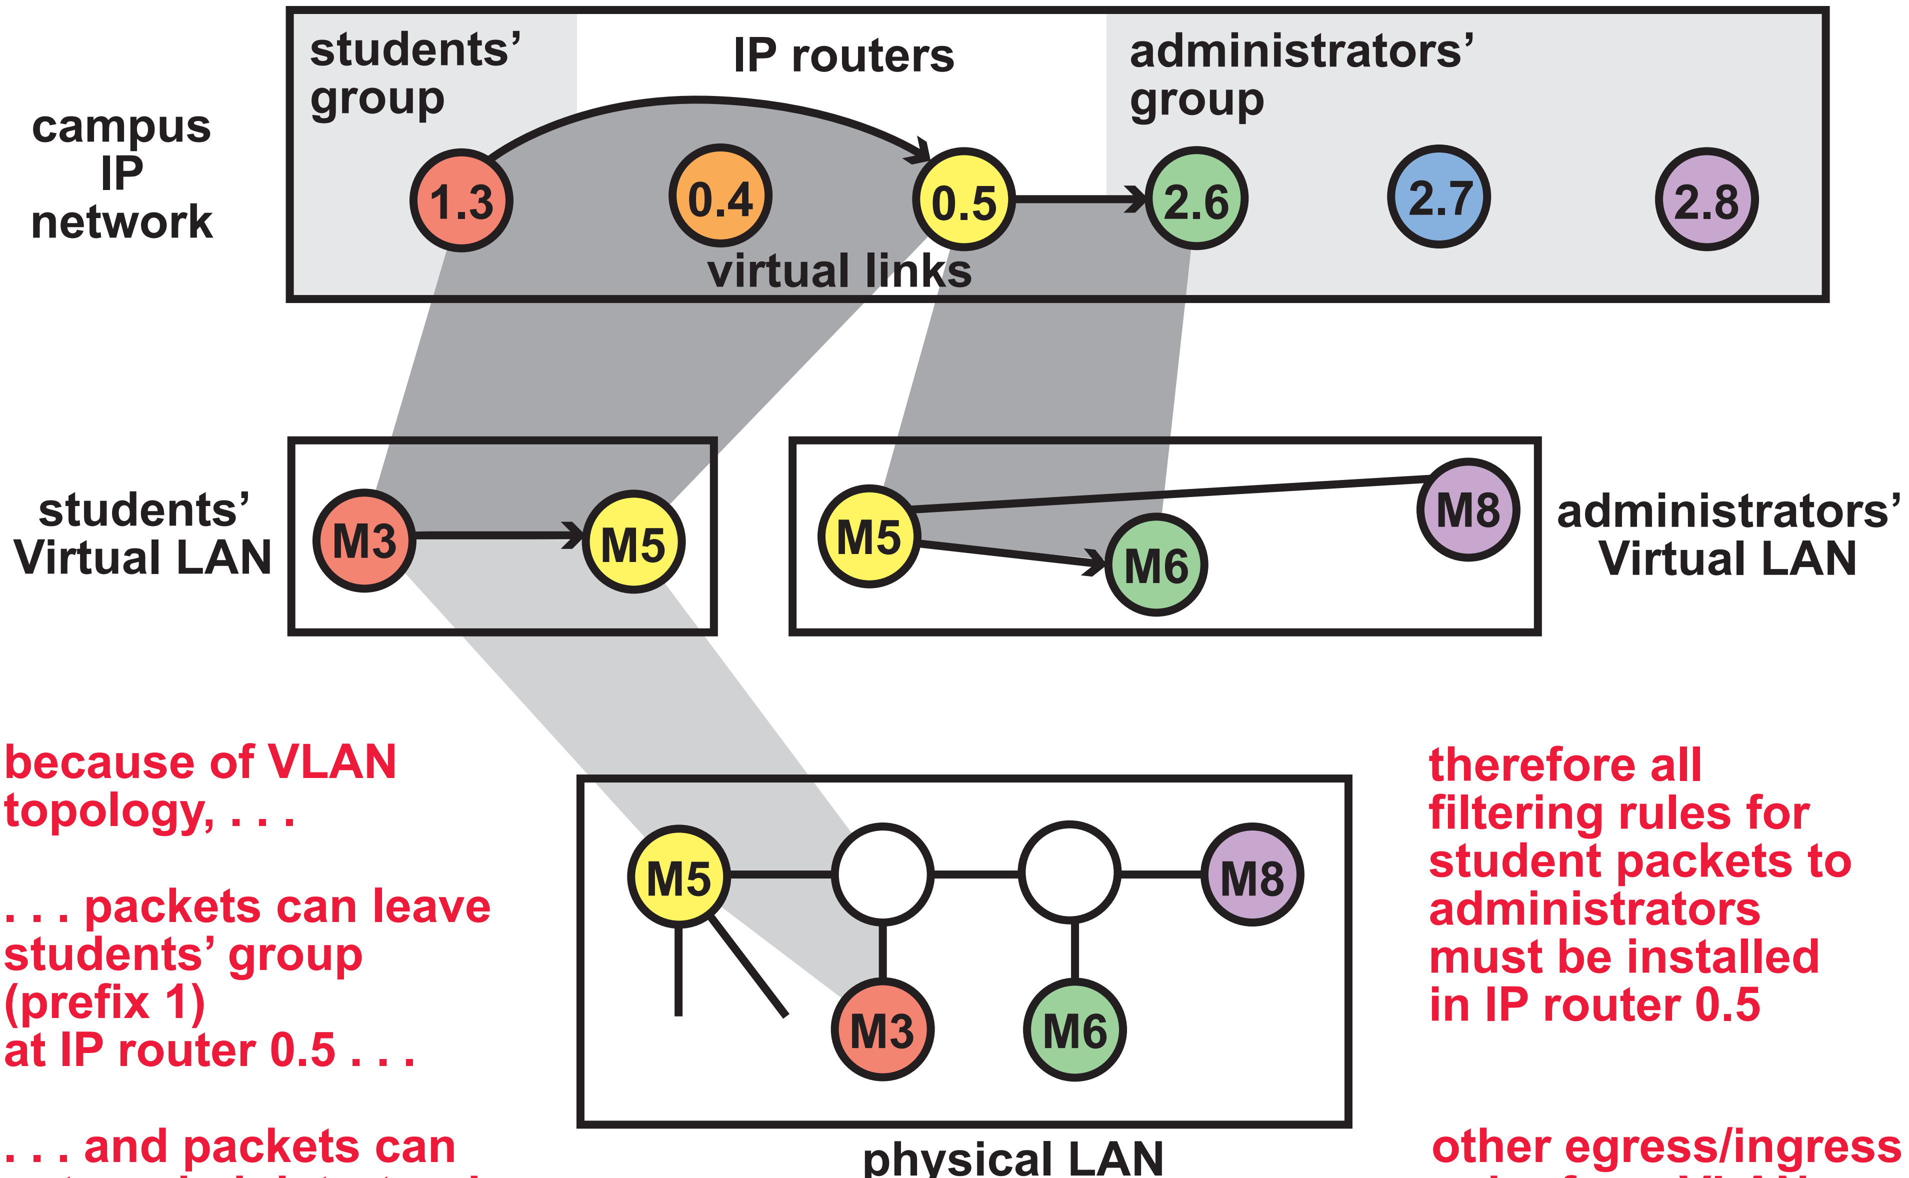

because of VLAN topology, . . .

. . . packets can leave students' group (prefix 1) at IP router 0.5 . . .

. . . and packets can enter administrators' group (prefix 2) at IP router 0.5

therefore all filtering rules for student packets to administrators must be installed in IP router 0.5

other egress/ingress pairs from VLAN topology also require rule installation

THE OTHER COMPOSITION OPERATOR: BRIDGING

bridging allows services to be implemented by networks chained end-to-end

THE EASY WAY

networks have . . .
. . . same namespace
. . . same protocols
. . . globally unique names
. . . access to other network's routing and directories

this is how the networks of the public Internet are composed—they differ only in their administrative authorities

THE HARD WAY

all private IP networks re-use the same name space . . .

. . . so S cannot reach X with a simple session

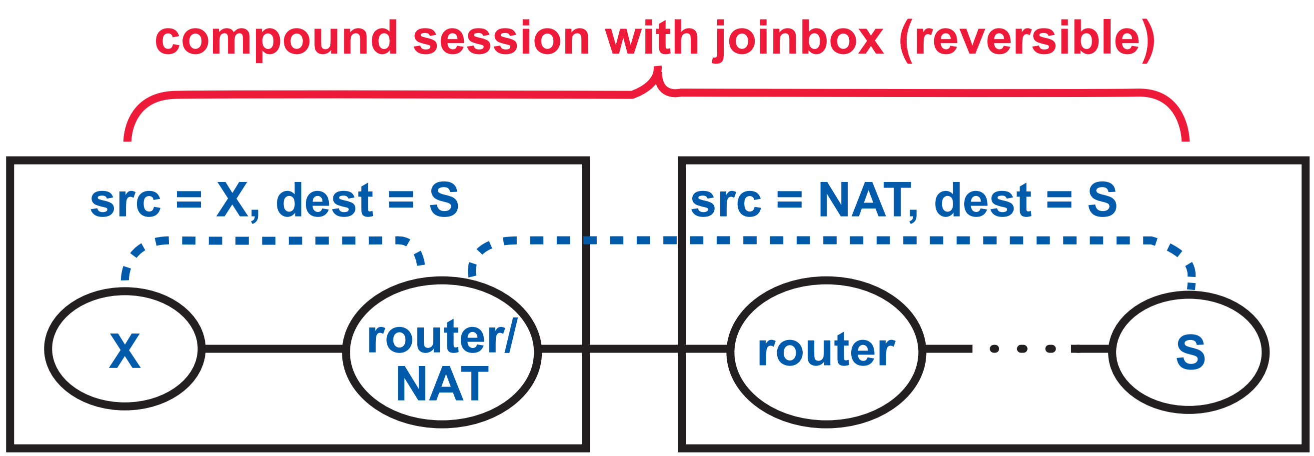

private network in coffee shop

public Internet

other barriers (unshared directories, different protocols) require more powerful joinboxes

A BASIC MODEL OF TRUST

a *member of a network* plays a role in that network; the role is trusted in specific ways

user's laptop is trusted in enterprise network (because it has secret credentials), but not in coffee shop (where is it an anonymous visitor)

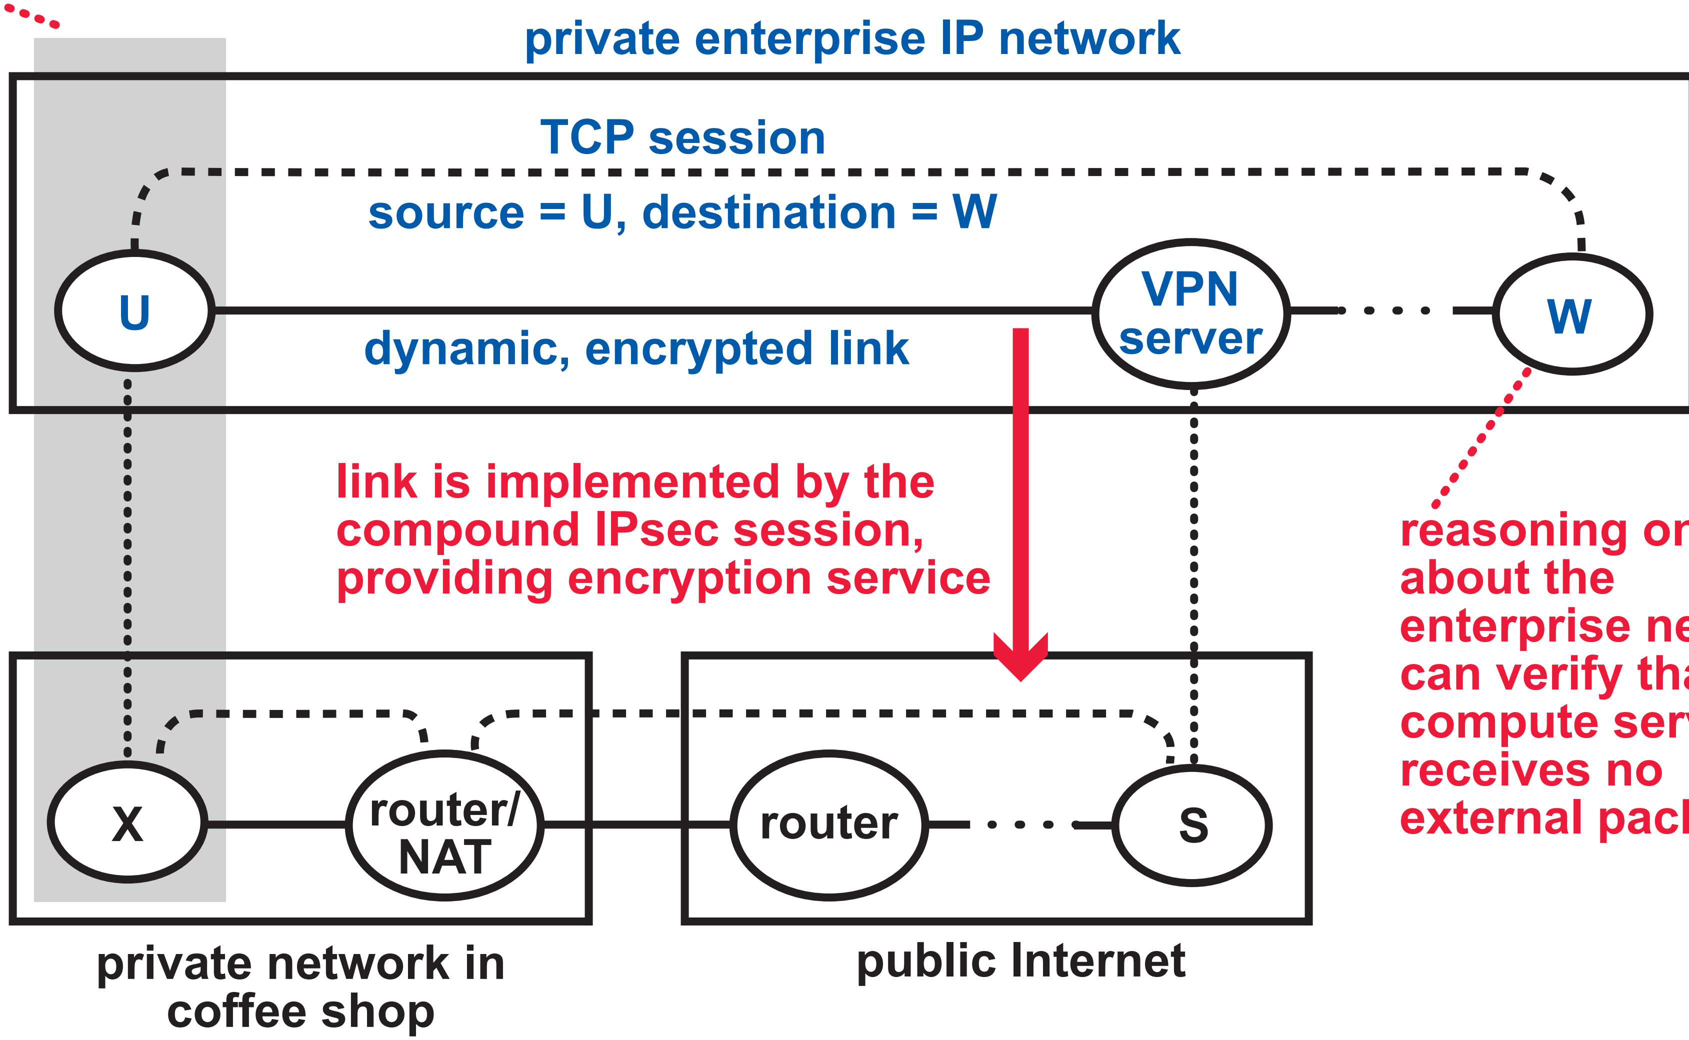

reasoning only about the enterprise network, can verify that the compute server W receives no external packets

MOBILITY WITH LISP Mobile Node

true mobility: a member has a persistent name by which it can be reached at any time, even if it moves during a session

true mobility is difficult to implement in the Internet, because IP addresses are location-dependent

most people get mobile service from cellular networks, which are expensive

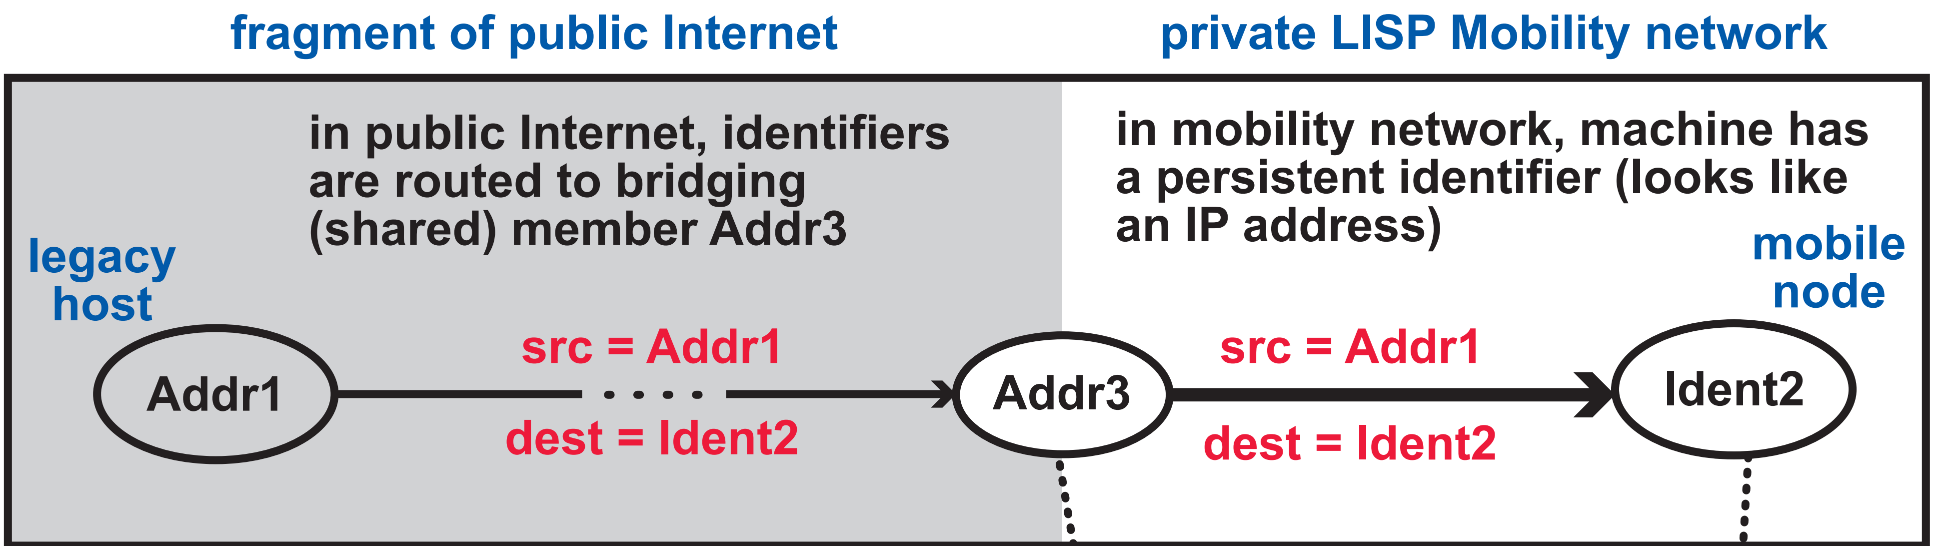

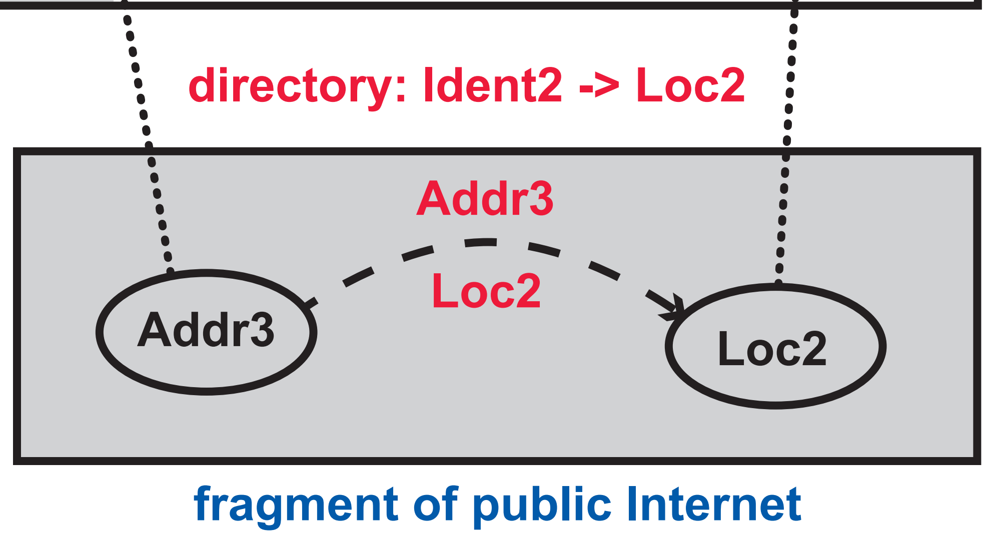

LISP Mobility is a cheaper way to implement mobility

the public Internet and mobility network can interoperate because they are bridged together

the mobility network is also layered on the public Internet!

THIS IS A COMMON PATTERN FOR INTEROPERATION OF SPECIAL NETWORKS WITH THE PUBLIC INTERNET

the “observable Internet” is constructed by bridging

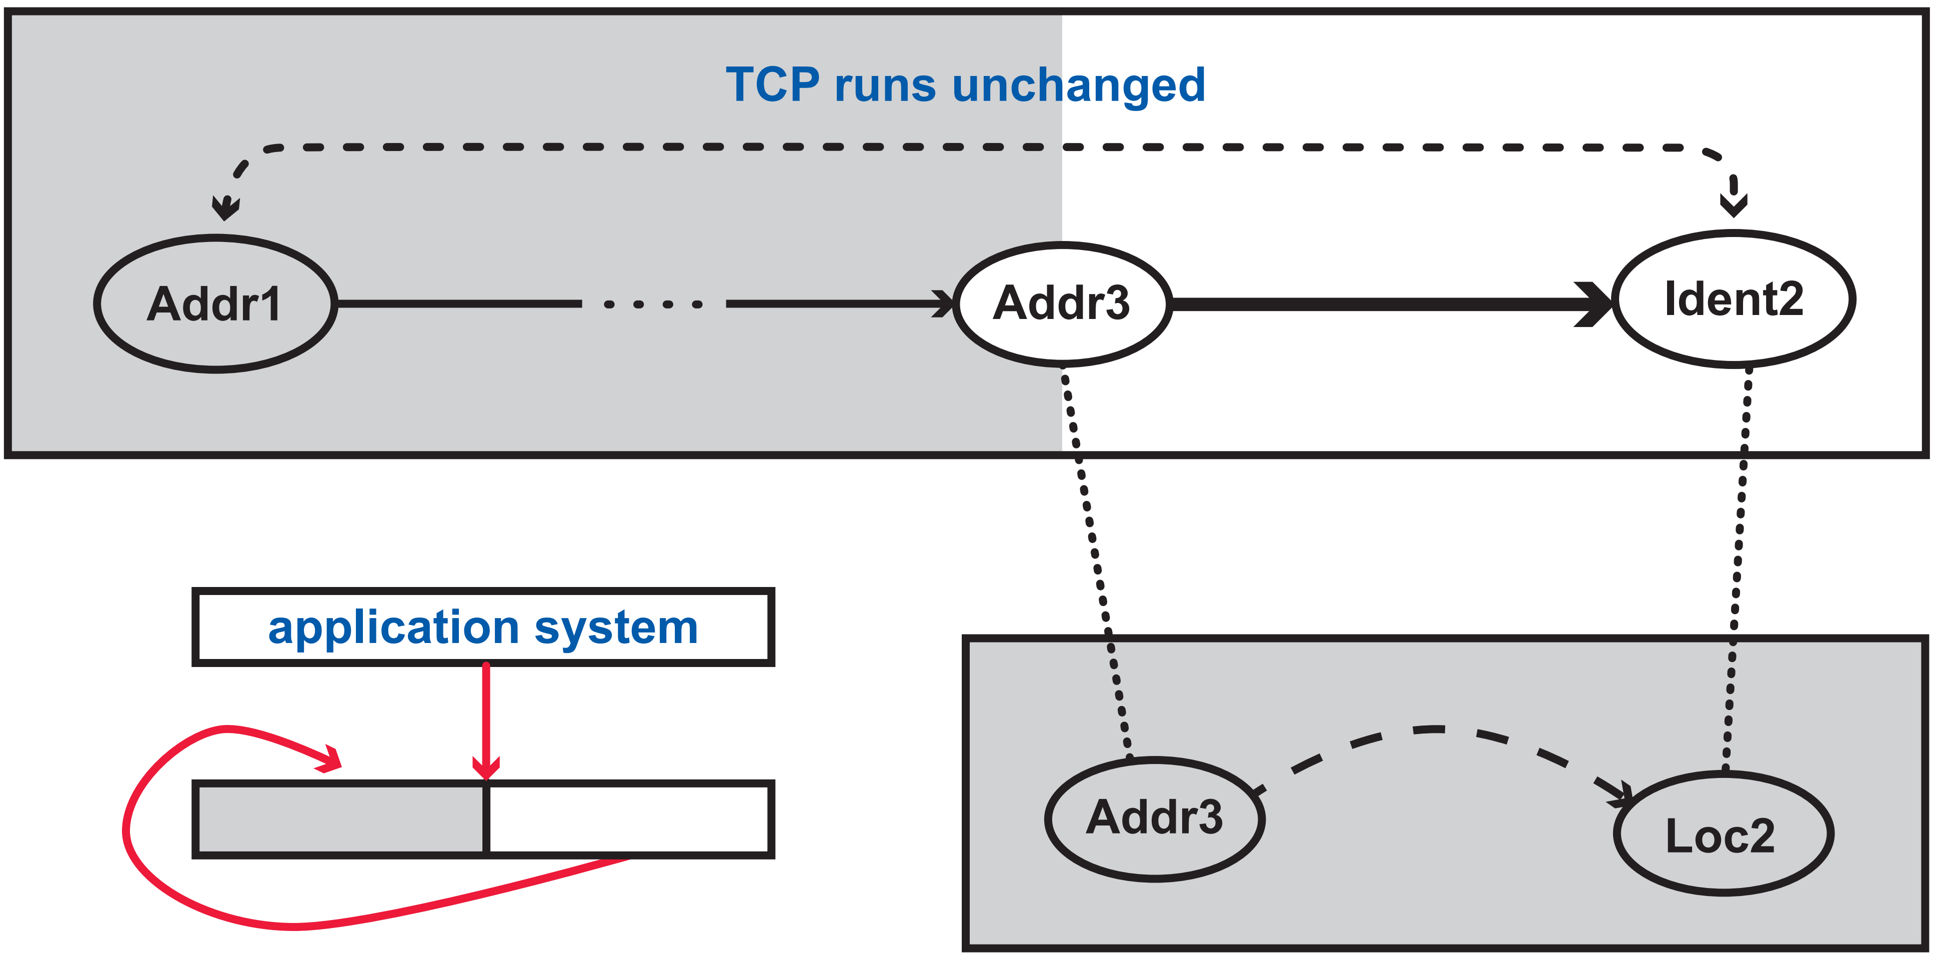

although the “usage hierarchy” of networks sometimes has cycles . . .

. . . a dependency graph of links and paths must not have cycles

EXAMPLES YOU HAVE SEEN

EXAMPLE	WHY IS THERE EXTRA COMPOSITION?	
campus network	need to implement two network structures (one for security, one for physical connectivity) whose topologies are not the same	
enterprise Virtual Private Network	need a network of which an employee’s laptop can be a trusted member with a private access link, even when it is located in a public place on an untrusted network	
LISP Mobility	device mobility	need to add to the public Internet a capability that is intrinsically difficult to implement with the Internet’s native architecture

. . . PLUS MANY OTHERS

SHOW THAT THE NEW COMPOSITIONAL MODEL IS VALID

EXAMPLE	WHY IS THERE EXTRA COMPOSITION?	
campus network	need to implement two network structures (one for security, one for physical connectivity) whose topologies are not the same	
enterprise Virtual Private Network	need a network in which an employee’s laptop can be a trusted member with a private access link, even when it is located in a public place on an untrusted network	
LISP Mobility	device mobility	need to add to the public Internet a capability that is intrinsically difficult to implement with the Internet’s native architecture
Secure Overlay Services	security for a small club that excludes all others	
Named Data Networking	need to experiment with a new architecture—nothing like the Internet—designed for content distribution	
VL2	use of all available bandwidth and switch capacity	need to implement cloud services efficiently
SIMPLE	insertion of functional middleboxes into end-to-end paths	

NOW WE NEED FORMAL THEORY FOR THE COMPOSITIONAL MODEL, IN SUPPORT OF . . .

RE-USE OF SOLUTION PATTERNS

- emphasizing how networks are similar (even when they have different purposes) leads to recognition of re-usable solution patterns
- once a pattern is identified, all artifacts (e.g., code, proofs) can potentially be re-used or generated

EFFECTIVE OPTIMIZATIONS

- the important optimizations move functions up (virtualization) or down (hardware acceleration) in the composition hierarchy
- need compositional reasoning to optimize in the best way
- need automated transformations for safe optimization

INTERNET INTEROPERATION AND EVOLUTION

- composition allows the Internet to interoperate with new concepts and then evolve toward them
- with recognition of this reality, the process can be made easier

VERIFICATION OF TRUSTWORTHY SERVICES

- there is increasing demand for trustworthy services
- composition is so ubiquitous that service verification is impossible without compositional reasoning

A NEW INTERNET STORY

OLD

There is a single Internet (not counting administrative boundaries) which cannot be replaced.

Because it does not meet all current and projected requirements, we must seek to add a never-ending list of new features to it.

Because its complexity is growing continually, we must work ever harder to find ways to secure and verify it.

NEW

The Internet will continue to evolve by means of new networks and new compositions.

These are easy to add . . .

. . . locally (campus networks, cloud computing) . . .

. . . or at high levels of the composition hierarchy (mobility, distributed systems), . . .

. . . and slower to disseminate when both global and low in the composition hierarchy (IPv6).

By studying and emphasizing composition, we can make evolution faster, easier, and better.

LESSONS FOR US

FORMAL METHODS

- research in formal methods is not just tool development

the model really matters,

and the right model

is not always obvious

NETWORKING

- networks used to be dominated by hardware—now, like all other complex systems, networks are software systems
- networking today is overwhelmed by complexity, . . .

  . . . and network researchers/practioners have few solutions to this problem
- the boundary between networks and distributed systems is artificial

OPPORTUNITY

- networking is now essential to civilization . . .

  . . . and the technology must mature to become trustworthy
- with the right model for managing complexity and focusing attention on real issues, . . .

  . . . formal methods can make a big contribution to this progress