NET WORKING?

S. Keshav
University of Waterloo

DIMACS
May 21, 2012

The definition of a USM is joint work with M. Karsten, S. Prasad, and O. Beg

A typical exchange

I can't get to CNN.com, can you?

Yes, I can…

Oh…

50 ways to lose your connection

- Link failure
- Router failure
- NAT table overflow (SIGCOMM 2011)
- Wire unplugged
- Network congestion
- DNS server failure
- Mis-configured firewall
- Incorrect browser setting
- Personal firewall misconfiguration
- Error in antivirus program
- Wireless AP failure
- Transient routing fault
- …

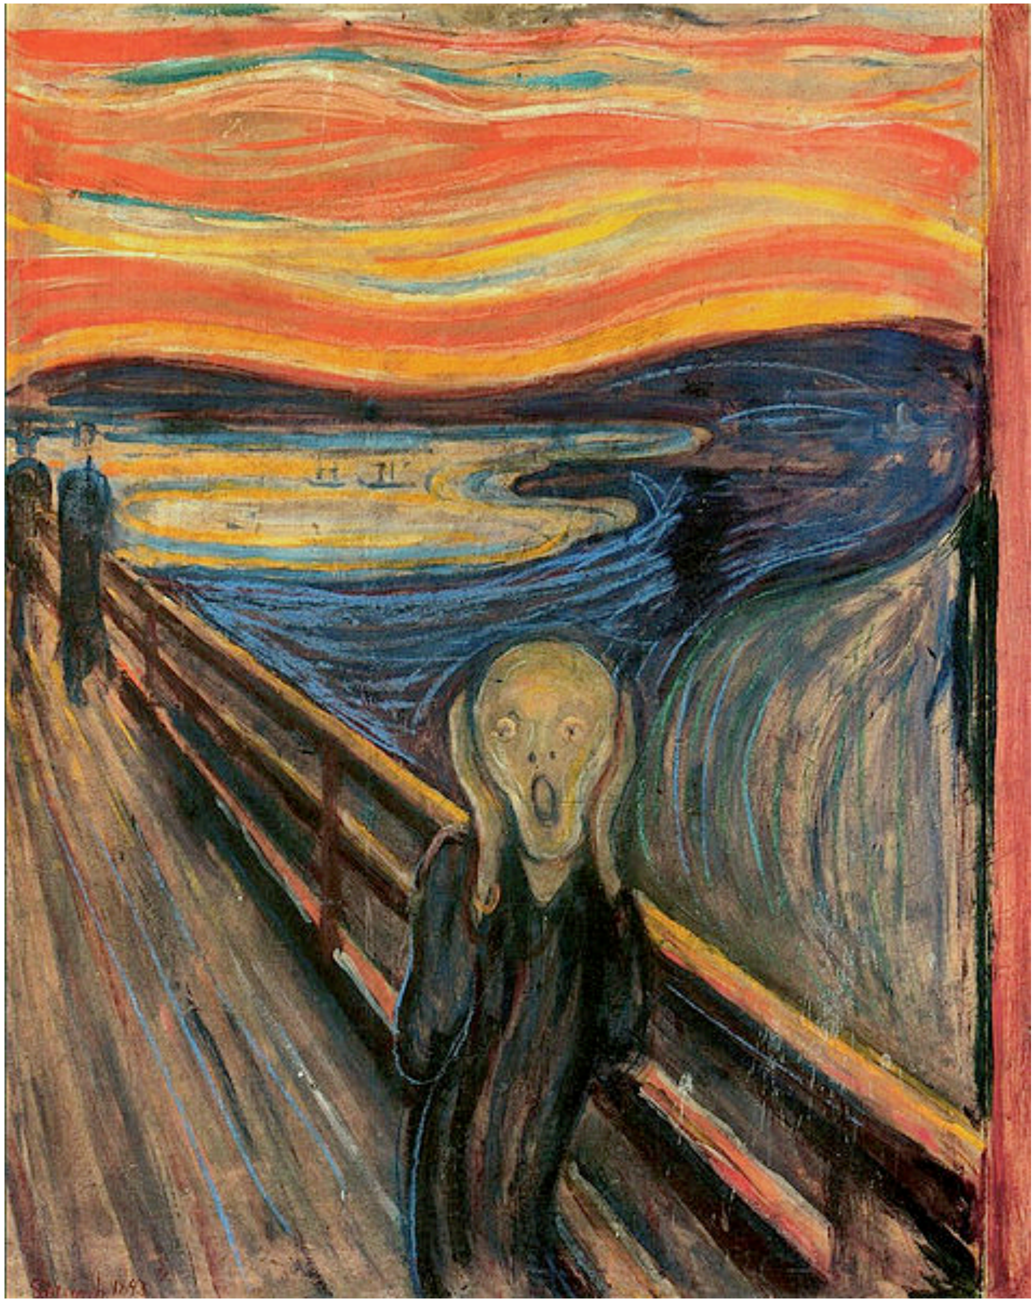

Intuition

- Ability to establish an end-to-end connection depends on many underlying systems functioning ‘correctly’

- How to model this?

Universal Switching Machines

Universal Switching Machine

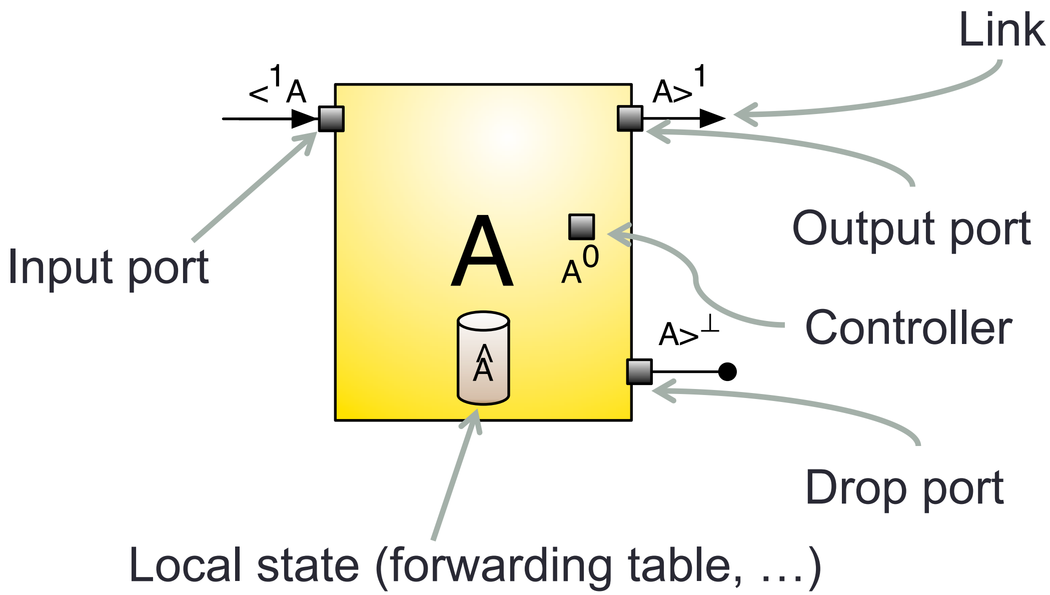

Actions

1. forward, with modified headers
2. drop
3. consume
4. generate

Packet

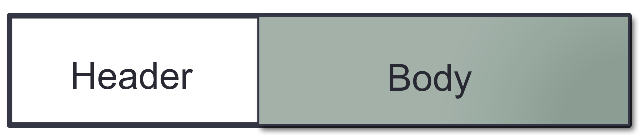

details in “An Axiomatic Basis for Communication, Karsten et al, Proc SIGCOMM 2007

A network of USMs

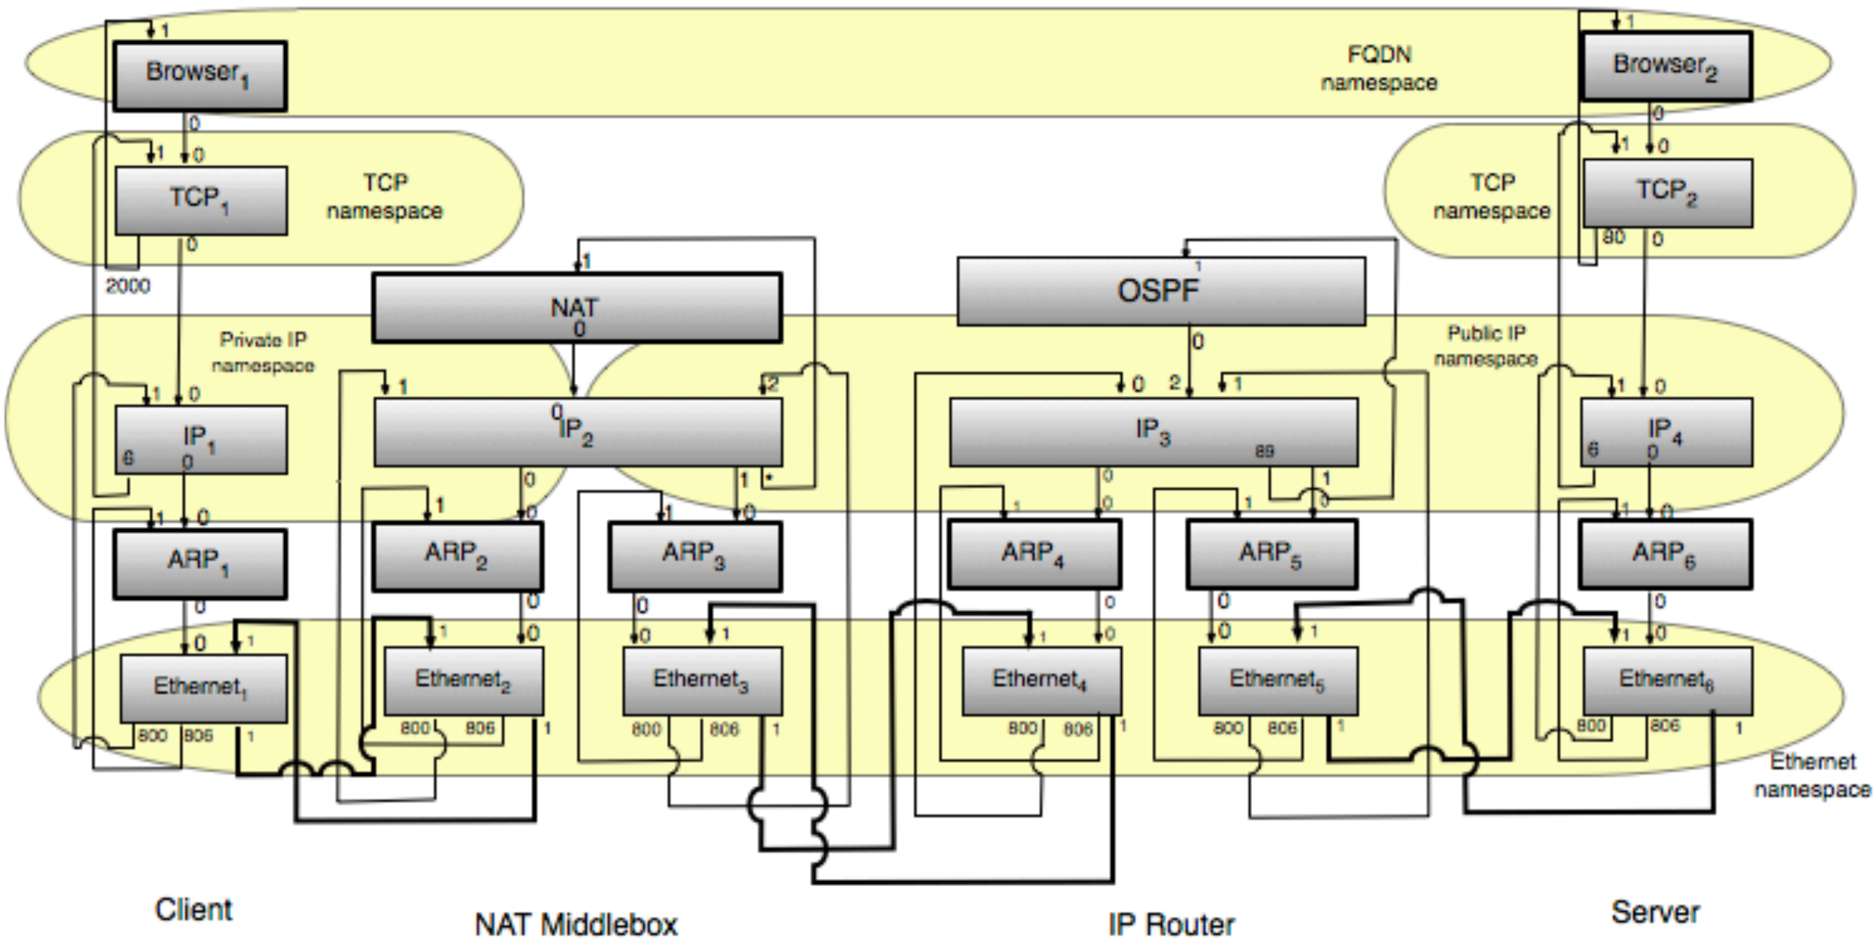

Fundamental questions

- Reachability
  - will a packet from USM A get delivered to USM B?

- What does ‘delivered to B mean’?
- What destination name should A use?

Naming

- What port should A use to reach B?
- What state at intermediate USMs is required?
- How is this state created?

State creation

Naming

What's in a name? That which we call a rose
By any other name would smell as sweet

The name of the song is called "HADDOCKS' EYES."'

'Oh, that's the name of the song, is it?' Alice said, trying to feel interested.

'No, you don't understand,' the Knight said, looking a little vexed. 'That's what the name is CALLED. The name really IS "THE AGED AGED MAN."'

'Then I ought to have said "That's what the SONG is called"?' Alice corrected herself.

'No, you oughtn't: that's quite another thing! The SONG is called "WAYS AND MEANS": but that's only what it's CALLED, you know!'

'Well, what IS the song, then?' said Alice, who was by this time completely bewildered.

'I was coming to that,' the Knight said. 'The song really IS "A-SITTING ON A GATE": and the tune's my own invention.'

…which one will respond to ‘Rose!’

Name for self

- Let A's name for itself be 'a'
  - A accepts packets with 'a' in header
  - Similar to an accepting state in a Turing Machine

- More precisely:

  The bit string

  which, when present in the destination field of a packet,

  when the packet is present at the i^{th} input port of USM A

  causes the packet to be delivered to the controller port of USM A

  is a name for A at that port

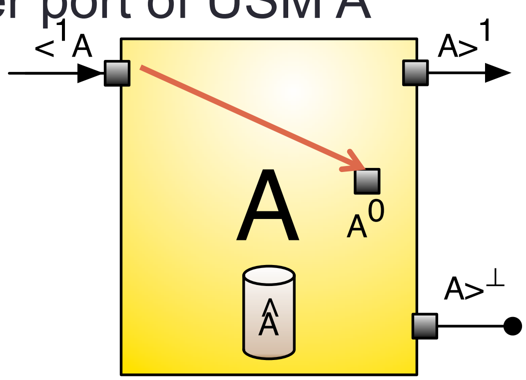

Definite and broadcast names

- Assume that every USM accepts the name ‘*’
  - allows broadcast
- A definite name for A is a name for it other than ‘*’

Namespace

- At any USM A,
  - (non-unique) name for B at A is written b_a
    - can be a source route
- The namespace at A is the set of definite names it has for every other USM
  - b_a, c_a, d_a, e_a …

Some comments

- Names and addresses are treated alike
- A USM can only be sure of its own name for itself
  - We have to assume that USMs do not lie about their names

Aside

- For global reachability, there must exist at least one shared global namespace G

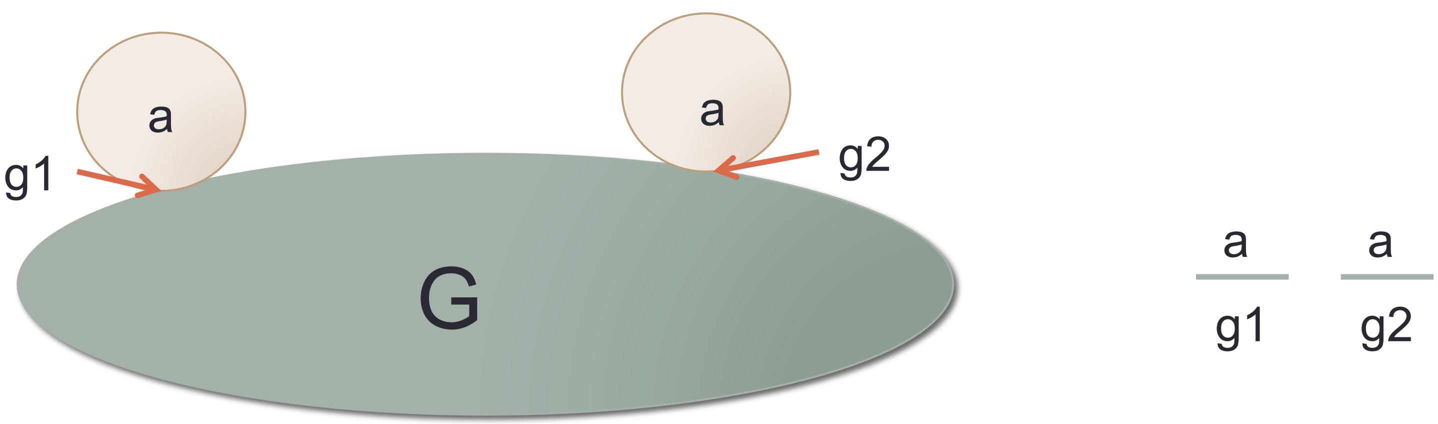

Stacked names

This can be used to set up temporary names (e.g. VCID) or for translation

Hypothesis

Reachability in a network of USMs corresponds to computability in a UTM

State creation

Three problems

- How does a USM learn of the definite name for another USM?
- Given a name, what output port to use?
- How to install state in intermediate USMs?

Routing using announcements simultaneously solves all three problems!

Name announcements

- Tells recipient of the existence of a USM with a particular name
  - assuming a bidirectional link, where it came from tells the recipient what path to take
    - abstracts distance-vector routing

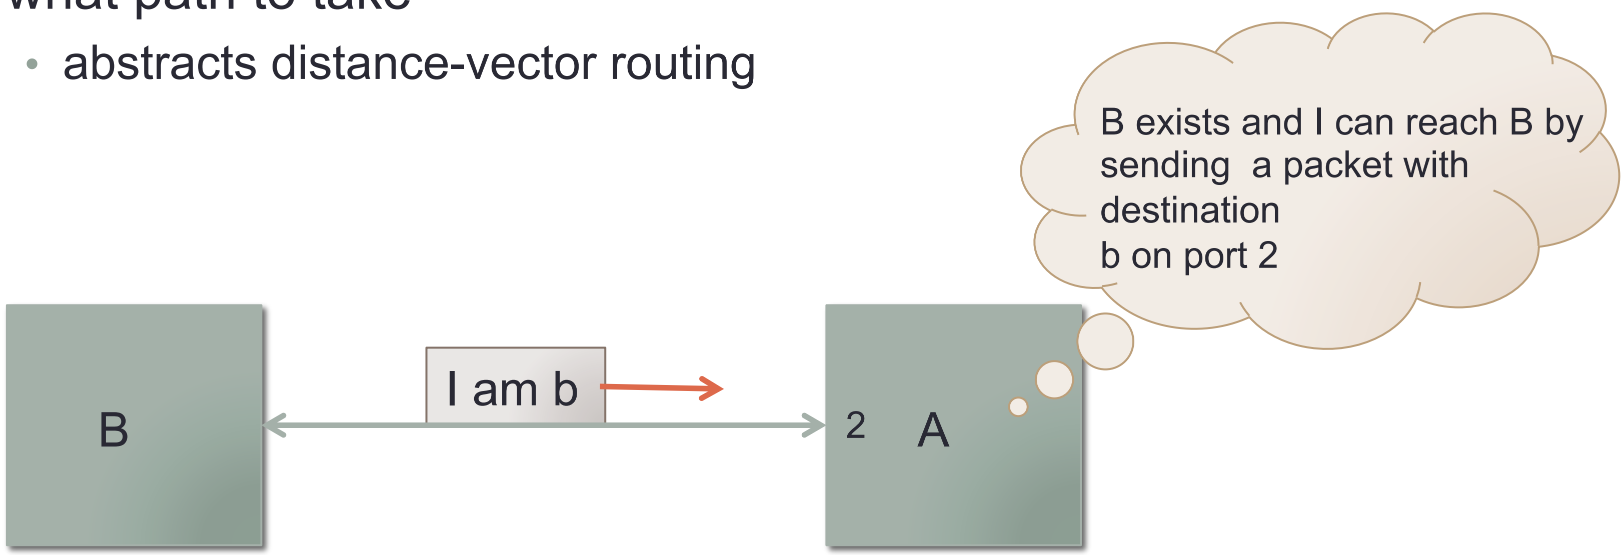

Source address is an implicit name announcement

Extension 1

A broadcasts announcements to all its output ports

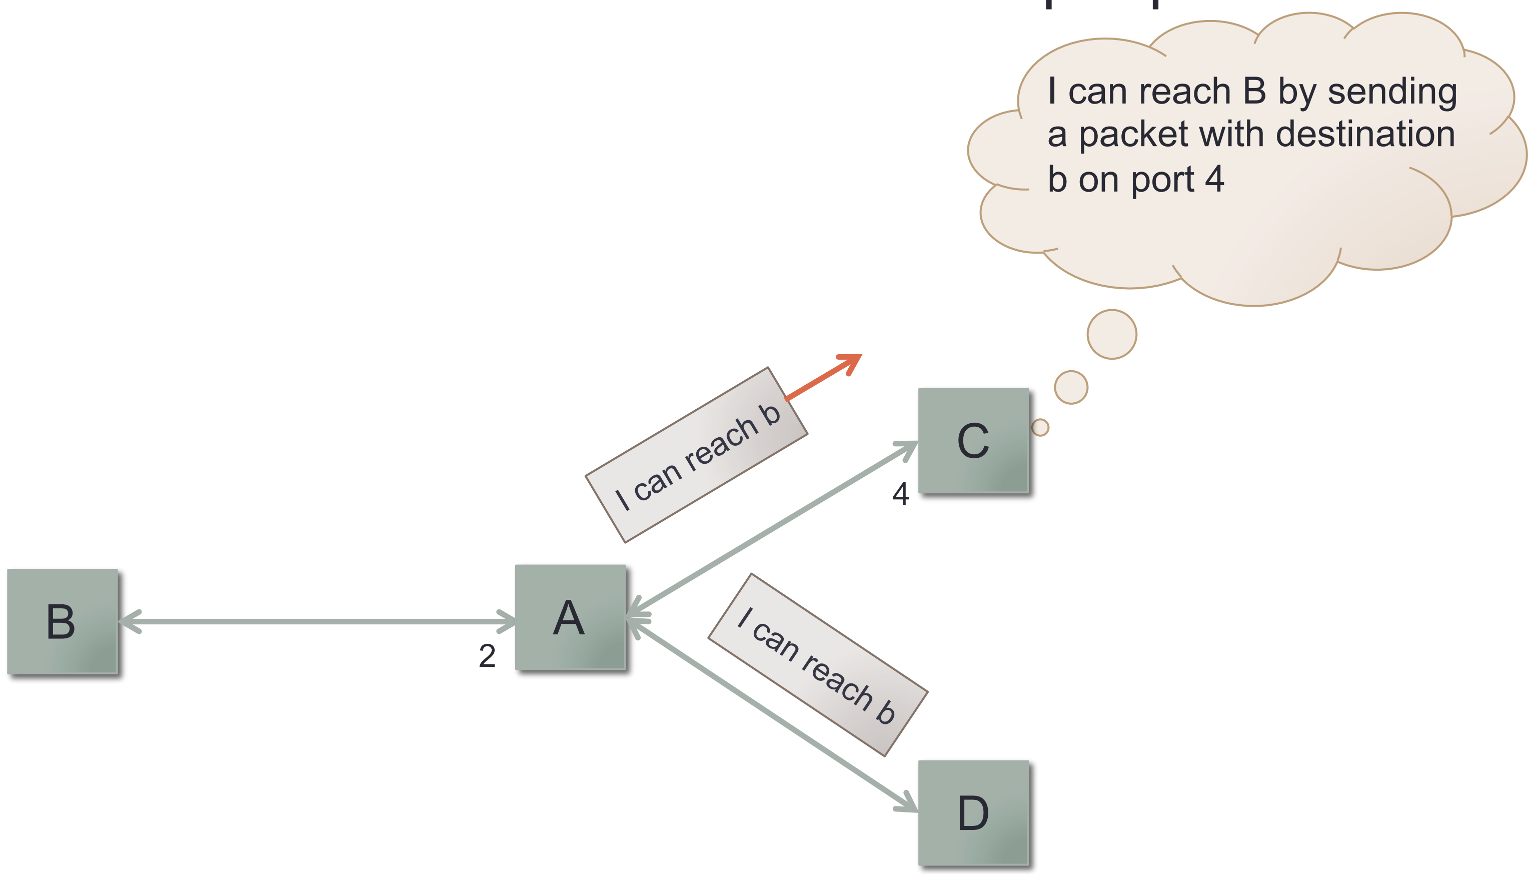

Extension 2

Announcements carry costs

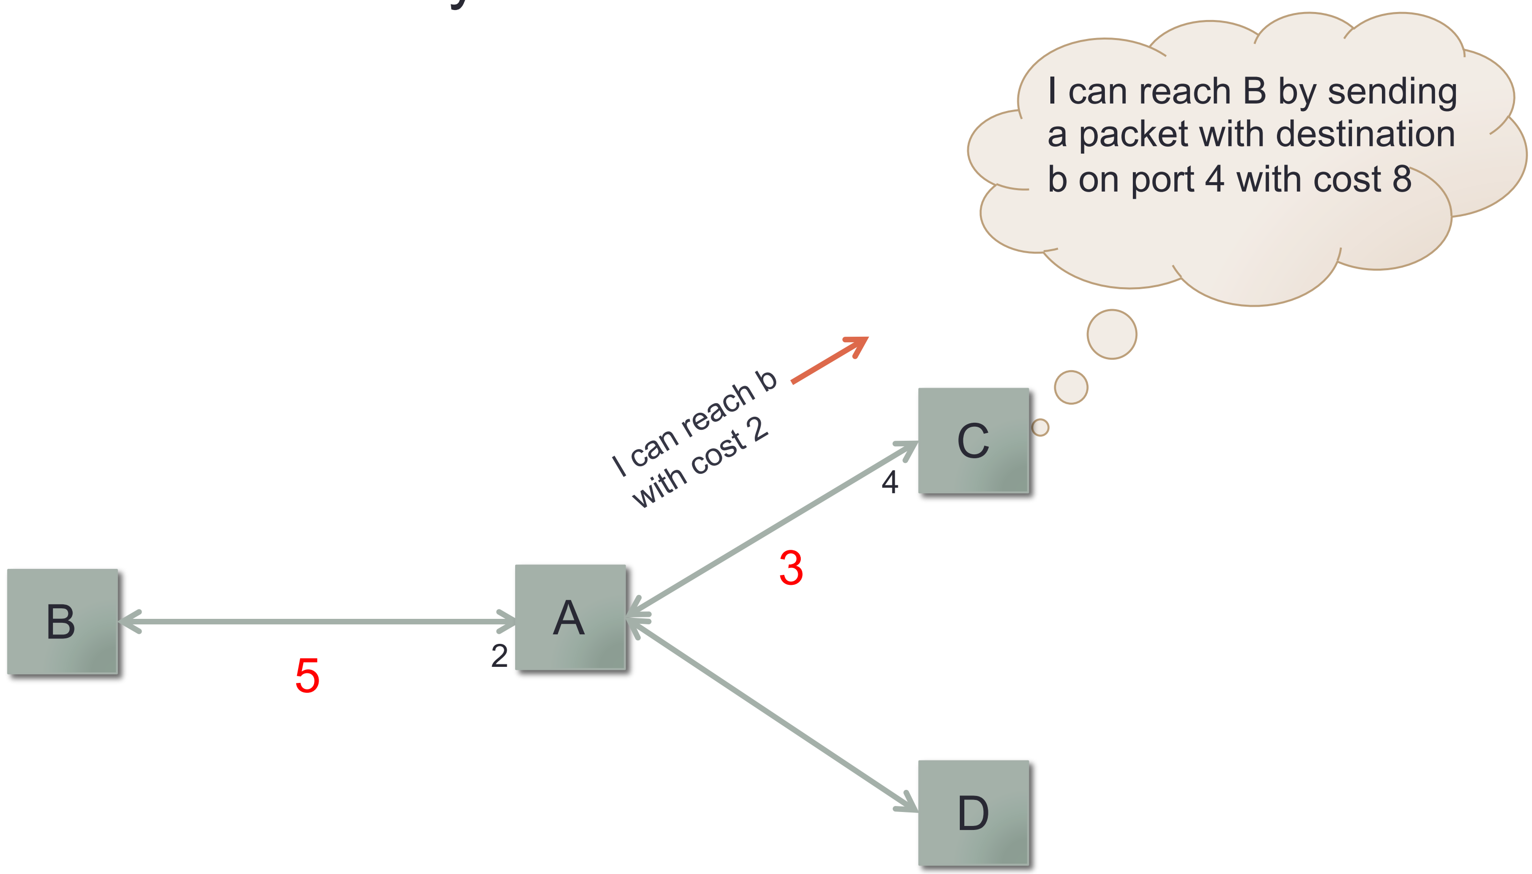

Naturally induces a routing algebra

Link announcements

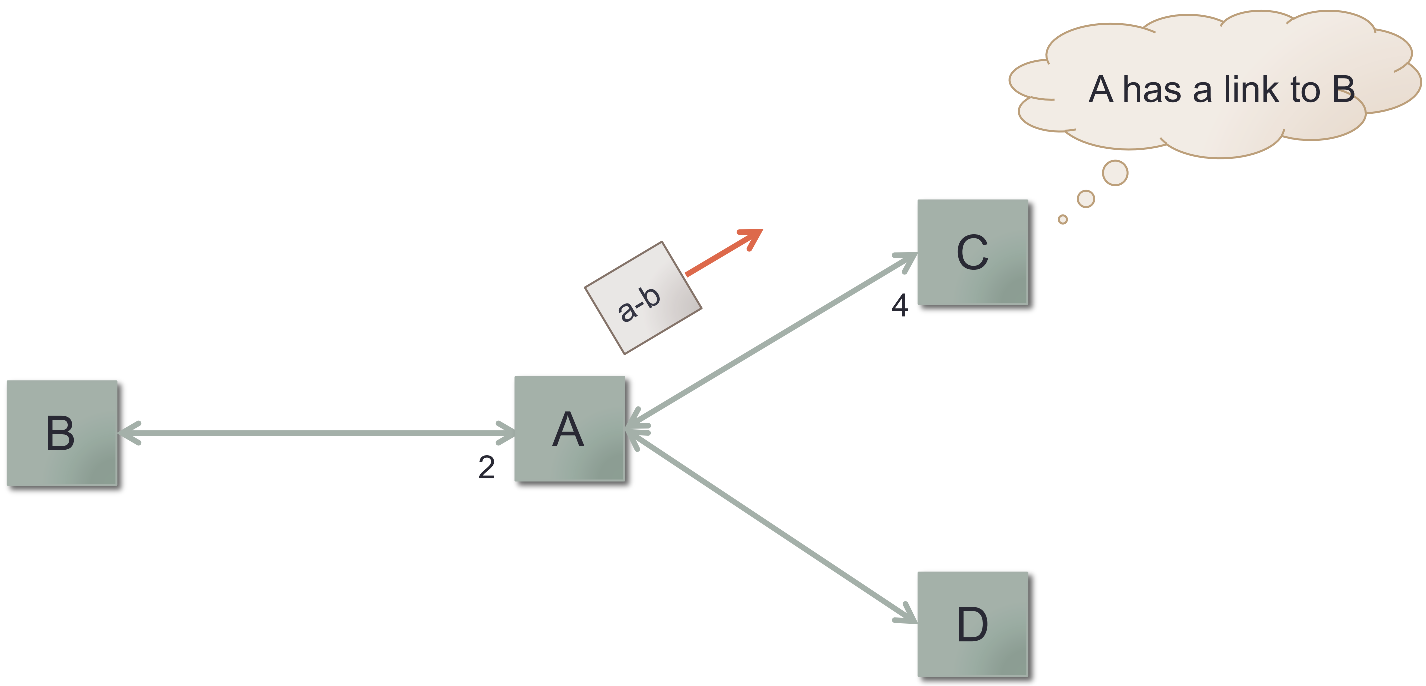

Similar extensions allow population of the forwarding table

Reasoning about net working

A logical approach

- What predicates must hold true for connectivity to be achieved between two USMs?
- Protocols relate logic to networking in the same way that a Von Neumann machine relates Lambda calculus to computing

Approach

- Consider progressively more complicated networks

Simplest possible network

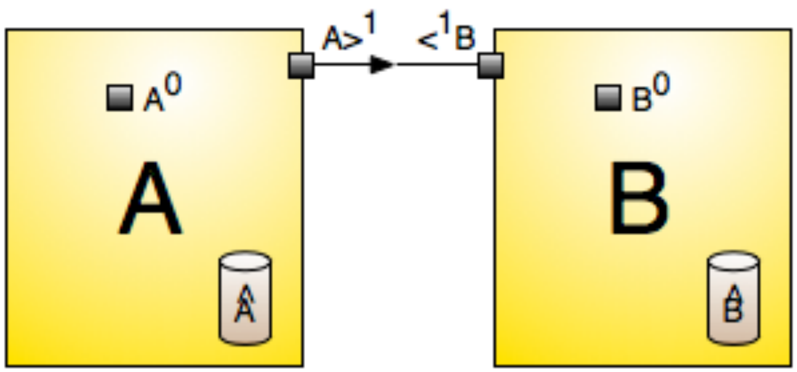

Reachability of B from A requires that

- B is up
- Link is up
- (b,1) is in A's state (exogenously)

Two USMs on a bidirectional link

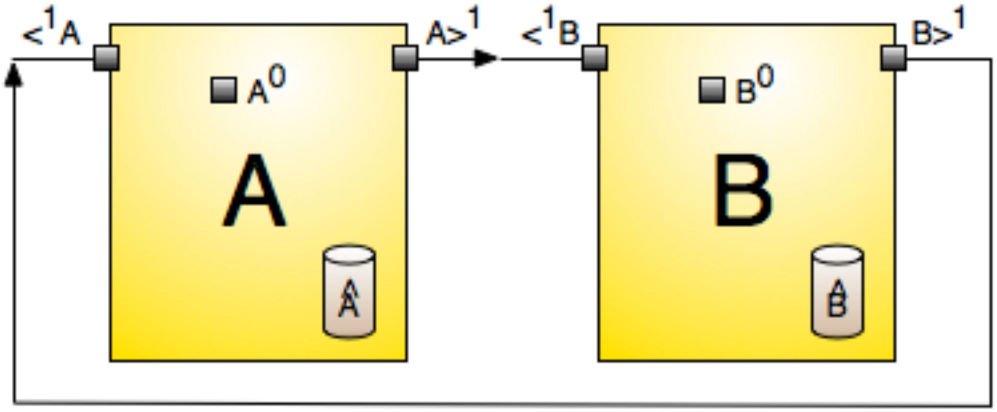

Reachability of B from A requires that

- A,B are up
- Both links are up
- Exogenously introduced state in A and B

  or

  Name announcements generated and processed

 or

  Link announcements generated and processed

One bidirectional broadcast link

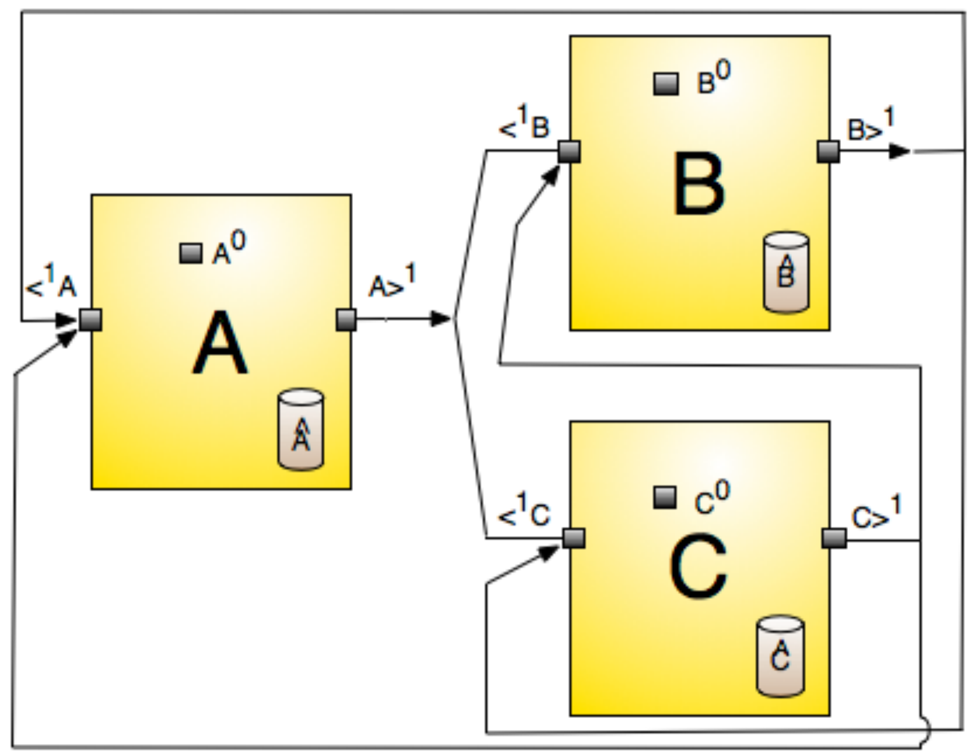

- Nearly identical pre-conditions as with a single bidirectional link

Chains

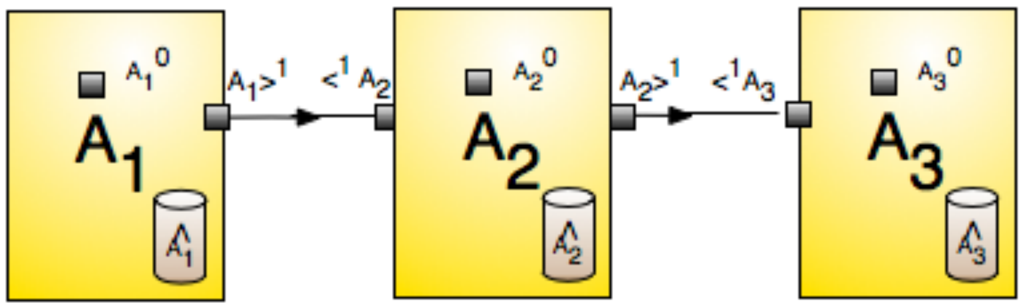

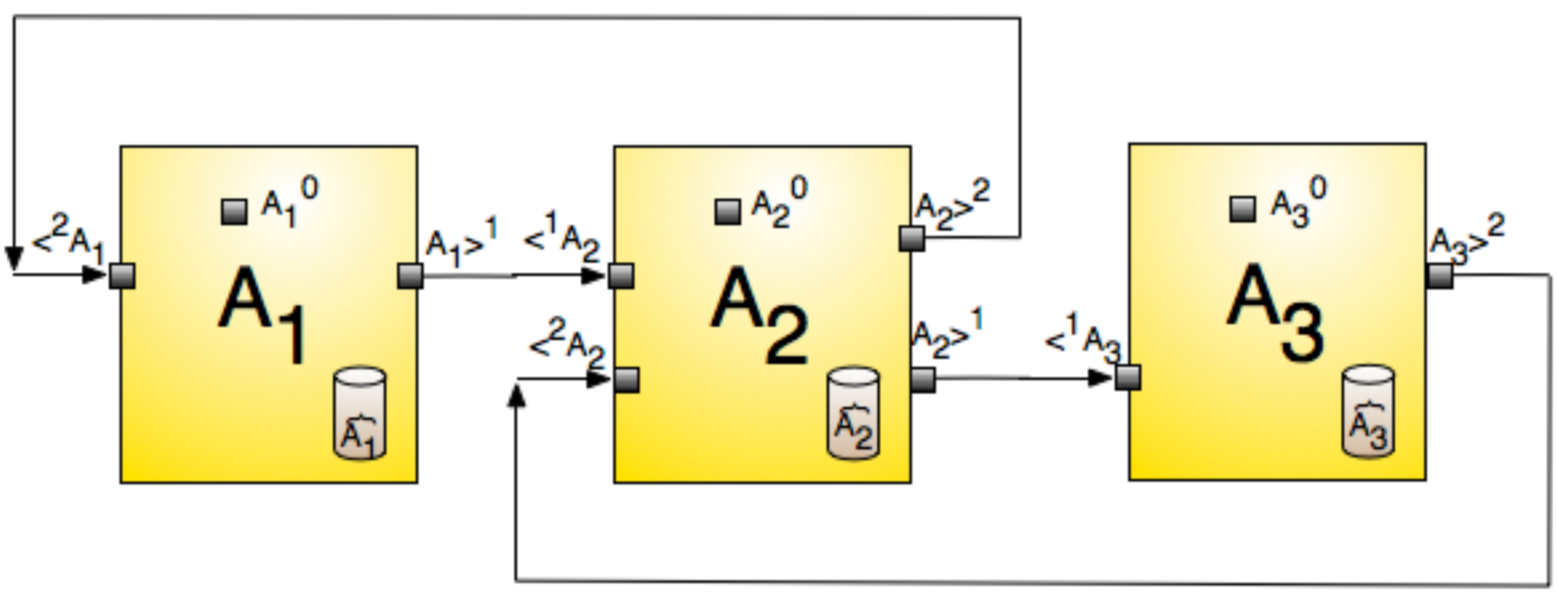

Cycles

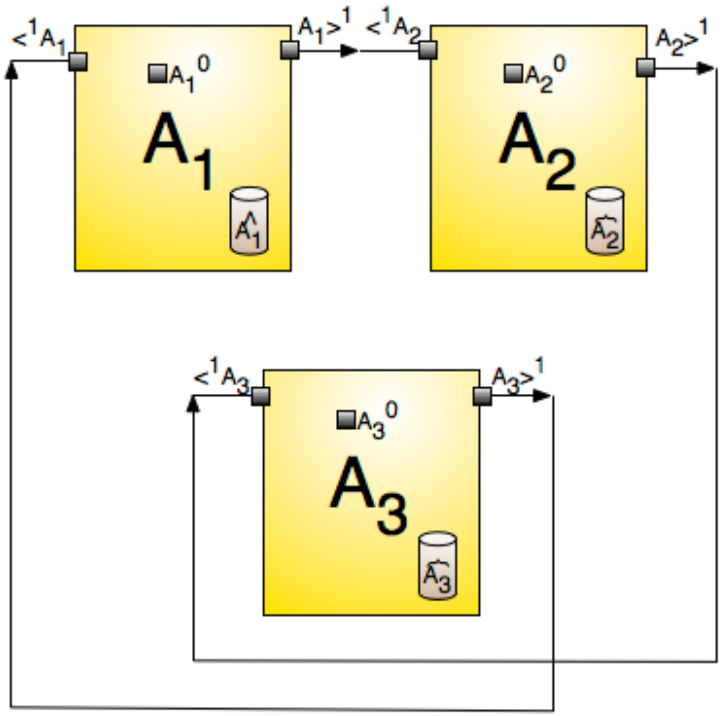
$<^1A_1$
$A_1>^1$
$<^1A_2$
$A_2>^1$
A_1^0
A_1
$\hat{A}_1$
A_2^0
A_2
$\hat{A}_2$
$<^1A_3$
$A_3>^1$
A_3^0
A_3
$\hat{A}_3$

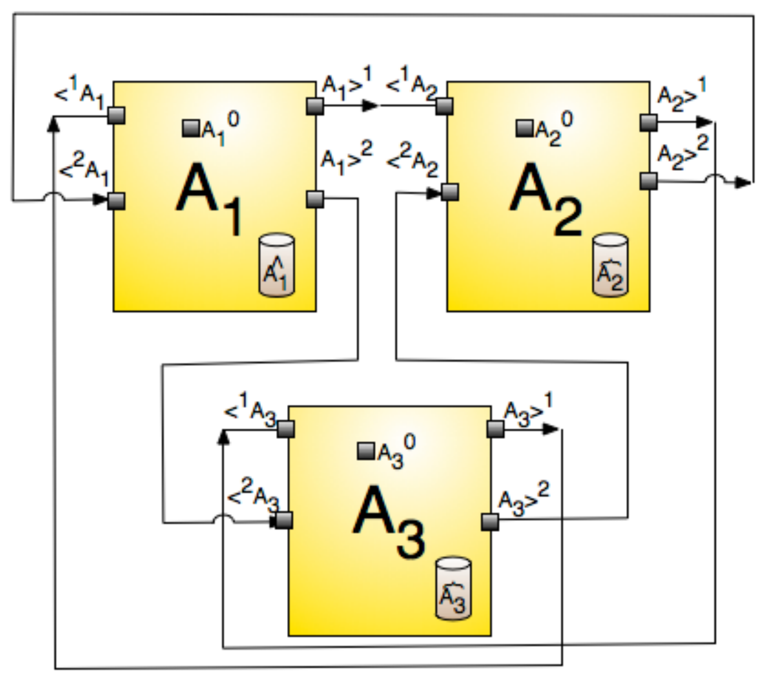
$<^1A_1$
$<^2A_1$
$A_1>^1$
$A_1>^2$
$<^1A_2$
$<^2A_2$
$A_2>^1$
$A_2>^2$
A_1^0
A_1
$\hat{A}_1$
A_2^0
A_2
$\hat{A}_2$
$<^1A_3$
$<^2A_3$
$A_3>^1$
$A_3>^2$
A_3^0
A_3
$\hat{A}_3$

Most general network

- General bidirectional mesh with multiple namespaces, policies, and broadcast links
  - a generalization of the prior results
- Can model the Internet including all middleboxes
- Work still in progress…

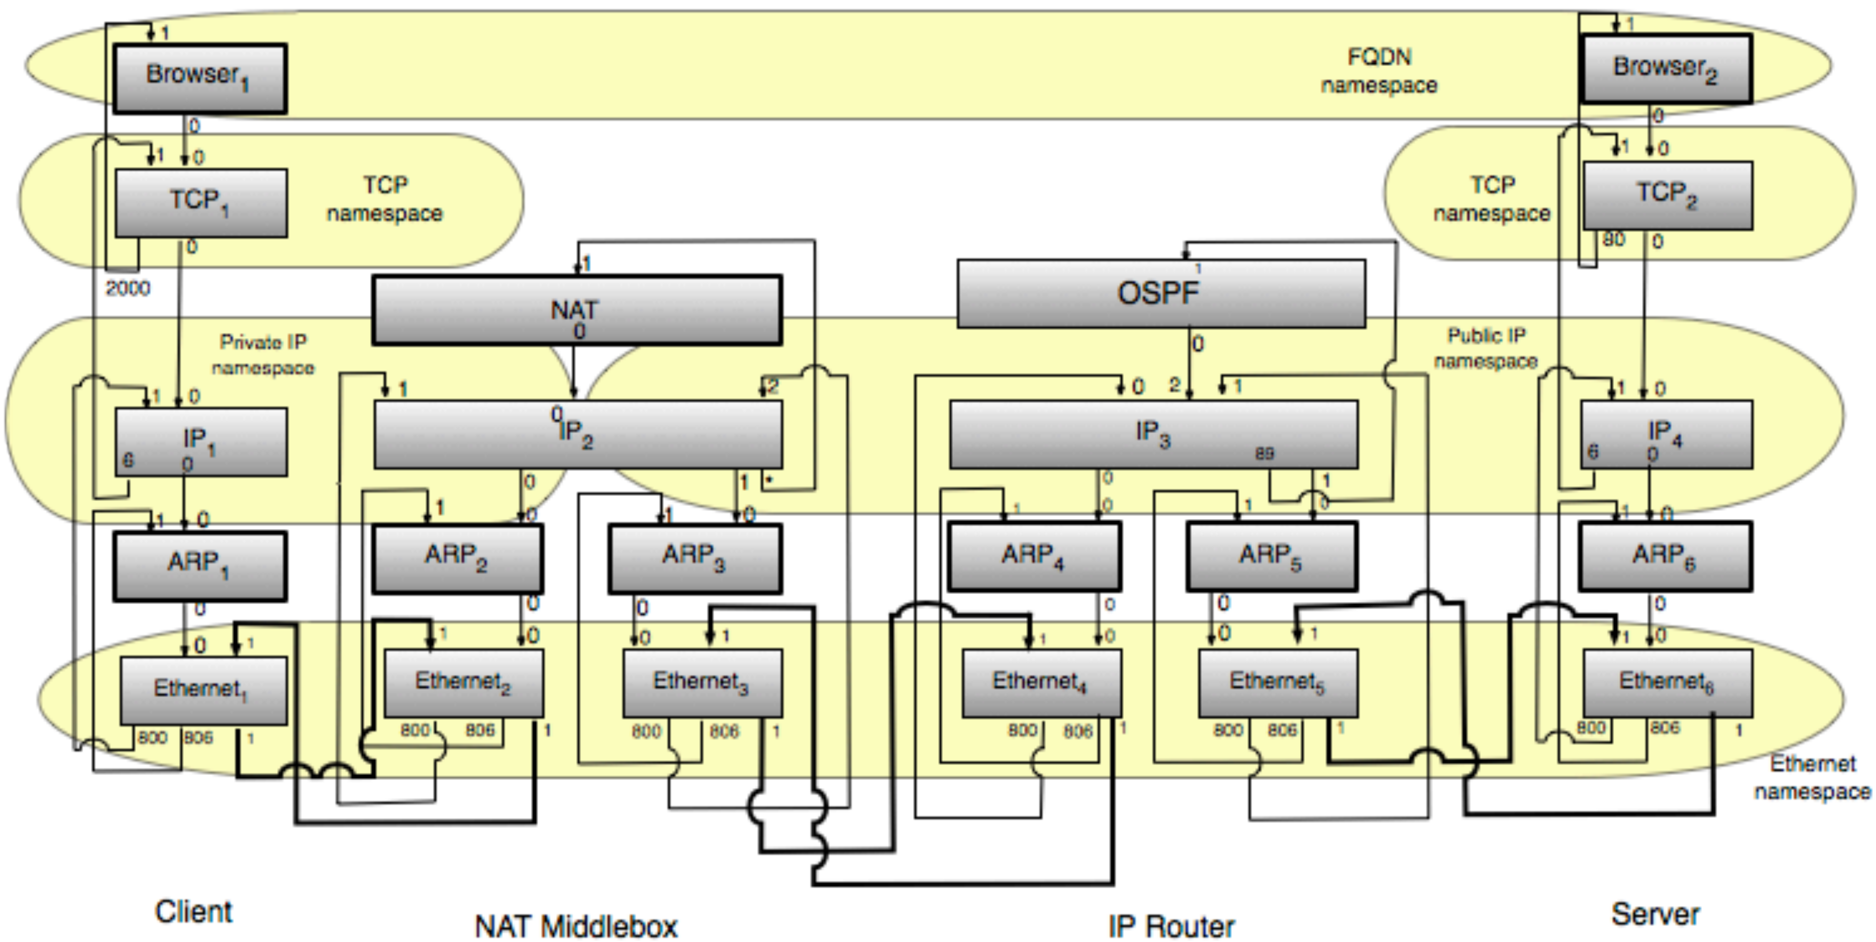

Gateways are shown with bold outlines

Related work

- Naming and Binding
  - Compositional Bindings in Network Domains, Zave, LNCS 2006
  - Axiomatic Basis for Communication, Karsten et al, Proc. SIGCOMM 2007
- Declarative networking
  - Declarative Routing: Extensible Routing with Declarative Queries, Loo et al, Proc. SIGCOMM 2005
  - Towards Declarative Network Verification, Wang et al. Proc. Int. Symp. on Practical Aspects of Declarative Languages, 2009
  - Declarative Networking, Loo and Zhou, Morgan and Claypool, 2012
- Formal analysis of routing
  - Metarouting, Griffin et al, Proc. SIGCOMM 2005
  - A Model of Internet Routing using Semi-Modules, Billings et al, Springer LNCS 2009
  - FSR: Formal Analysis and Implementation Toolkit for Safe Inter-domain Routing, Wang et al, ACM ToN, 2012
- Reasoning about network properties
  - Nexus authorization logic, Shneider et al, ACM Trans. Information and System Security, 2011
  - NetQuery: A Knowledge Plane for Reasoning about Network Properties, Shieh et al, Proc. SIGCOMM 2011
- Automatic generation of OpenFlow configurations
  - Frenetic: A network programming language, Foster et al, Proc. SIGPLAN, 2011
  - A Compiler and Run-time System for Network Programming Languages, Monsanto et al, POPL 2012

Potential outcomes

- Automatic validation of network configurations
  - using automatic theorem proving
- Building self-diagnosis into the system
  - Frenetic/NCore
  - NetQuery

Conclusions

- Determining the availability of an end-to-end path is a complex problem
  - Predicate logic offers a way out
- Universal Switching Machine framework brings together several disparate threads
- General theory of networking relates USM and UTM