A Recursive Network Architecture

Joe Touch
USC/ISI

What makes an architecture new?

- Shaking the Hourglass (CCW 08)
  - All exchanges are 1 packet
  - Collosograms > RTT*delay
  - No LANs? (all L2 was pt-pt)
- What defines success?
  - fixing what's 'broken'
  - doing something new/different
  - the Internet / circuits as a degenerate case

Information Sciences Institute

Motivation

- Desire to support new capabilities
  - Interlayer cooperation, dynamic layer selection, layering created by virtualization
- Desire to support emerging abstractions
  - Overlay layers don't map to 1-7
  - Support for recursive nodes (BARP, LISP, TRILL)
- Desire to coordinate services in diff. places
  - Security, soft-state, pacing, retransmission

Information Sciences Institute

Observations

- Networking is *groups of interacting parties*
  - Groups are heterogeneous
  - All members want to interact
  - Groups can be dynamic (*i.e.*, virtual)
- Need an architecture that supports:
  - Heterogeneity
  - Interaction
  - Virtualization

Heterogeneity leads to layering

- M different interacting parties need
  - M^2 translators

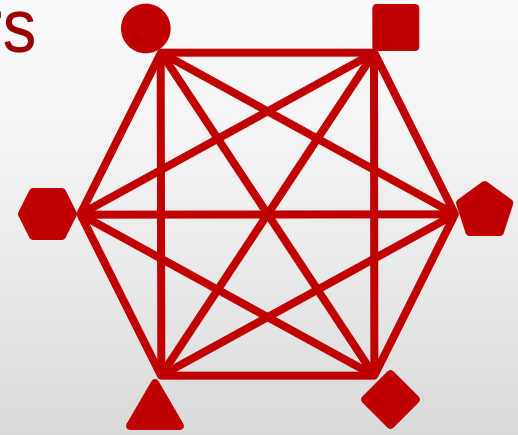

or

  - M translators + common format

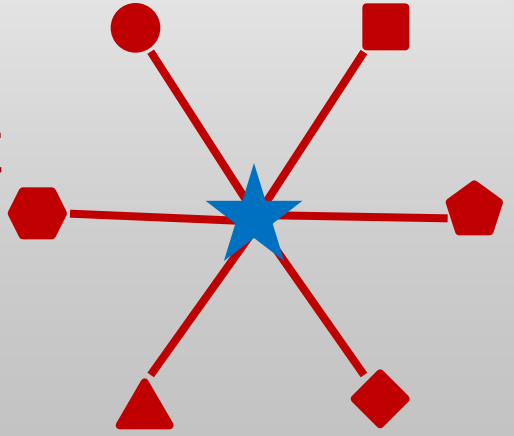

... *i.e.,* a layer

Interaction leads to forwarding

- N parties need
  - N^2 circuits

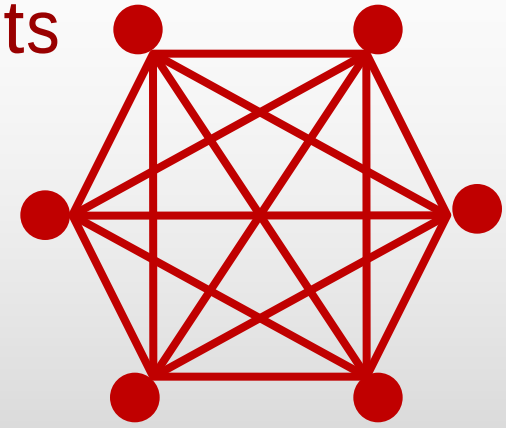

or

  - O(N) links + forwarding

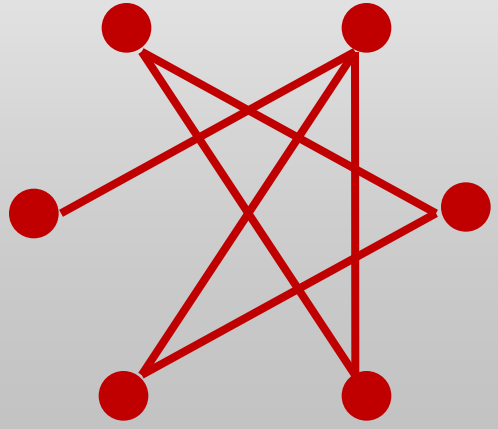

Information Sciences Institute

Virtualization leads to recursion

- N parties want to group in arbitrary, dynamic ways.

  ... such groups are inherently virtual

... and virtualization is inherently recursive

Information Sciences Institute

Recursion also supports layering and forwarding

- Layering (left)
  - Heterogeneity via O(N) translators
  - *Supported by successive recursive discovery*
- Forwarding (right)
  - N^2 connectivity via O(N) links
  - *Supported by successive iterative discovery*

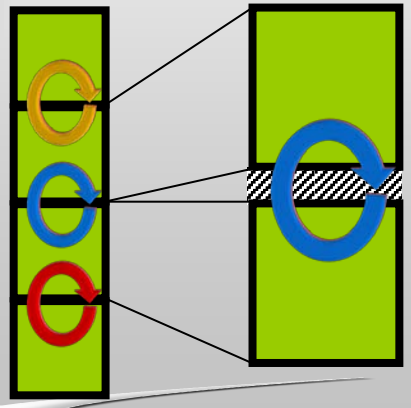

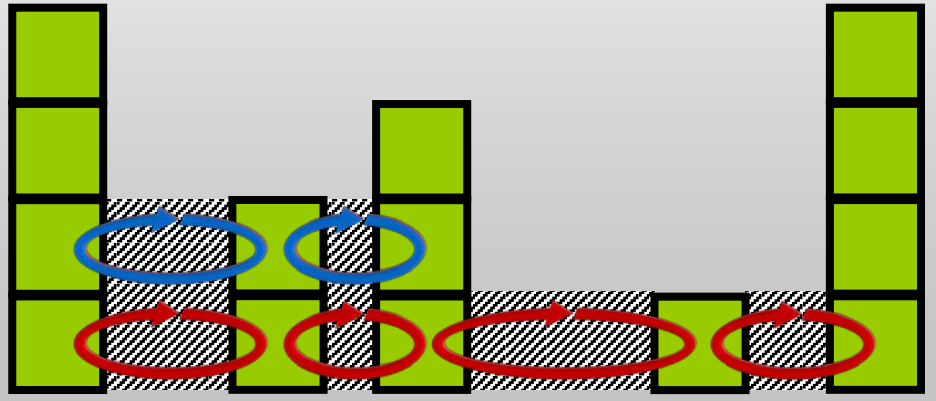

Information Sciences Institute

What makes this an architecture?

- General template (metaprotocol + MDCM)
  - Instantiates as different layers or forwarding
- Abstraction for virtualization
  - Tunnel as link
  - Partitioned router as virtual router
  - Partitioned host + internal router as virtual host
- Abstraction for recursion
  - Recursive router implemented as a network of vrouters with vhosts at the router interfaces

Information Sciences Institute

What does RNA enable?

- Integrate current architecture
  - 'stack' (IP, TCP) *vs.* 'glue' (ARP, DNS)
- Support needed improvements
  - Recursion (AS-level LISP, L3 BARP, L2 TRILL)
  - Revisitation
- Supports "old horses" natively
  - Dynamic 'dual-stack' (or more)

Recursive Internet Architecture

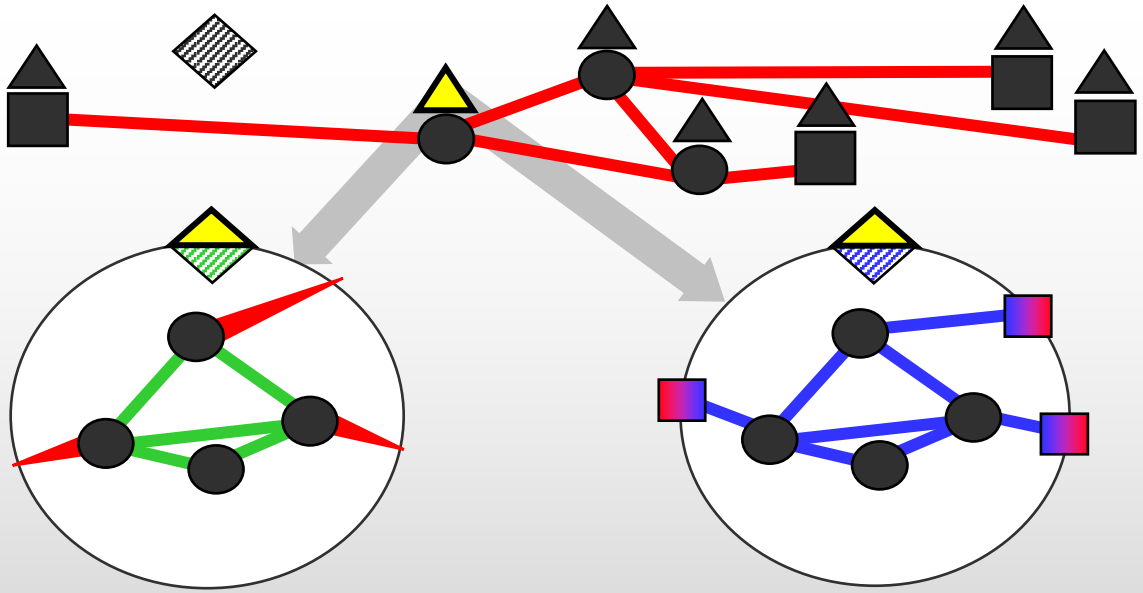

Control / deployment ***Network***

- Recursion as a router
  - L3 = BARP (X-Bone), LISP (IRTF)
  - L2 = Rbridges/TRILL

RNA Metaprotocol

- Template of basic protocol service:
  - Establish / refresh state
  - Encrypt / decrypt message
  - Apply filtering
  - Pace output via flow control
  - Pace input to allow reordering
  - Multiplex/demultiplex
    - includes switching/forwarding

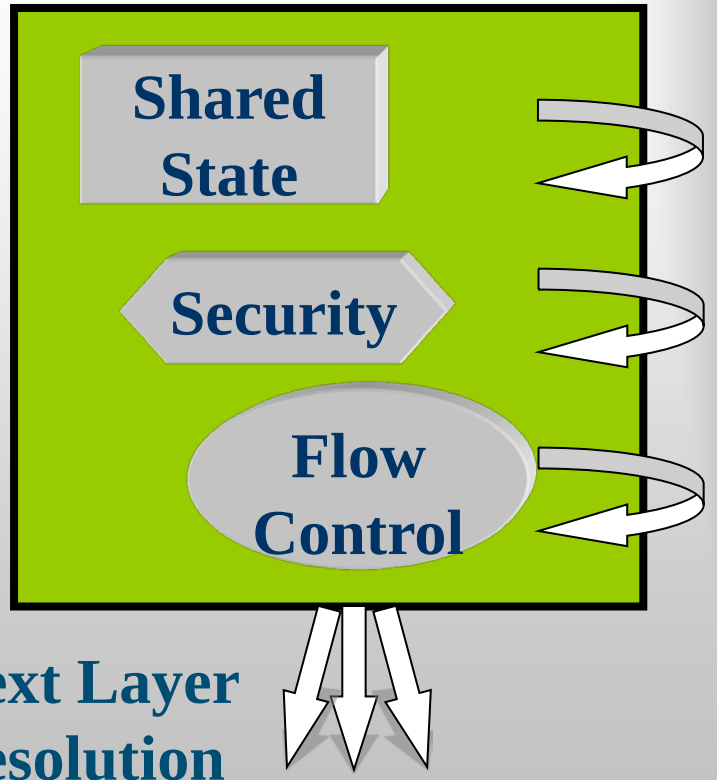

MDCM

Structured template w/plug-in functions

- Layer address translate/resolution
  - ARP, IP forwarding lookup
  - BARP/LISP/TRILL lookup
- Layer alternates selection
  - IPv4/IPv6, TCP/SCTP/DCCP/UDP
- Iterative forwarding
  - IP hop-by-hop, DNS recursive queries

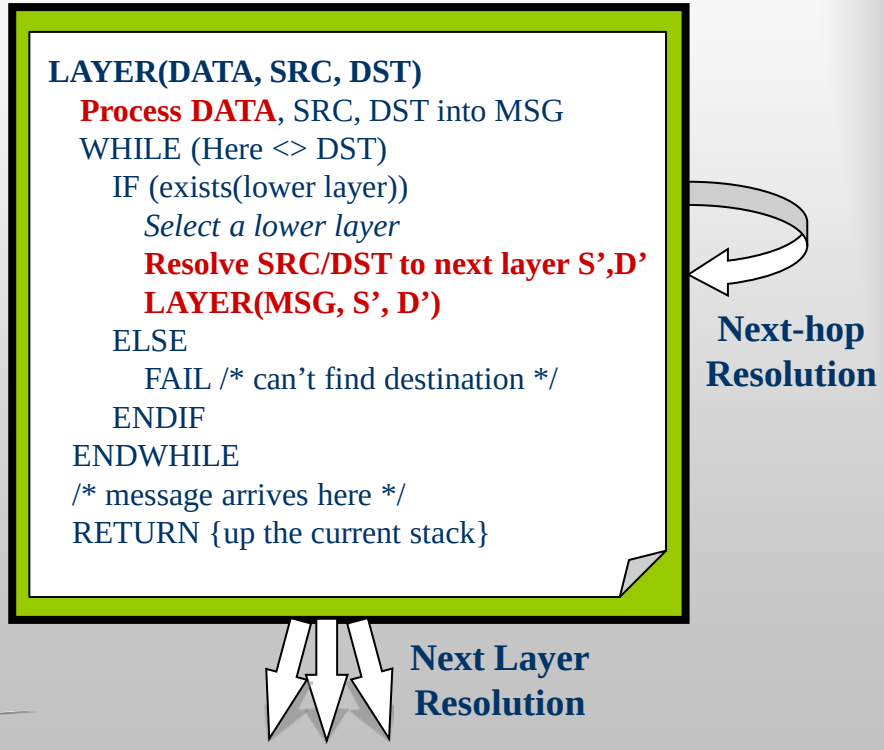

Information Sciences Institute

Related Work

- Recursion in networking
  - X-Bone/Virtual Nets, Spawning Nets, TRILL, Network IPC, LISP
  - *RNA natively includes resolution and discovery*
- Protocol environments
  - Modular systems: Click, x-Kernel, Netgraph, Flexible Stacks
  - Template models: RBA, MDCM
  - *RNA adds a constrained template with structured services*
- Context-sensitive components
  - PEPs, Shims, intermediate overlay layers, etc.
  - *RNA incorporates this into the stack directly*
- Configurable über-protocols
  - XTP, TP++, SCTP
  - *RNA makes every layer configurable, but keeps multiple layers.*

Information Sciences Institute

Conclusions

- Virtualization requires recursion
- Recursion supports layering
- Recursion supports forwarding

One recurrence to bind them all…

- ***Recursion is a native network property***
  - Integrates and virtualization, forwarding and layering **in a single mechanism**

Information Sciences Institute

Internet Architecture

Accused of ossification, but:

- Ossification = stability
- Flexibility is abundant:
  - Shim layers:
    - HIP, SHIM6, IPsec, TLS
  - Muxing layers:
    - SCTP, RDDP, BEEP
  - Connections:
    - MPLS, GRE, IKE, BEEP, SCTP
  - Virtualization:
    - L2VPN, L3VPN/X-Bone/RON/Detour, L7-DHTs

Net Arch - Assumptions

- **Internet-Compliant Architecture**
  - Hosts add/delete headers
  - Routers transit (constant # headers)
- **Supports New Capabilities**
  - Concurrence (multiprocessing)
  - Revisitation (multiple roles in one net)
  - Recursion (to hide topology and/or mgt.)

Information Sciences Institute

Virtual Networks

- ***Internet-like***
  - Internet = routers + hosts + links
  - VIs = VRs + VHs + tunnels
  - Full architecture (vs. VPNs, PP-VPNs, etc.)
- ***All-Virtual***
  - Supports VNs on VNs
  - "Reality" is undecidable
- ***Recursion-as-router***
  - Some of VRs are VI networks
- ***See Globecom 1998 (running code 2000)***
  - 15 layers deep, 800 wide, app. deploy, P2P integration

Information Sciences Institute

Recursion requires new layers – where? Why?

- Wedge between (IPsec, left) or replicate (virtualization, right)

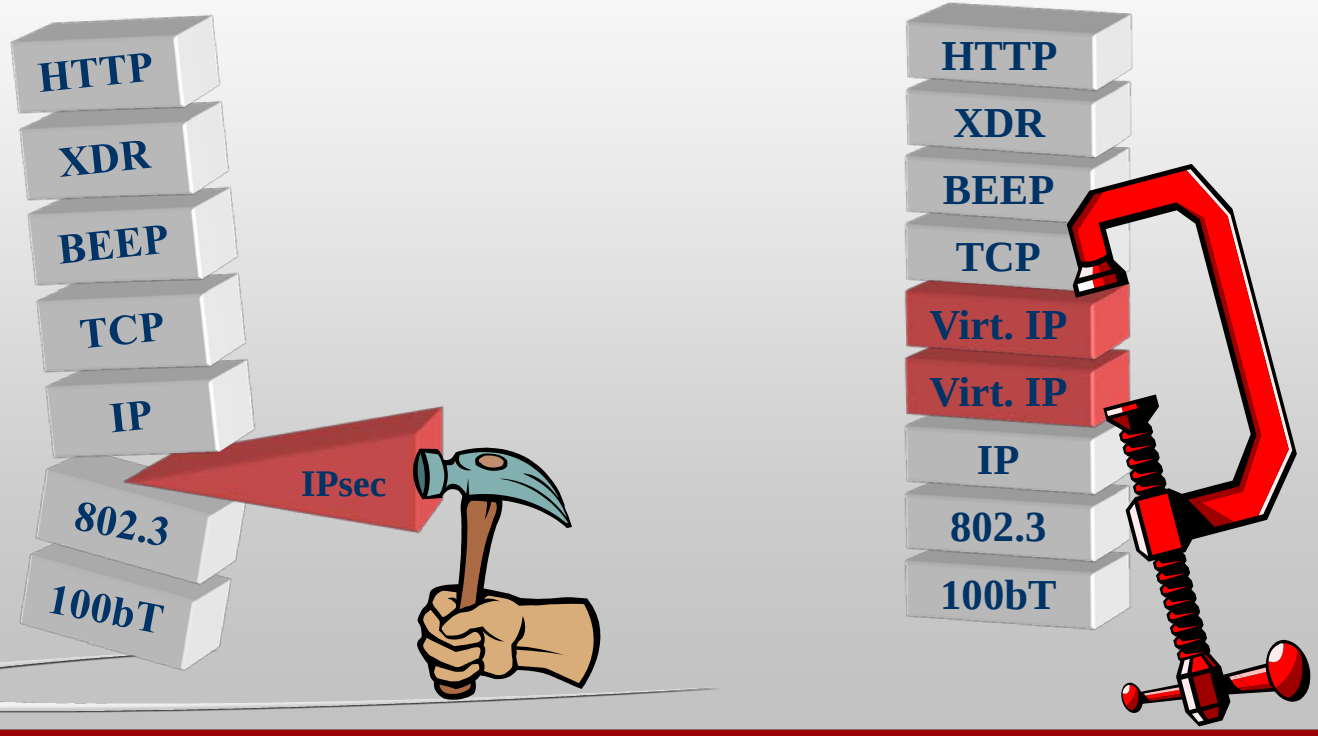

RNA Stack (2006)

- One MP, many instances
  - Needed layers, with needed services
  - Layers limit scope, enable context sensitivity
  - Scope defined by reach, layer above, layer below

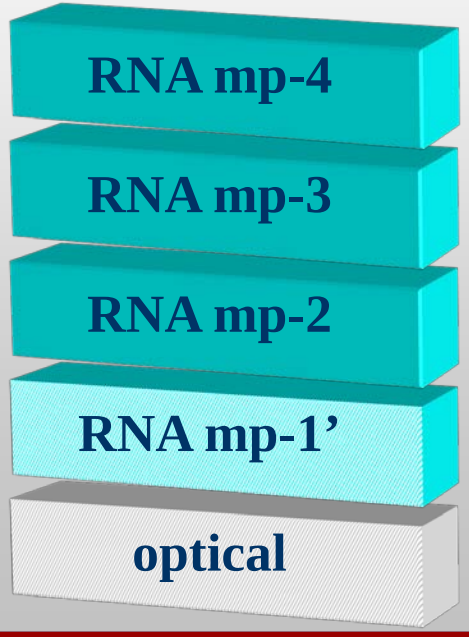

Information Sciences Institute

Click Implementation