



# Recursion in Networking

---

Joe Touch

Postel Center Director

USC/ISI

Research Assoc. Prof.

USC CS and EE/Systems Depts.



USC **Viterbi**  
School of Engineering



# Preview

---

- Recursion is a fundamental network property
  - Not just a software engineering artifact
  - Enables new capabilities for fault tolerance, DOS protection, and routing isolation
  - Unifies “planes” – data, control, mgt, security
  - Unifies forwarding, layering, and resolution
  - May hint to a broader science of networking



# Outline

---

- Recursion is key
- Background on X-Bone VNs
- RNA
  - Intro.
  - Implementation issues
  - Related work
- Summary



# Recursion is key

---



# What makes an architecture new?

---

- "Shaking the Hourglass"
  - All exchanges are 1 packet
  - Collosograms  $>$  RTT\*delay
  - No LANs? (L2 is only pt-pt)
- What defines success?
  - Fixing what's 'broken'
  - Doing something new/different
  - The Internet / circuits as a degenerate case



# 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



# 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



# Shannon Channel

- Two preselected parties
  - Homogenous endpoints



- Unidirectional channel
  - Preselected sender, preselected receiver





# What is communication?

---

- Shannon: shared bits
  - Between fixed endpoints, known *a priori*
- Shared bits between two parties
  - How do we find the party to talk to?



# What SCs Ignore

---

- What if you're not directly connected?
  - A) multihop
  - B) multilayer
- Why are multihop/multilayer interesting?
  - Scalable = multihop
  - Ubiquitous = multilayer
  - I.e., all scalable, ubiquitous comms!



# 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





# Layering leads to resolution

---

- IDs are local to a layer
  - Whether names, paths, locations
- Need to resolve IDs between layers
  - Google, DNS, ARP, LISP encap tables




# Interaction leads to forwarding

---


- N parties need

- $N^2$  circuits




*or*

- $O(N)$  links + forwarding



# Virtualization leads to recursion

- N parties want to group in arbitrary, dynamic ways.
  - ... such groups are inherently virtual
  - ... and virtualization is inherently recursive





**Control / deployment**      **Network**



# Recursion unifies layering, forwarding, & resolution

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







# What makes this an architecture?

---

- Abstraction for virtualization
  - Tunnel as link
  - Partitioned router as virtual router
  - Partitioned host + internal router as virtual host
- Abstractions for recursion
  - Recursive router implemented as a network of vrouters with vhosts at the router interfaces
  - Recursion within the protocol stack
- General templates (metaprotocol, ID tree)
  - Instantiates as different layers or forwarding



# X-Bone Virtual Nets

---




# Definition: Virtual Networks

---

- A network composed of:
  - Virtual routers = transit (unchanged header/circuit)
  - Virtual hosts = source/sink (add/delete hdr/ckt)
  - Virtual links = tunnels (links inside an existing path)  
*NB: VC nets are mature, so this focuses on virtual packet nets*
- Defining properties:
  - Provide a network
  - Requires existing network but distinct from it
  - Does *not* reach anywhere new



# Example: Ring



Base

Ring V/ON

*map*



# Motivation

---

*Provide to net users what an OS/VM provides:*

- Protection:
  - Isolate test/new protocols
- Concurrency:
  - Share infrastructure for testbeds
  - Share infrastructure for different services
- Abstraction:
  - Provide simpler topology
  - Provide application-specific topology




# Other Benefits....

---

- Scale
  - Scalable deployment and management
- Automation
  - Safe, consistent configuration
- Emulation
  - Support tests and legacy systems
- On-the-fly modification
  - Tune & move without affecting hosts/routers


# Recursion-as-Router

- **Sub-overlays look like routers**
  - L3 version of *rbridges* (*IETF TRILL WG*)
  - Similar to *LISP*



# Recursion requires new layers - where? Why?

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







# Challenges of Layering

---

- Which to add...
  - IPv4/IPv6, TCP/DCCP/SCTP
- When to add...
  - Security, muxing, cong. control
- Real vs. virtual
  - What's the difference?



# RNA Intro.

---



# Motivation for RNA

---

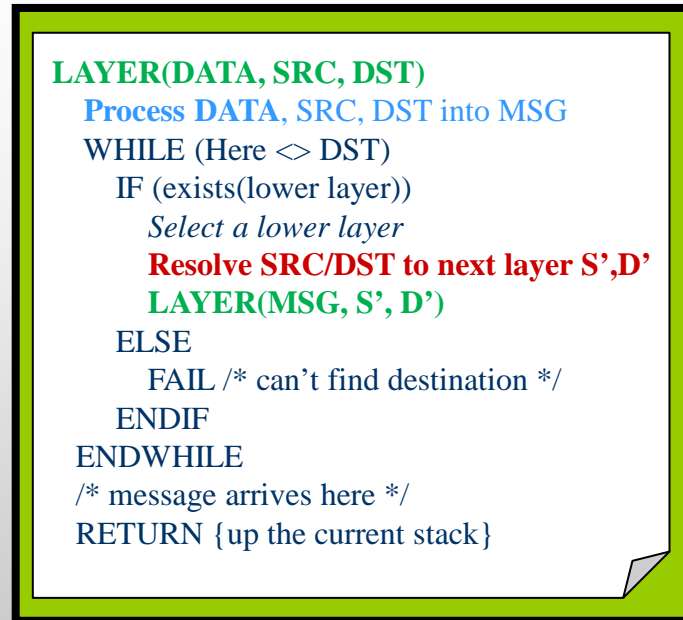
- Layers of a stack becoming more similar
  - Security, soft-state, pacing, retransmission
- Desire to support new capabilities
  - Interlayer cooperation, dynamic layer selection
- Desire to support emerging abstractions
  - Overlay layers don't map to 1-7
  - Support for recursive nodes (BARP, LISP, TRILL)

*Is layering more than a coding artifact?*



# One module to reuse

- “Resolve” unifies:
  - 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
- “Process data” unifies:
  - Shared state, security, management
  - Flow control, error control




Next-hop  
Resolution

Next Layer  
Resolution


# RNA Stack

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




# Retain layering

- One metaprotocol, many instances
  - Needed layers, with needed services
  - Layers limit scope, enable context sensitivity
  - Scope defined by reach, layer above, layer below
  - Resolution connects the layers (red/green)



# Scope defines a layer

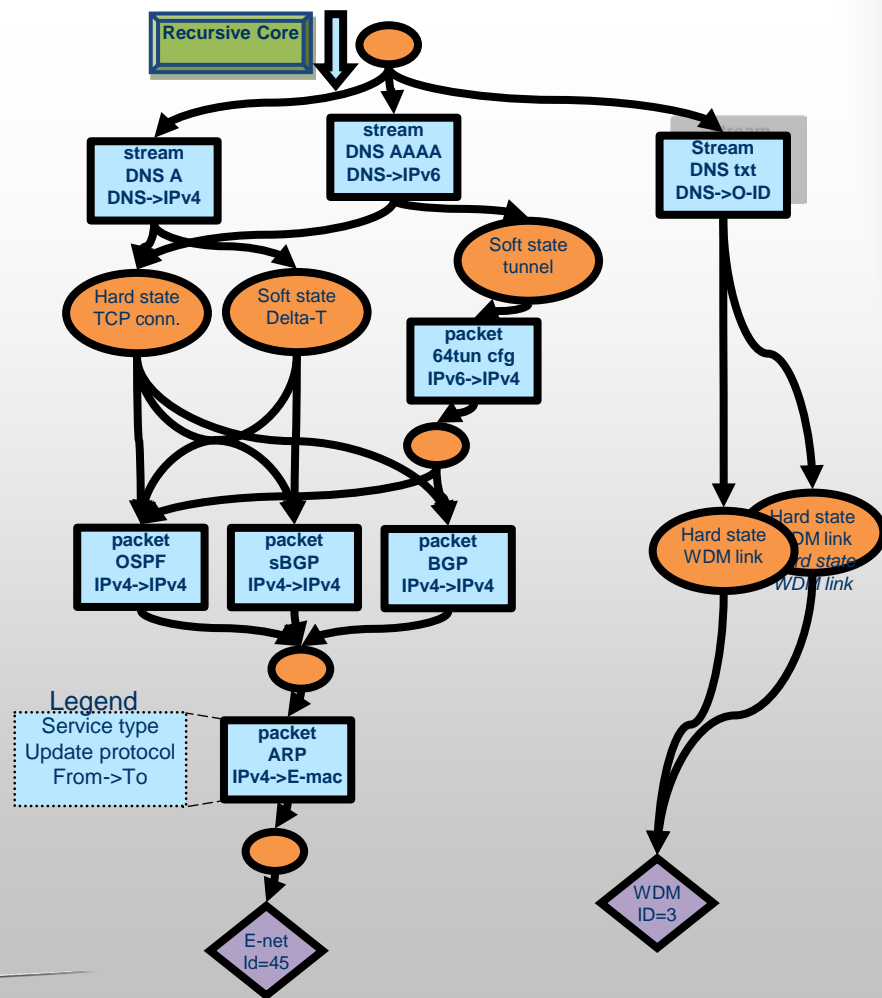
- Its endpoints
  - A “hop” @layer N = E2E extent of layer N-1
- The layer above
  - What services this layer provides
- The layer below
  - What services this layer requires
- E.g.: Shared state at diff. layers for diff. services
  - Application binding
  - Transport delivery
  - Net security



*The difference is scope*

# IDs constrain structure

- Tree of ID spaces
  - Link at resolvers
- State inbetween
  - Connections, provisioning
- Table management
  - ID use coordination
  - Routing
  - Resolution







# What makes this an architecture?

---

- Basic components
  - Metaprotocol + MDCM, ID space tree, etc.
  - 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



# What does RNA enable?

---

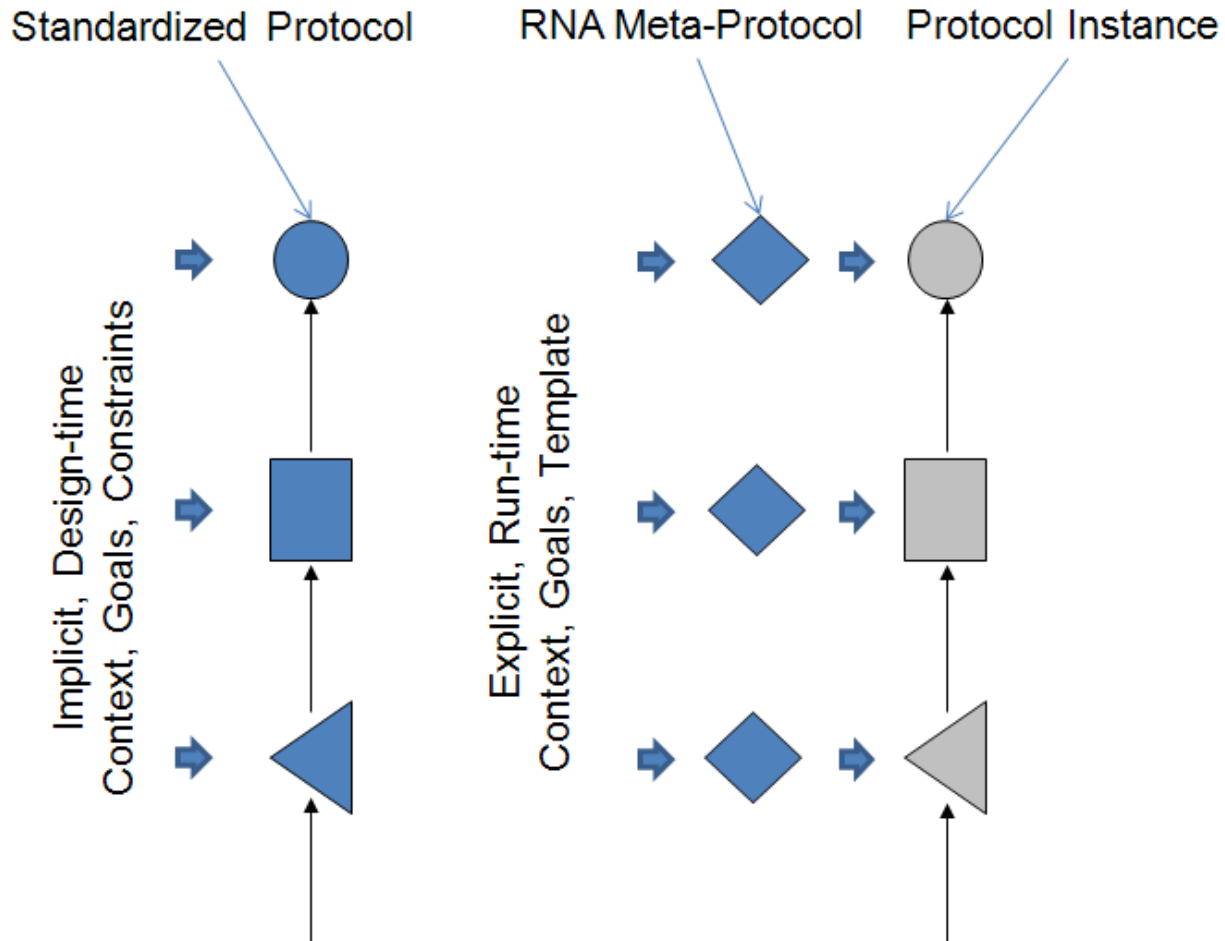
- Integrate current architecture
  - 'stack' (IP, TCP) *vs.* 'glue' (ARP, DNS)
  - Forwarding and layering
  - Unify data, control, management, security planes
- Support needed improvements
  - Recursion (AS-level LISP, L3 BARP, L2 TRILL)
  - Revisitation (X-Bone)
  - Concurrence (VPNs, multipath TCP)
- Supports "old horse" challenges natively
  - Dynamic 'dual-stack' (or more)



# Implementation Issues

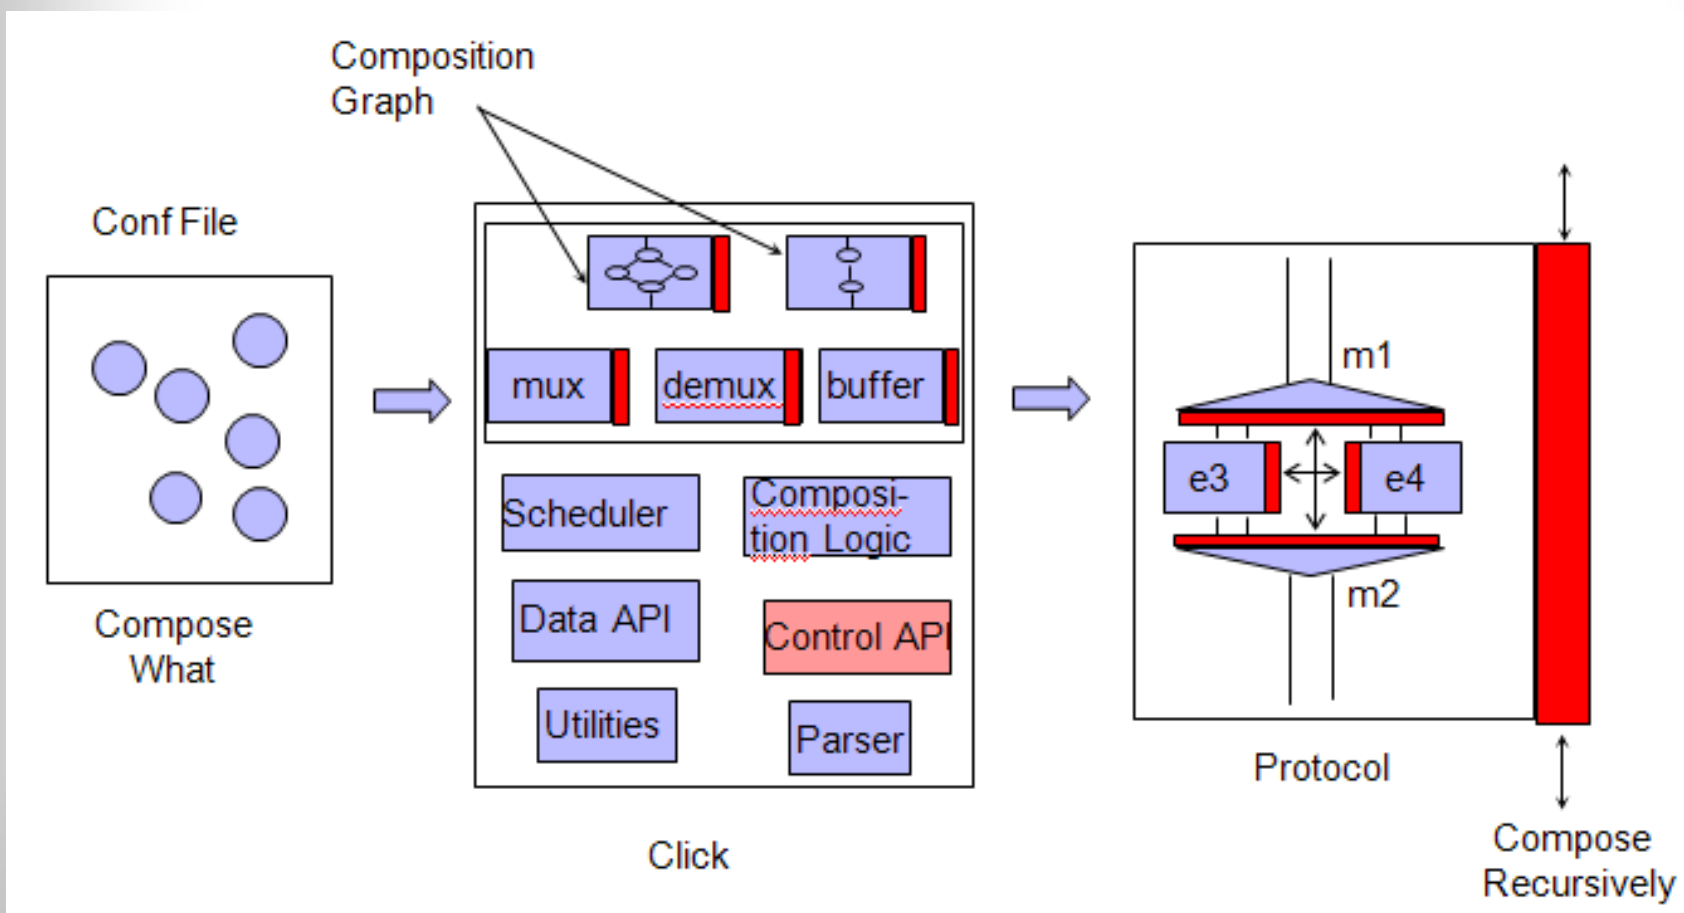
---

# Instantiation



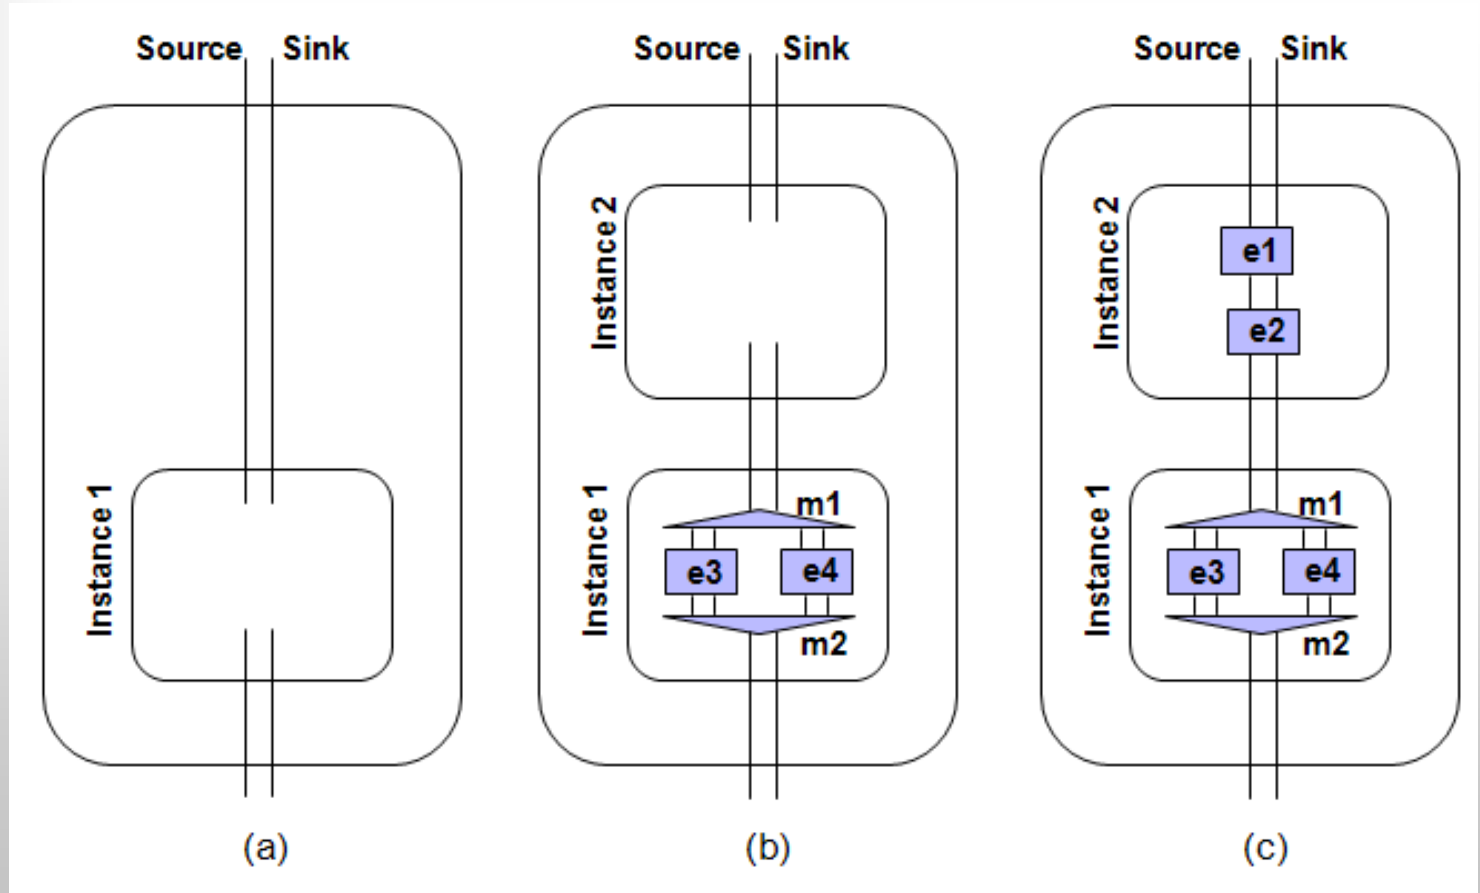


# Click Implementation



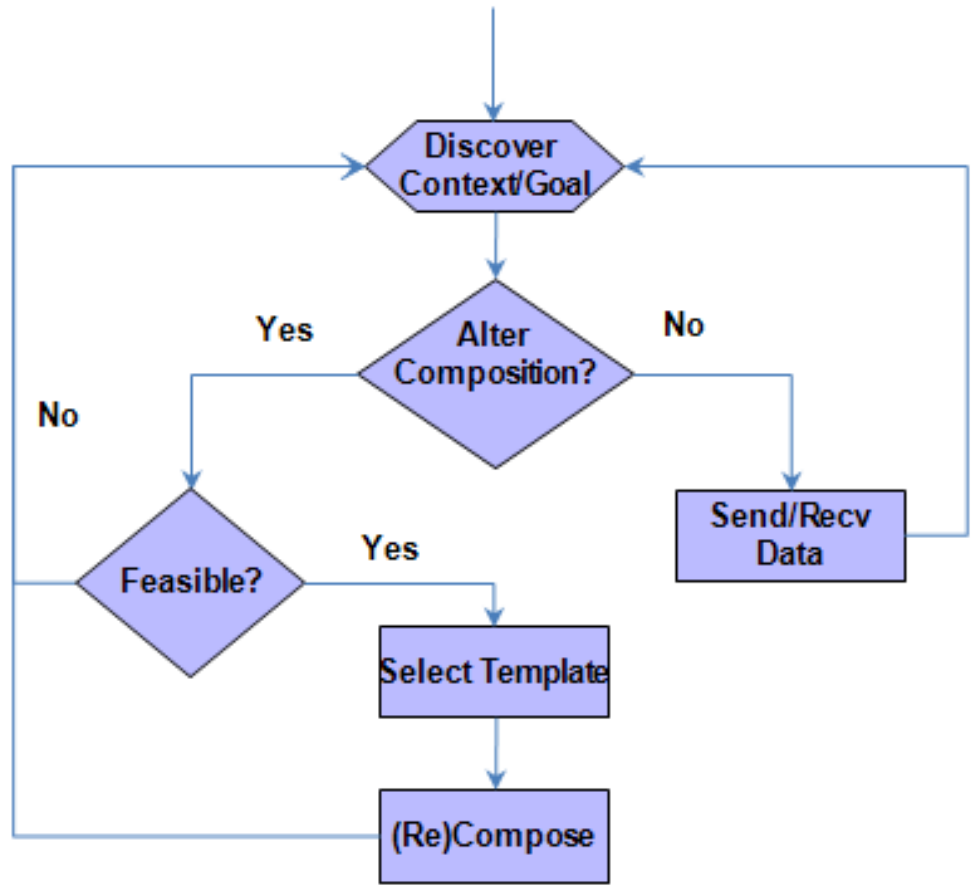


# Building a Stack





# Composition Process





# Other Components

---

- MP design
  - What's inside the "box"
  - Interlayer coordination
  - Context sensitivity, environment tuning
- Dynamic negotiation protocol
  - Cross-layer negotiation, IETF TAE
- Composeable/recursive extensions
  - Network management/SLAs
  - Security (user/infrastructure)
  - Non-comm services (storage, computation)
- Integrated optimization
  - Caching, precompute/prefetch
  - Pinning, dampening





# Related Work

---



# Related Work Summary

- 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.*



# RNA and Network IPC

---

- Similarities
  - Recursive protocol stack
  - Unified communication mechanism
  - Focus on process-to-process interaction
- Differences
  - RNA uses MDCM to define IPC as combining a Shannon-style channel with namespace coordination
  - RNA provides a detailed (and demonstrated) mechanism that achieves unification and recursion
  - RNA supports both recursion and forwarding in a single mechanism



# Conclusions

---



# Recursion uniquely enables...

---

- Integrates data, control, mgt, security
  - All are different ways of managing state inbetween resolutions
  - State can be shared – TCP RTT, NM liveness, BGP timers, etc. are all the same info.
- Integrates routing and resolution
  - Both are just ways to manage the tables
- Integrates provisioning and conn. mgt
  - Provisioning is at layer N is just a new connection at layer N-1



# Summary

---

- Recursion is an integral part of networking
  - Falls out of multiparty communication
- Recursion is a native part of layering
  - Whether IP/ethernet, or LISP (IP/IP), or TRILL (ether/ether)
- Recursion allows us to keep layering
  - Layering is critical to constrain scope



# Conclusions

---

- Virtualization requires recursion
- Recursion supports layering, forwarding, resolution
- Recursion integrates data, control, mgt, security

*One recurrence to bind them all...*

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



# Credits

---

- ID tree and related issues
  - Christos Papadopolous and Dan Massey CSU
- MDCM
  - Yu-Shun Wang
- RNA
  - Yu-Shun Wang, Venkata Pingali
- Naming unification
  - Venkata Pingali
- Virtual networking (X-Bone *et al.*)
  - Lars Eggert, Yu-Shun Wang, Greg Finn, Steve Hotz, Oscar Ardaiz-Villanueava, Norihito Fujita