Principles for Measurability in Protocol Design

From:
ACM SIGCOMM
CCR Vol. 47 Issue 2

Mark Allman (ICSI)
***Robert Beverly** (NPS)*
Brian Trammell (ETHZ)

August 24, 2017

ETH zürich

Principles for Mea… in Prot…

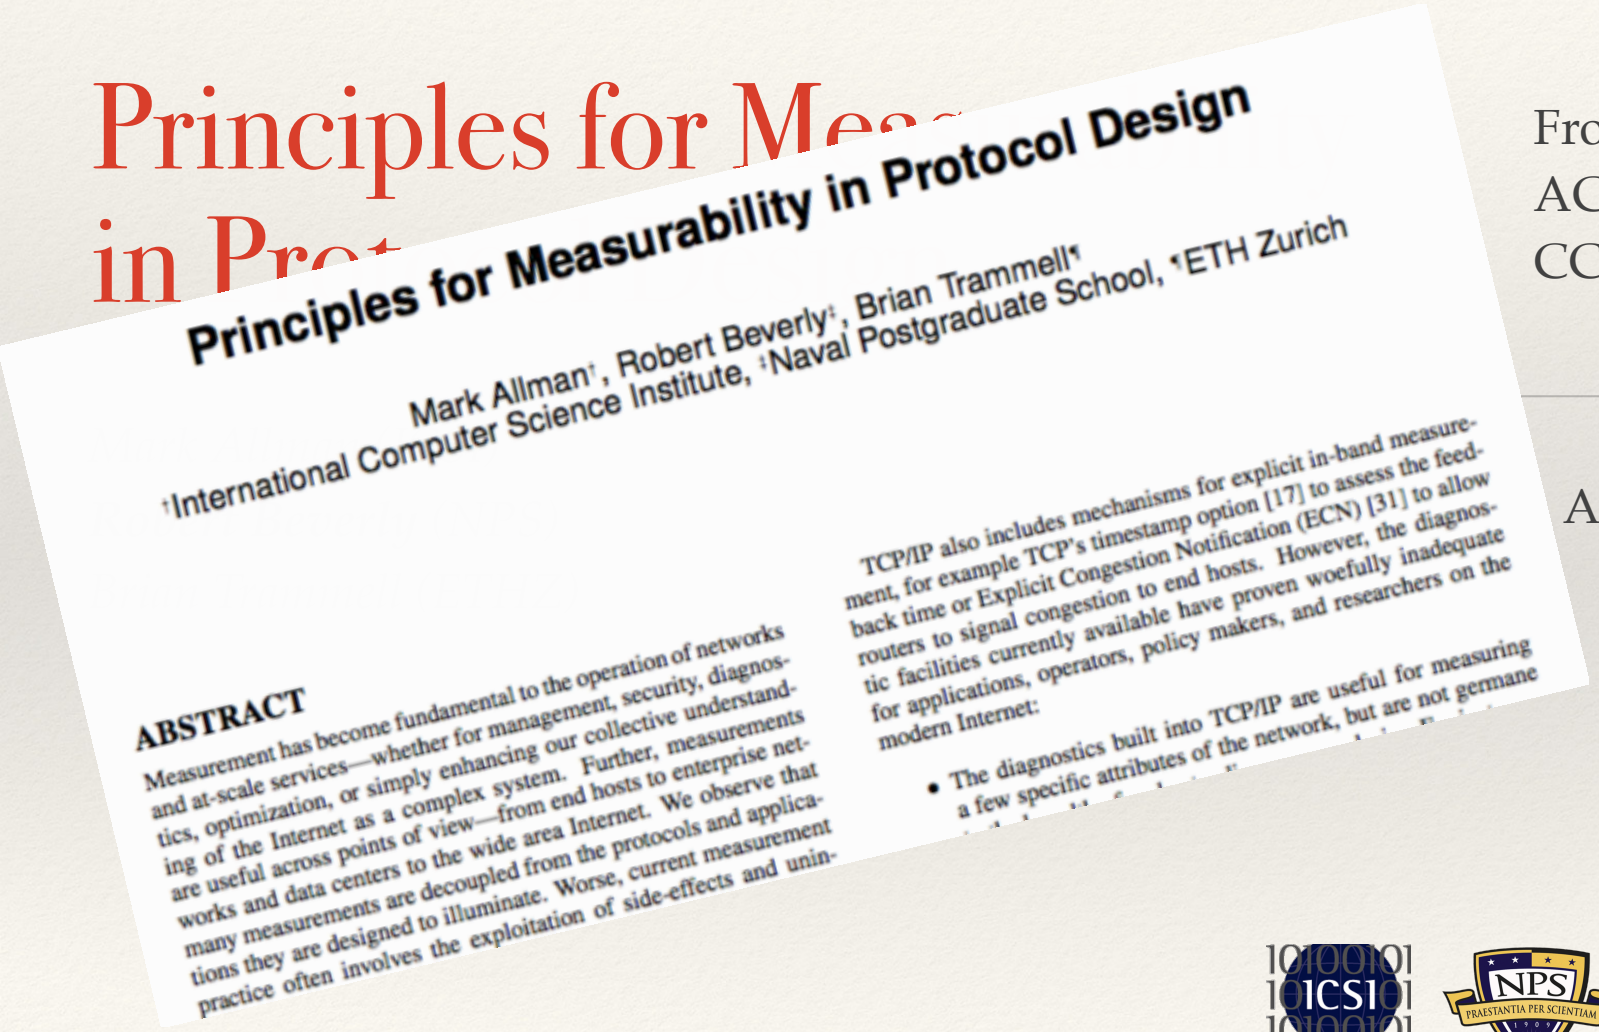

Principles for Measurability in Protocol Design

Mark Allman[†], Robert Beverly[‡], Brian Trammell[¶]
[†]International Computer Science Institute, [‡]Naval Postgraduate School, [¶]ETH Zurich

ABSTRACT

Measurement has become fundamental to the operation of networks and at-scale services—whether for management, security, diagnostics, optimization, or simply enhancing our collective understanding of the Internet as a complex system. Further, measurements are useful across points of view—from end hosts to enterprise networks and data centers to the wide area Internet. We observe that many measurements are decoupled from the protocols and applications they are designed to illuminate. Worse, current measurement practice often involves the exploitation of side-effects and unin-

TCP/IP also includes mechanisms for explicit in-band measurement, for example TCP's timestamp option [17] to assess the feedback time or Explicit Congestion Notification (ECN) [31] to allow routers to signal congestion to end hosts. However, the diagnostic facilities currently available have proven woefully inadequate for applications, operators, policy makers, and researchers on the modern Internet:

- The diagnostics built into TCP/IP are useful for measuring a few specific attributes of the network, but are not germane

From:
ACM SIGCOMM
CCR Vol. 47 Issue 2

August 24, 2017

Network Measurement

- Fundamental to network operation, application performance, and policy (not just research)
- But, today:
  - Minimal support from stack
  - Measurements rely on brittle hacks, unintended features, and inferences

Result: Important questions are hard

- E.g.:
  - What's the best path to route traffic?
  - What is the capacity or utilization of a link?
  - How do networks interconnect?
  - What AS operates a given router?

Even simple inferences difficult!

- What's the delay between two hosts

Even simple inferences difficult!

- What's the delay between two hosts
  - (Per-protocol traffic differentiation, path vs. host delay, asymmetry)

Even simple inferences difficult!

- What's the delay between two hosts
  - (Per-protocol traffic differentiation, path vs. host delay, asymmetry)
- What are the endpoints in a communication?

Even simple inferences difficult!

- What's the delay between two hosts
  - (Per-protocol traffic differentiation, path vs. host delay, asymmetry)
- What are the endpoints in a communication?
  - (NATs, CGNs, aliases, IPv6)

Even simple inferences difficult!

- What's the delay between two hosts
  - (Per-protocol traffic differentiation, path vs. host delay, asymmetry)
- What are the endpoints in a communication?
  - (NATs, CGNs, aliases, IPv6)
- How did packets arrive at a remote destination?

Even simple inferences difficult!

- What's the delay between two hosts
  - (Per-protocol traffic differentiation, path vs. host delay, asymmetry)
- What are the endpoints in a communication?
  - (NATs, CGNs, aliases, IPv6)
- How did packets arrive at a remote destination?
  - (order? modified? mangled? path? queued?)

Reconsidering Measurability

What if we re-think the stack with **measurability** as a **first-class** component?

Principles for Measurability

P1. Explicit

P2. In-band

P3. Consumer bears cost

P4. Provider retains control

P5. Visible

P6. Cooperative

Principles for Measurability

P1. Explicit

Remove ambiguity
Transparency encourages adoption

P2. In-band

P3. Consumer bears cost

P4. Provider retains control

P5. Visible

P6. Cooperative

Principles for Measurability

P1. Explicit

P2. In-band

Faithfully capture measurement objective

P3. Consumer bears cost

P4. Provider retains control

P5. Visible

P6. Cooperative

Principles for Measurability

P1. Explicit

P2. In-band

P3. Consumer bears cost

Measurement burden on consumer, not producer

P4. Provider retains control

P5. Visible

P6. Cooperative

Principles for Measurability

P1. Explicit

P2. In-band

P3. Consumer bears cost

P4. Provider retains control

Measurement producers can make conscious decisions what to expose

P5. Visible

P6. Cooperative

Principles for Measurability

P1. Explicit

P2. In-band

P3. Consumer bears cost

P4. Provider retains control

P5. Visible

P6. Cooperative

Measurements require visibility into forward / reverse paths and packet modification

Principles for Measurability

P1. Explicit

P2. In-band

P3. Consumer bears cost

P4. Provider retains control

P5. Visible

P6. Cooperative

Measurements must cooperate with routers, middleboxes, and infrastructure

Primitives

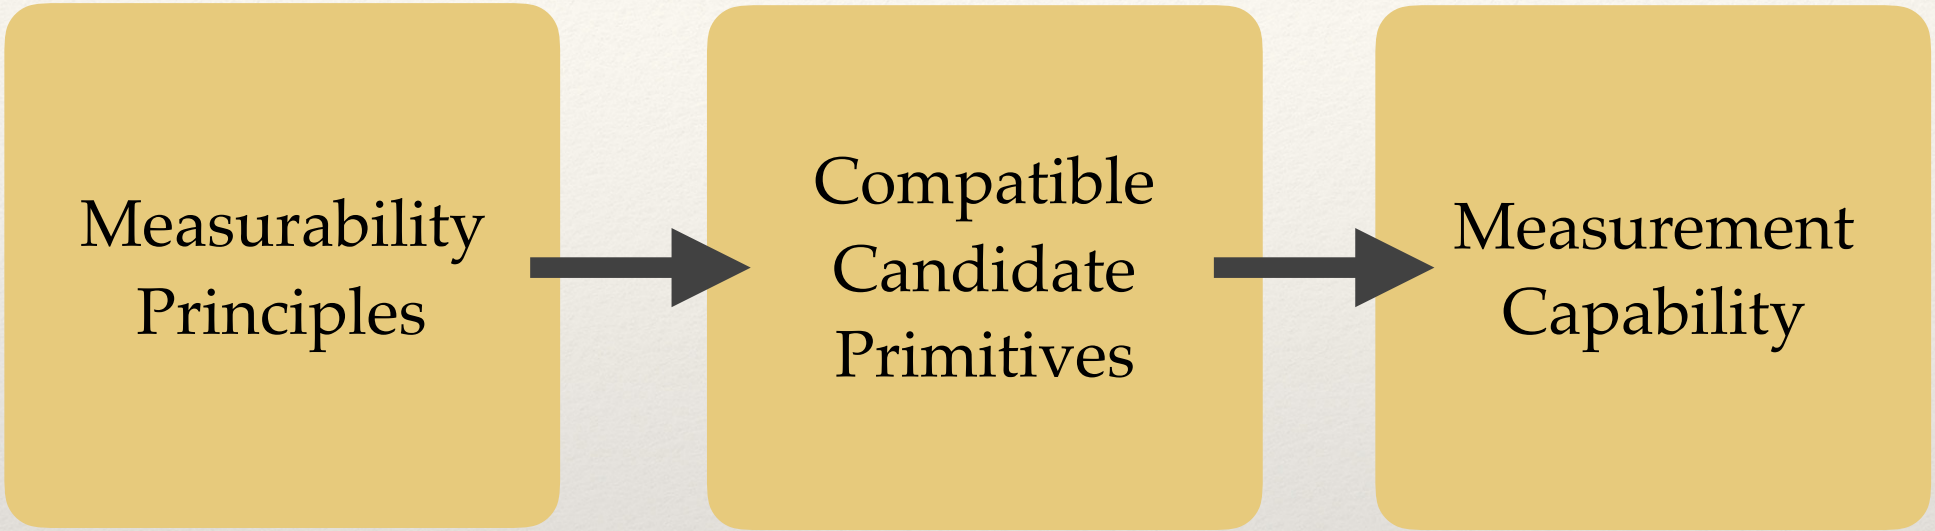
Measurability Principles
Compatible Candidate Primitives
Measurement Capability

Primitives

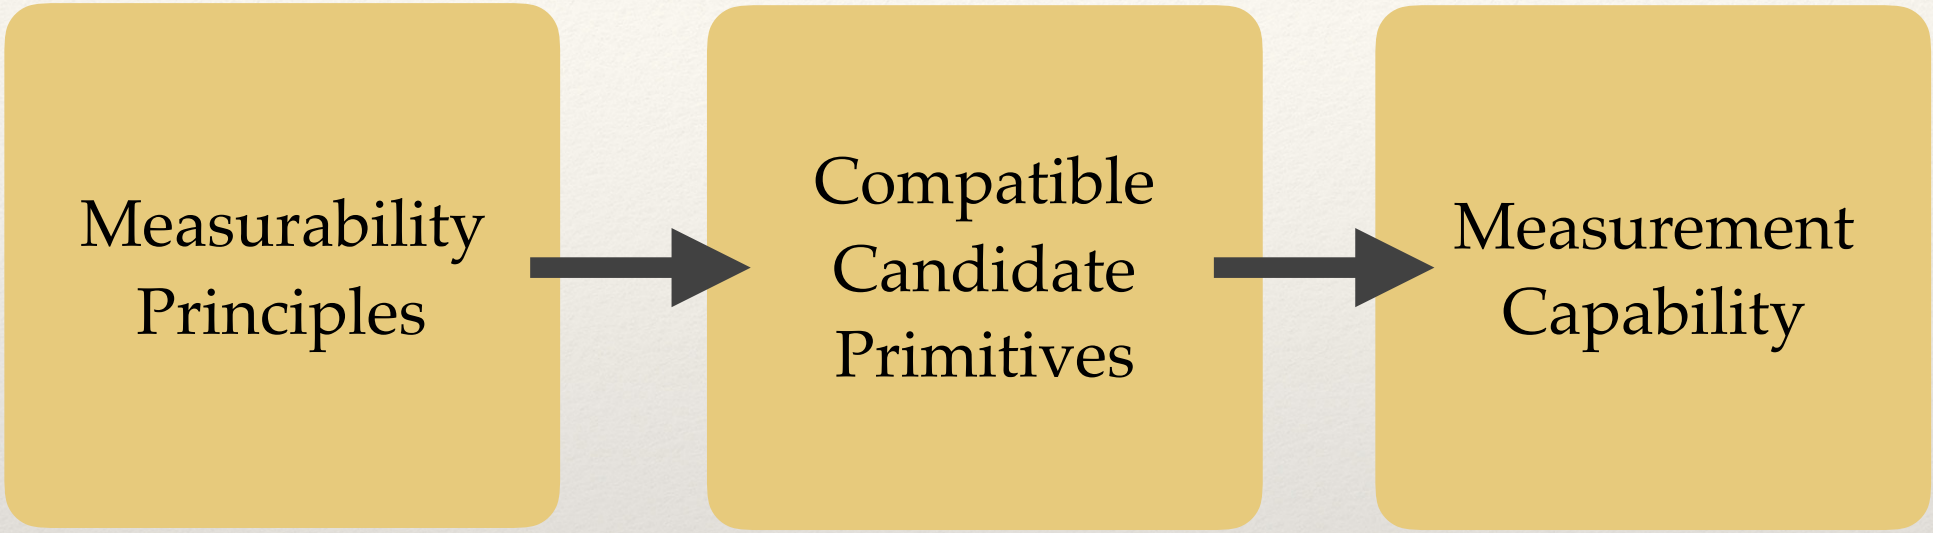

Imagine packets carry measurement meta-data.

What should that meta-data include?

Candidate Primitive: HostID

- Host ID:
  - Chosen randomly, included in packets
  - Removes IP address = host assumption
  - Remove NAT, load-balancer, IPv6, alias ambiguities that plague today's measurements

Candidate Primitive: HostID

- Host ID:
  - Chosen randomly, included in packets
  - Removes IP address = host assumption
  - Remove NAT, load-balancer, IPv6, alias ambiguities that plague today's measurements

How to apply principles to make HostID viable?

Candidate Primitive: HostID

- Host ID:
  - ID is ephemeral
  - Small ID space + change ID to prevent tracking
  - Large population requires observation over time, probabilistic inferences

Candidate Primitive: HostID

- Host ID:

P1. Explicit

  - ID is ephemeral
  - Small ID space + change ID to prevent tracking
  - Large population requires observation over time, probabilistic inferences

Candidate Primitive: HostID

- Host ID:
  - ID is ephemeral
  - Small ID space + change ID to prevent tracking
  - Large population requires observation over time, probabilistic inferences

P1. Explicit

P4. Provider retains control

Candidate Primitive: HostID

- Host ID:
  - ID is ephemeral
  - Small ID space + change ID to prevent tracking
  - Large population requires observation over time, probabilistic inferences

P1. Explicit

P4. Provider retains control

P3. Consumer bears cost

HostID: Active Inference

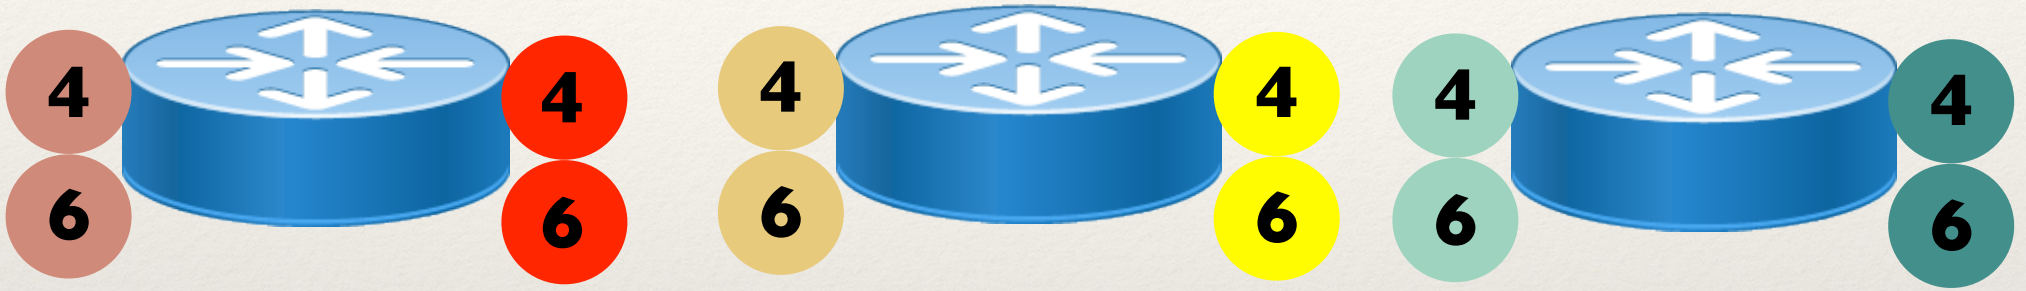

HostID: Active Inference

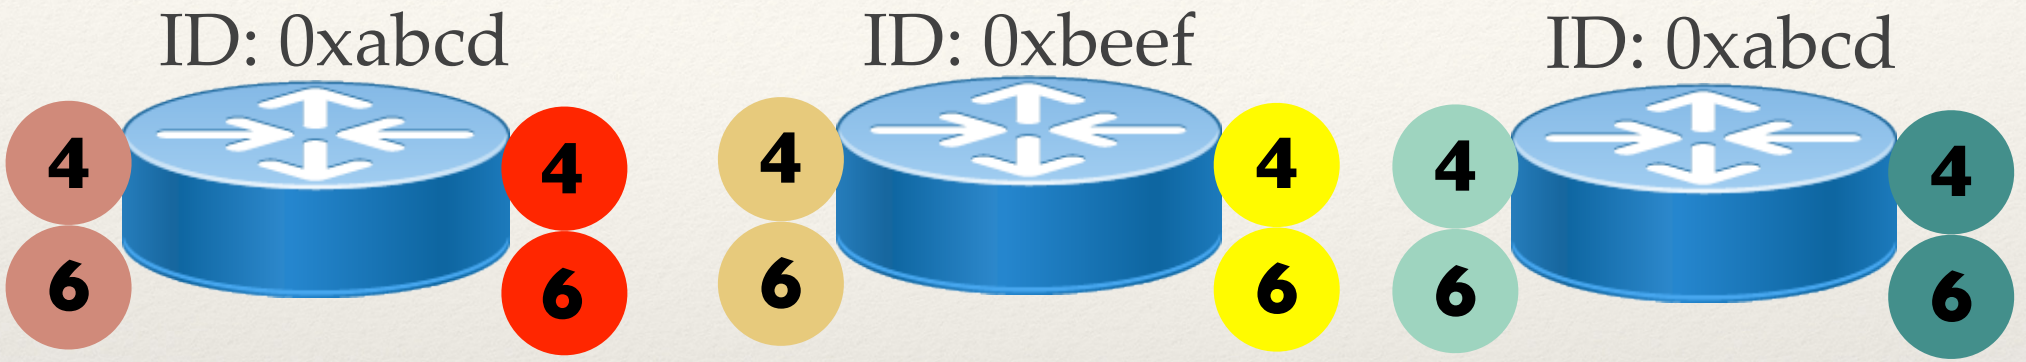

HostID: Active Inference

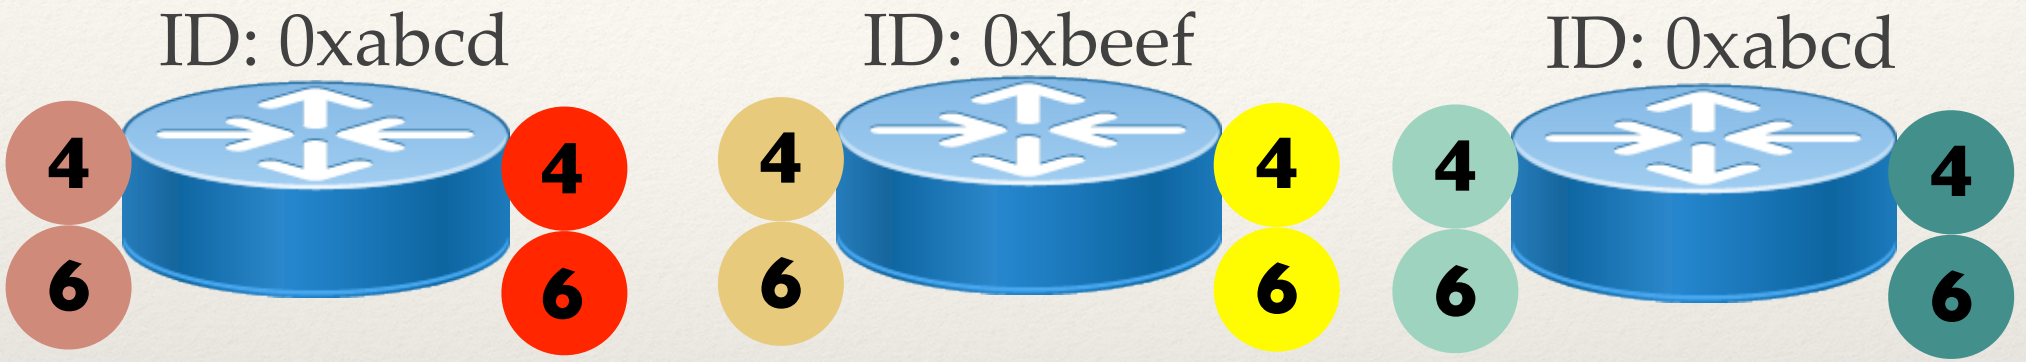

Router 1 Aliases:

Router 2 Aliases:

HostID: Active Inference

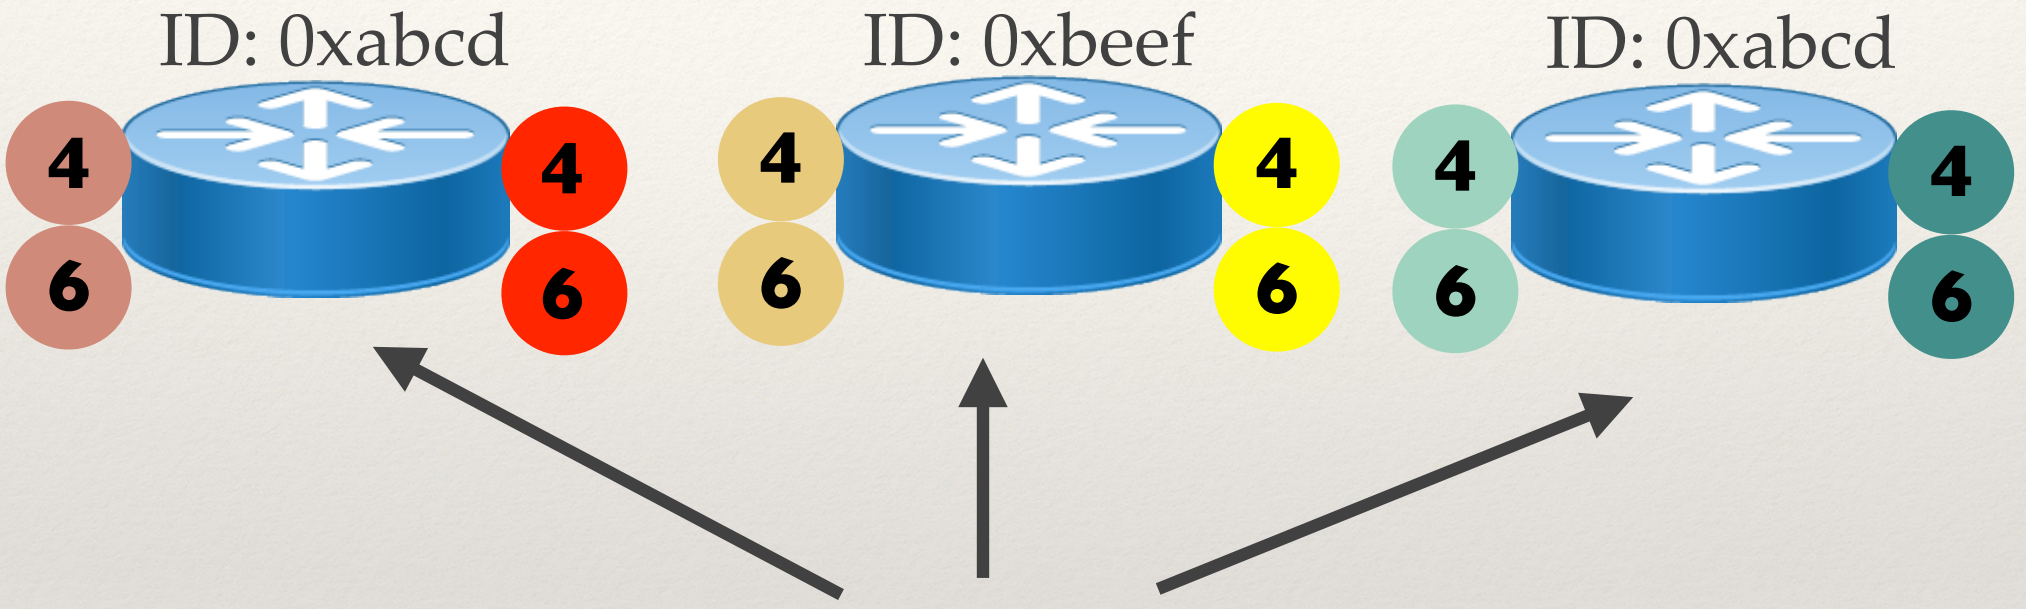

Router 1 Aliases:

Router 2 Aliases:

HostID: Active Inference

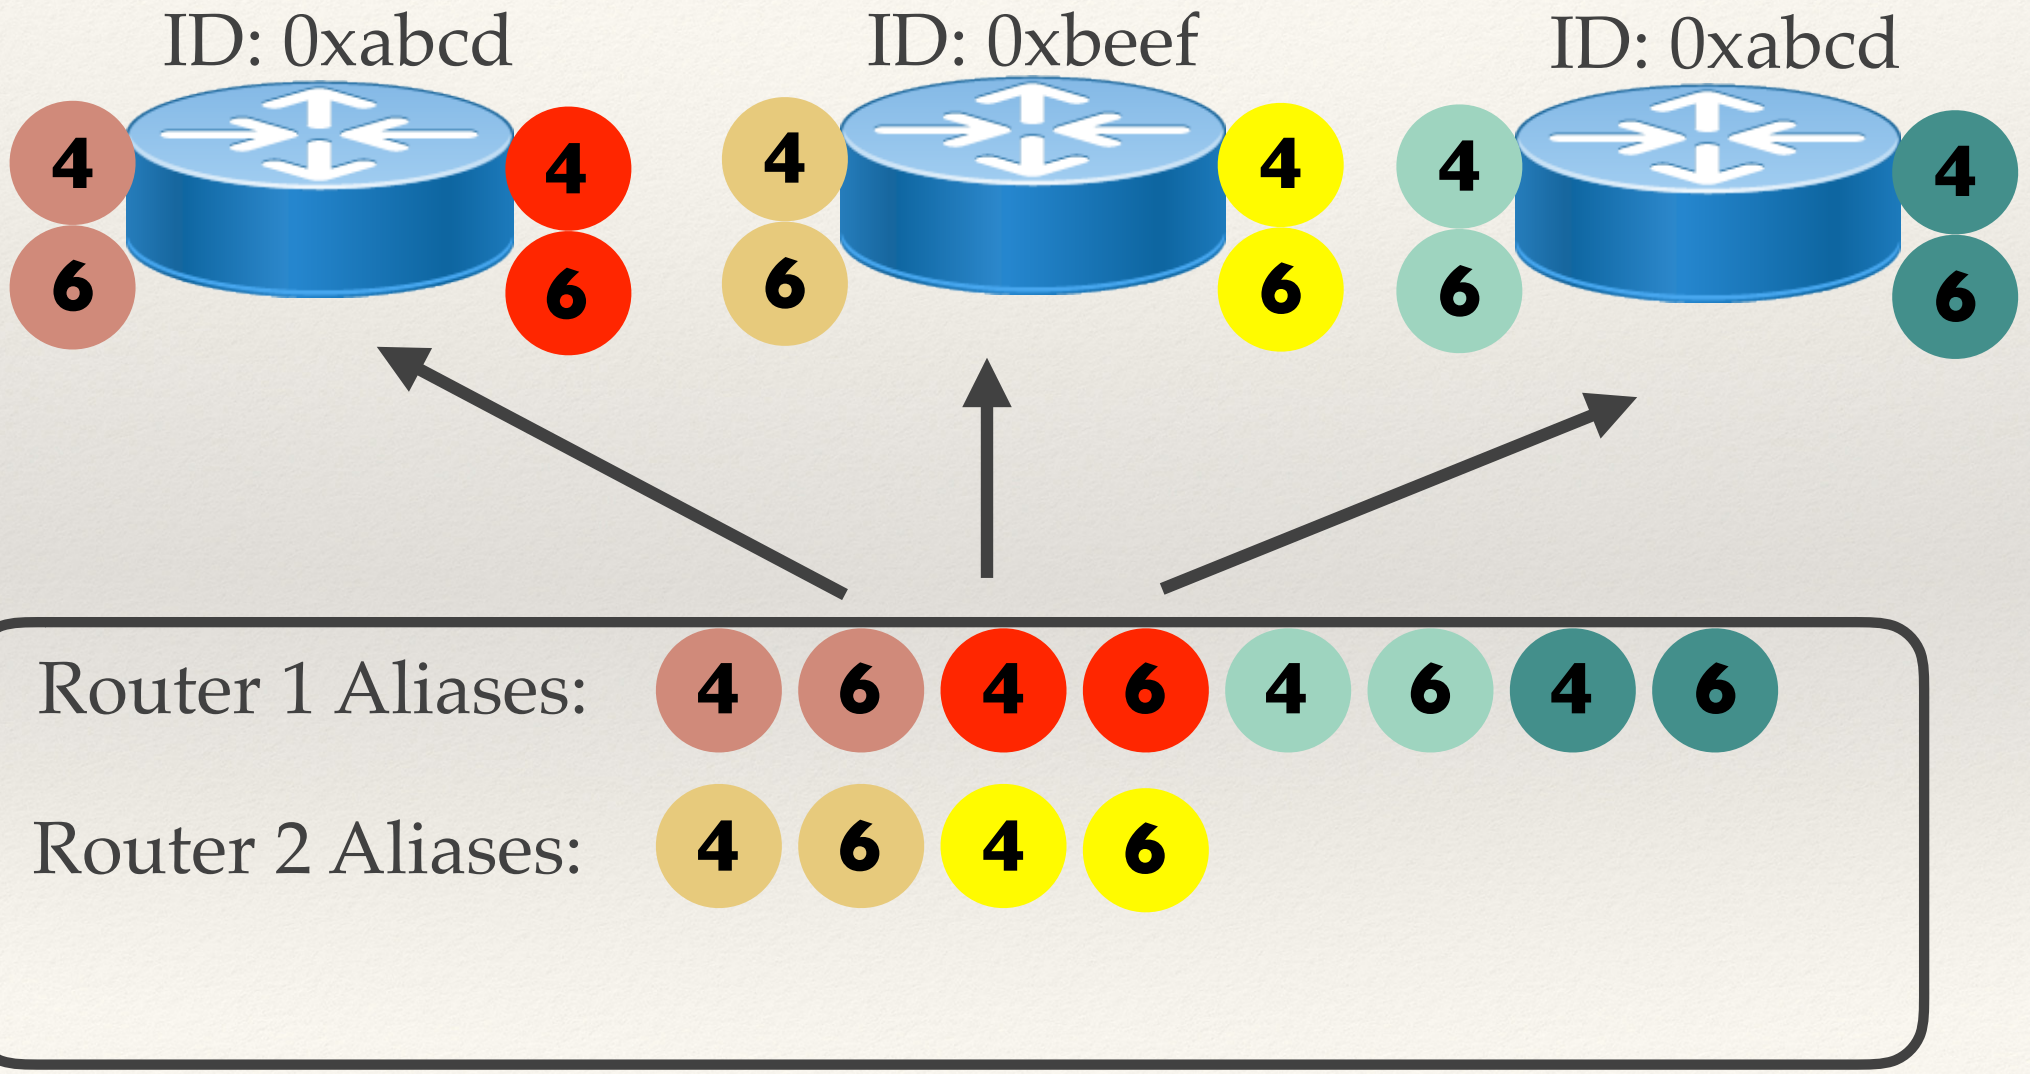

HostID

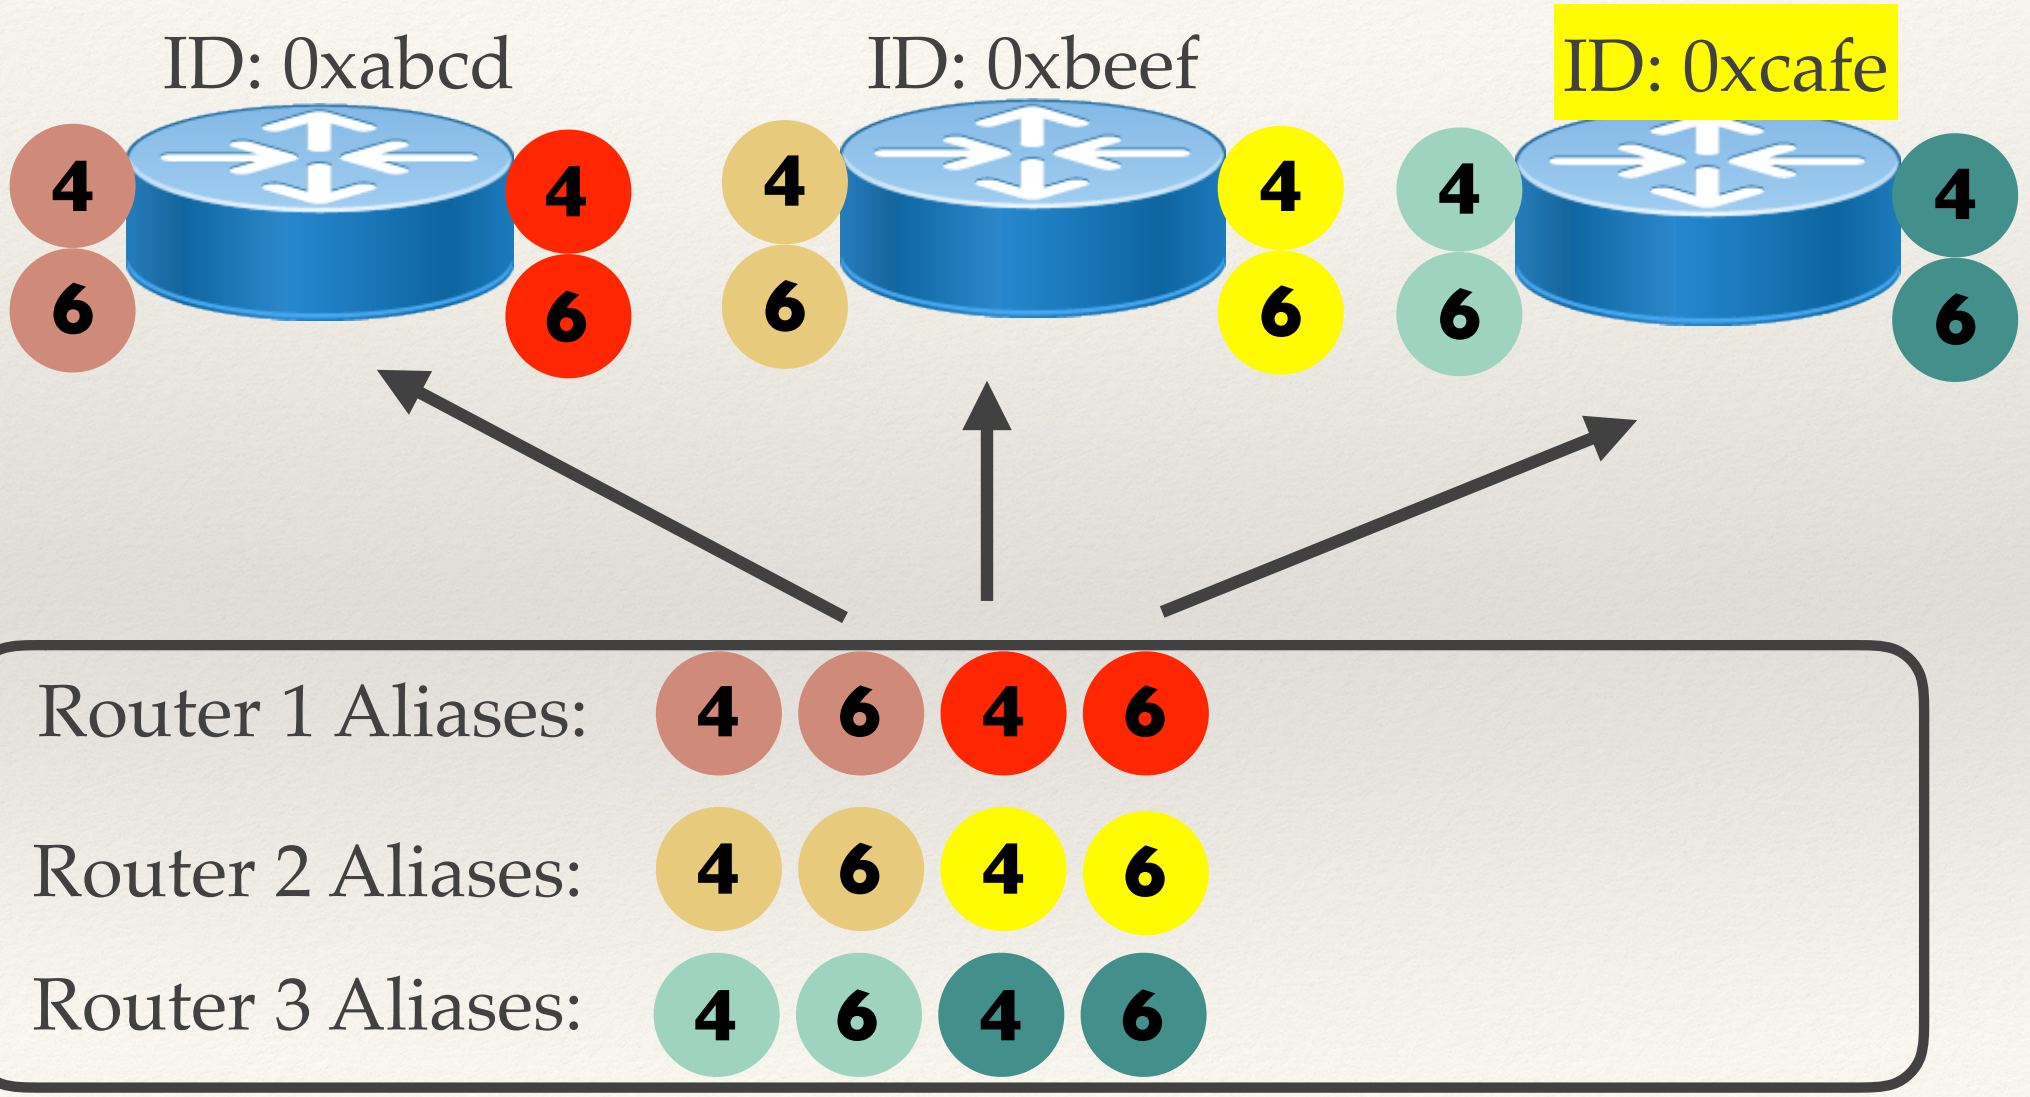
ID: 0xabcd
ID: 0xbeef
ID: 0xcafe
4
6
Router 1 Aliases:
Router 2 Aliases:
Router 3 Aliases:

HostID: Passive Inference

What are the end points in a communication?

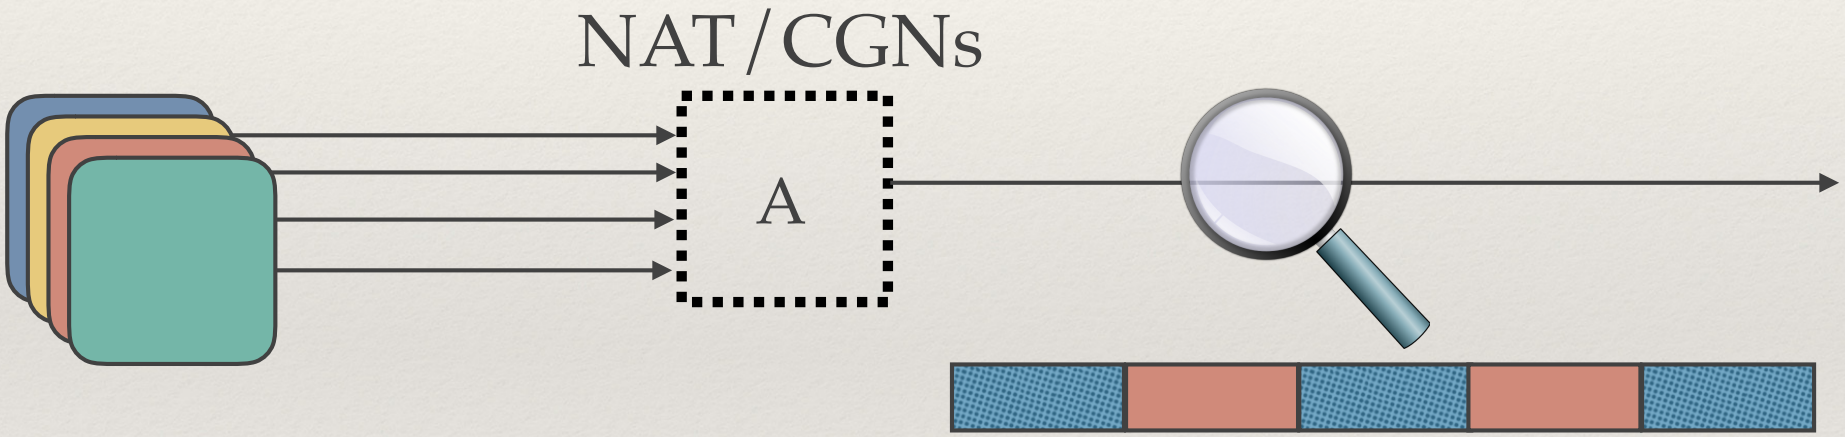

Candidate Primitive: Arrival Info

- How packets arrive at destination
- Nonce tuple (N_{xmit}, N_{sum}):
  - N_{xmit}: random, set by sender
  - N_{sum}: sum of received N_{xmit} values echoed back
  - Permits sender to reconstruct arrival stream

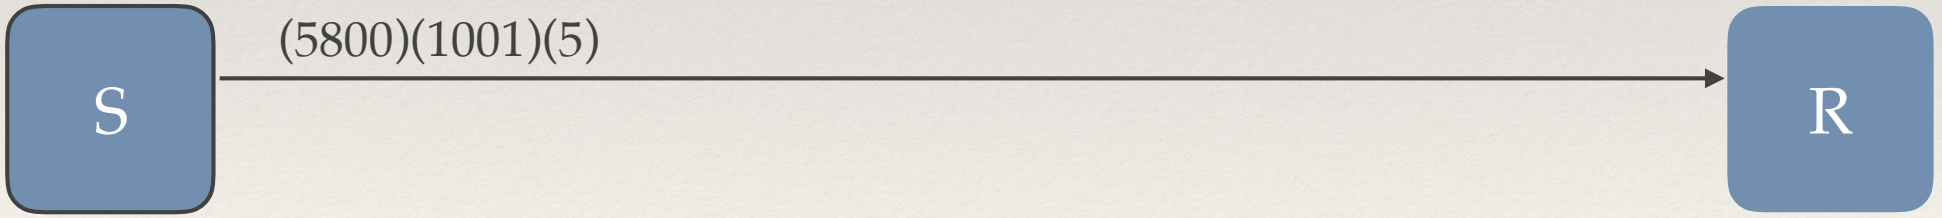

Candidate Primitive: Arrival Info

- How packets arrive at destination
- Nonce tuple (N_{xmit}, N_{sum}):
  - N_{xmit}: random, set by sender
  - N_{sum}: sum of received N_{xmit} values echoed back
  - Permits sender to reconstruct arrival stream

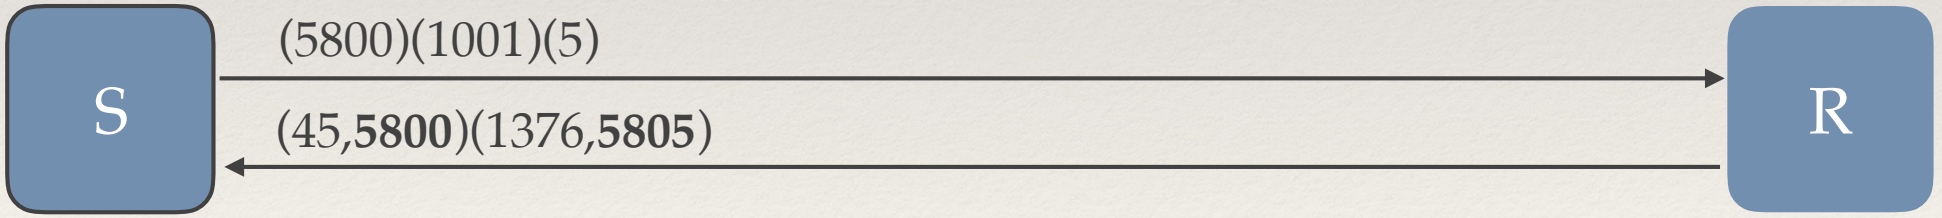

Candidate Primitive: Arrival Info

- How packets arrive at destination
- Nonce tuple (N_{xmit}, N_{sum}):
  - N_{xmit}: random, set by sender
  - N_{sum}: sum of received N_{xmit} values echoed back
  - Permits sender to reconstruct arrival stream

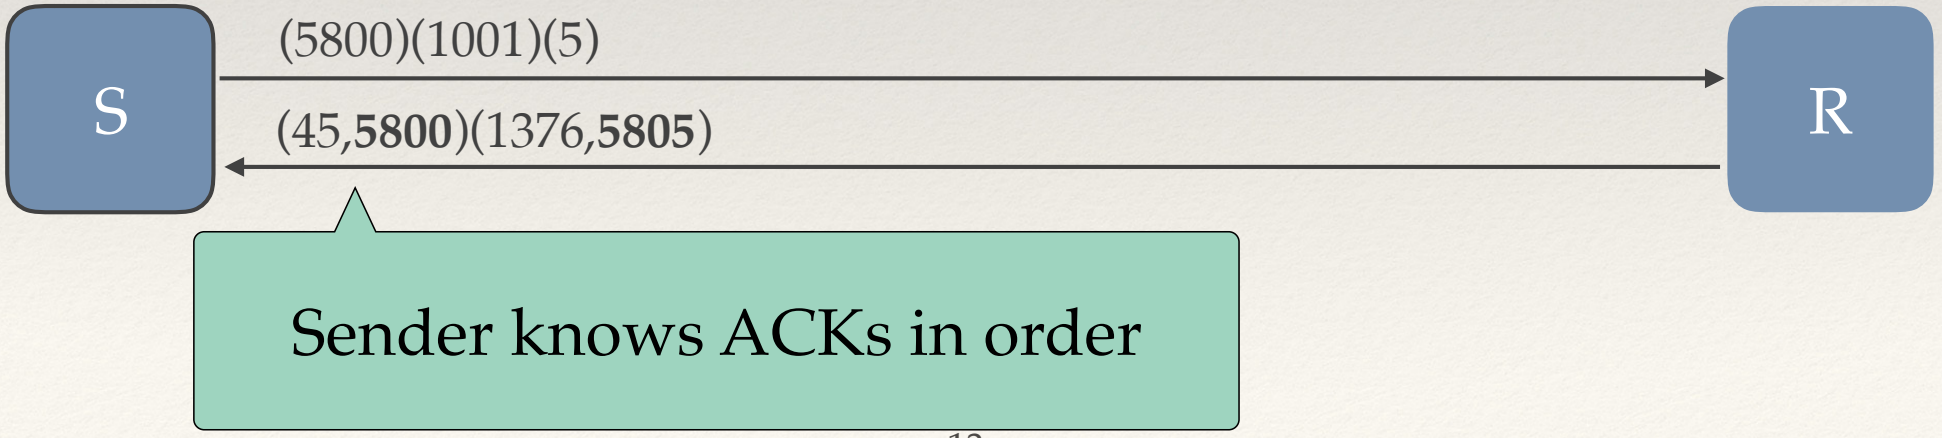

Candidate Primitive: Arrival Info

- How packets arrive at destination
- Nonce tuple (N_{xmit}, N_{sum}):
  - N_{xmit}: random, set by sender
  - N_{sum}: sum of received N_{xmit} values echoed back
  - Permits sender

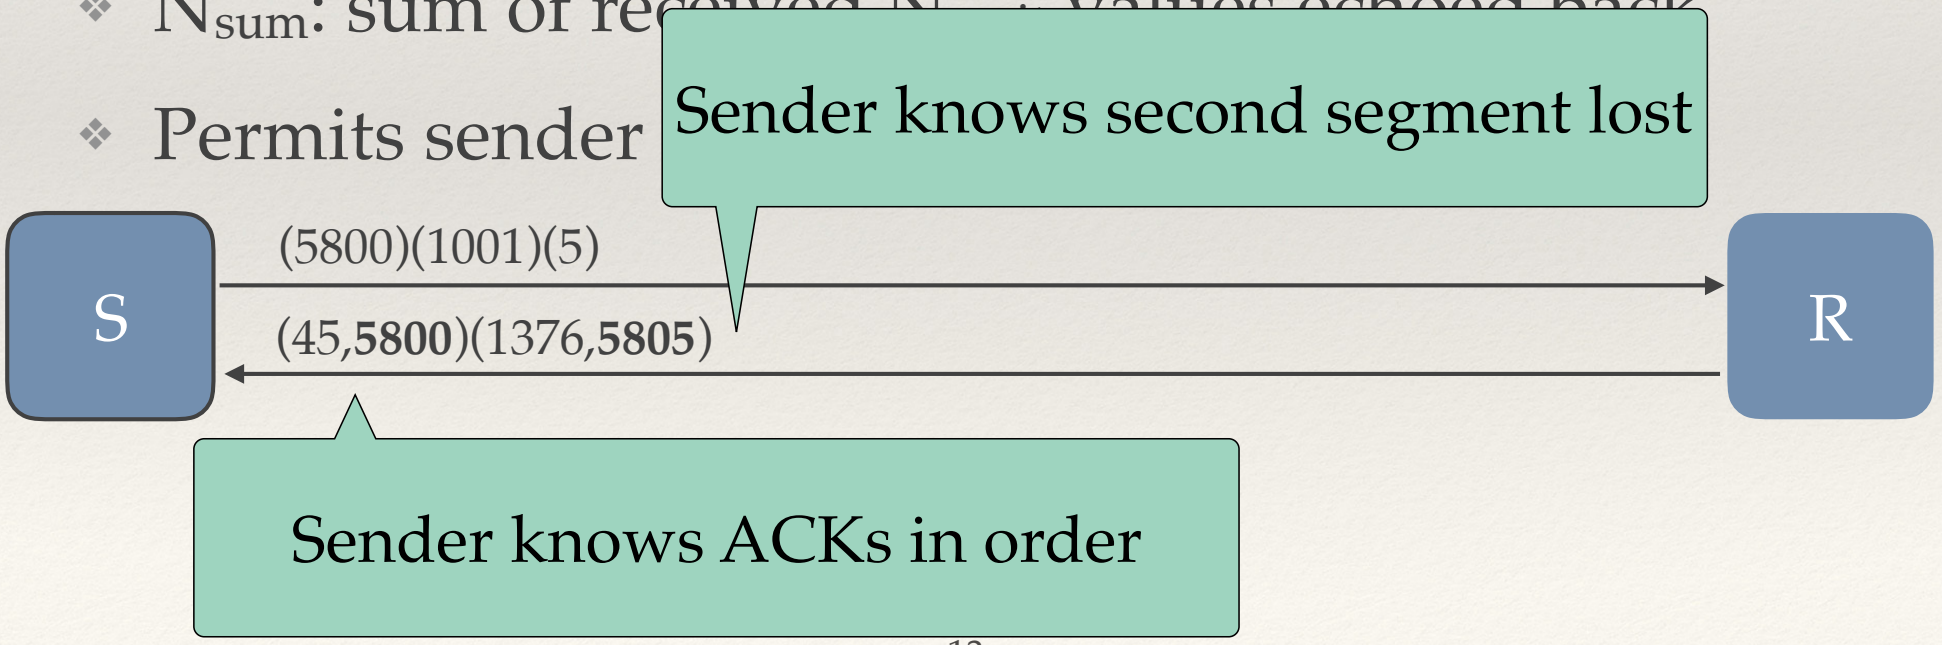

Candidate Primitive: Arrival Info

- How packets arrive at destination
- Nonce tuple (N_{xmit}, N_{sum}):
  - N_{xmit}: random, set by sender
  - N_{sum}: sum of rec[illegible] values echoed back
  - Permits sender

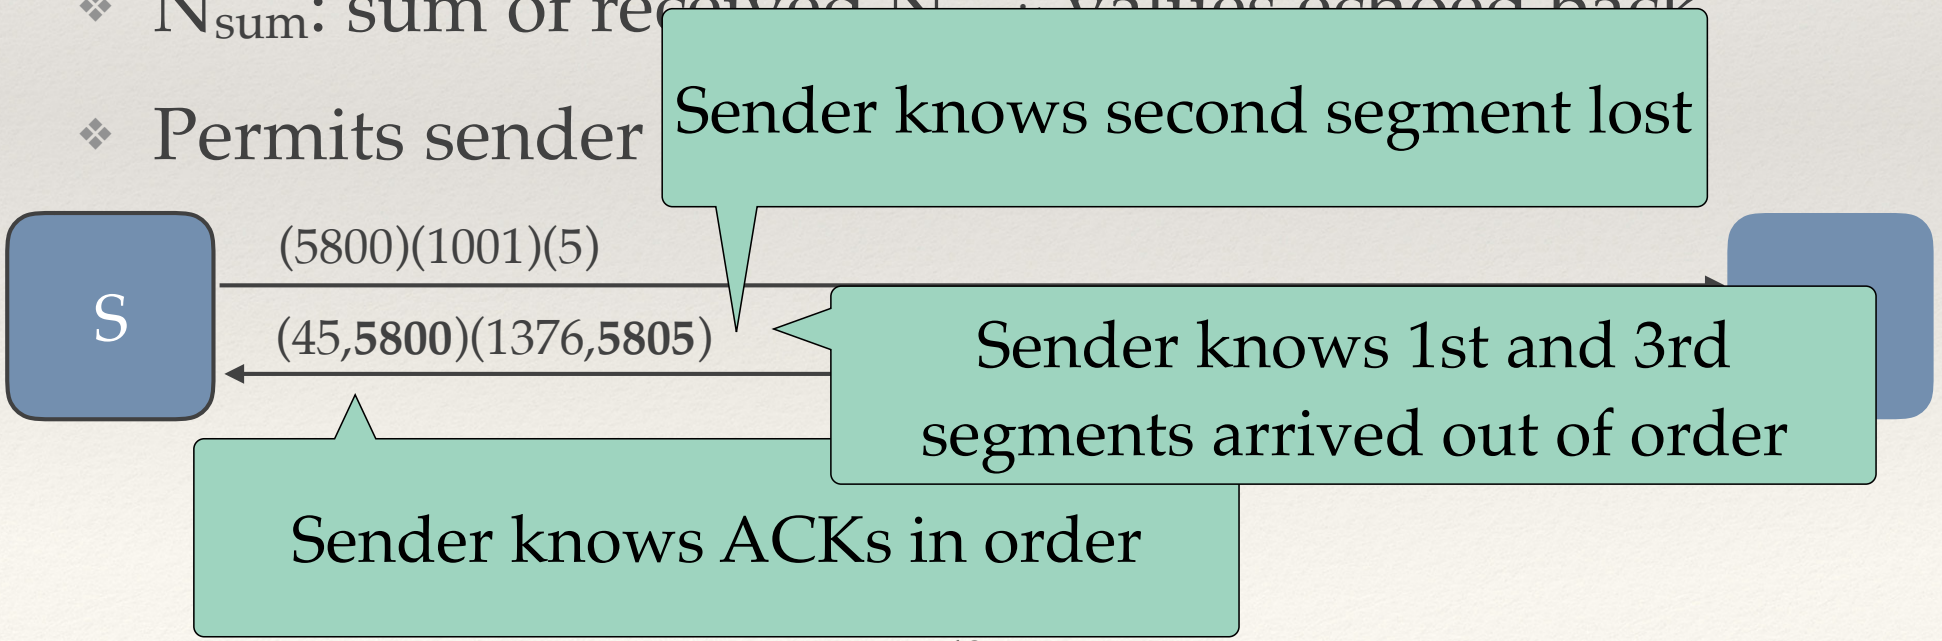

Candidate Primitive: Arrival Info

- How packets arrive at destination
- Nonce tuple (N_{xmit}, N_{sum}):
  - N_{xmit}: random, set by sender
  - N_{sum}: sum of received N_{xmit} values echoed back
  - Permits sender to reconstruct arrival stream

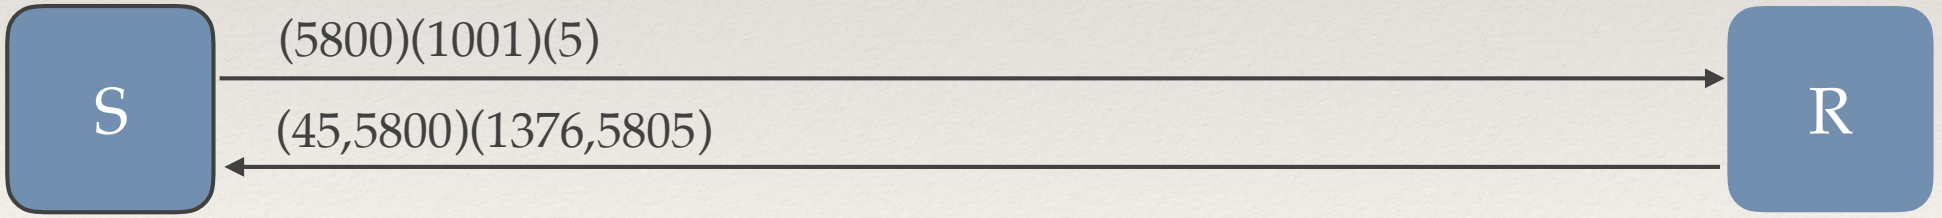

Candidate Primitive: Arrival Info

- How packets arrive at destination
- Nonce tuple (N_{xmit}, N
  - N_{xmit}: random, set
  - N_{sum}: sum of received N_{xmit} values echoed back
  - Permits sender to reconstruct arrival stream

P5: Visibility

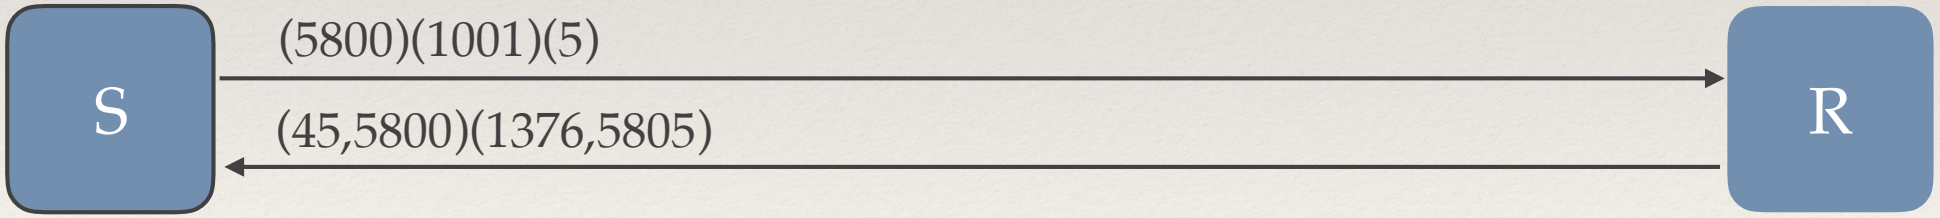

Candidate Primitive: Arrival Info

- How packets arrive at destination
- Nonce tuple (N_{xmit}, N
  - N_{xmit}: random, set
  - N_{sum}: sum of received N_{xmit} values echoed back
  - Permits sender to reconstruct arrival stream

P5: Visibility

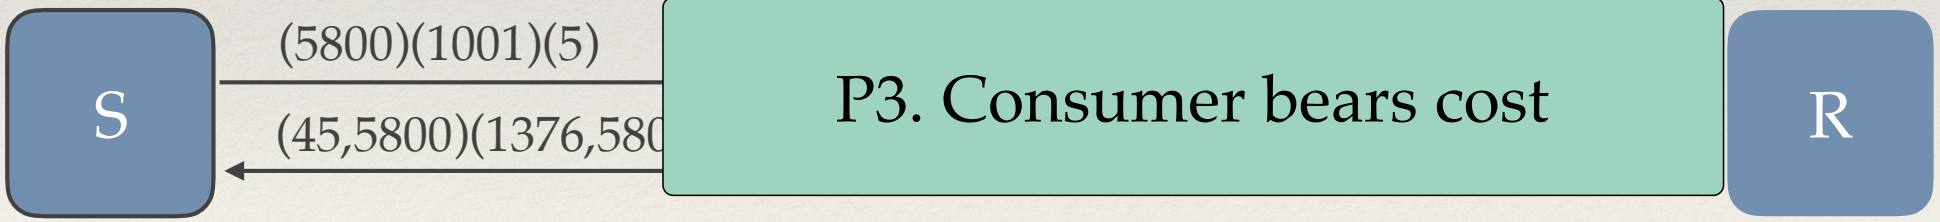

P3. Consumer bears cost

Network support

- Imagine increased cooperation and support from the network
  - Topology tuples
  - Path change tuples
  - Performance tuples
  - Accumulated performance tuples

Network support

- Imagine increased cooperation and support from the network
  - Topology tuples
  - Path change tuples
  - Performance tuples
  - Accumulated performance tuples

See paper for details on these

Conclusions

- Network **measurement critical**, we need better tools, and better tools **need better support from the network**
- Propose guiding **principles** for viable measurement
- Demonstrate candidate **primitives** that address long-standing, important real-world measurement problems
- Position paper: spur discussion, debate, and **inform protocol development**