



Resolving the Transport "Tussle"

Recursive InterNetwork Architecture

@
Computer Science
Boston U.



<http://csr.bu.edu/rina>

1


What this is (NOT) about

- ❑ NOT much about specific protocols, algorithms, interfaces, implementation
- ❑ It's about architecture, i.e., objects and how they relate to each other
- ❑ It's based on the **IPC model**, not a specific implementation

"Networking is inter-process communication"
--Robert Metcalfe '72

2

What's wrong with today's Internet?




- ❑ The **new brave world**
 - Larger scale, **more diverse** technologies
 - **New services**: content-driven, context-aware, mobile, socially-driven, secure, profitable, ...
- ❑ Custom **point-solutions**: No or little "science"
- ❑ Lots of problems: Denial-of-service attacks, unpredictable performance, hard to manage, ...

3

Questions?


- ❑ Is the Internet's architecture fundamentally broken that we need to “clean slate”?
 - Yes

- ❑ Can we find a new architecture that is complete, yet minimal? If so, what is it?
 - RINA? 

- ❑ Can we transition to it without requiring everyone to adopt it?
 - Yes

4


Internet's view: one big, flat, open net



- ❑ There's **no building block**
- ❑ The “hour-glass” model imposed a least common denominator
- ❑ Either didn't name what was needed or named the wrong things (i.e., interfaces)
- ❑ We exposed addresses to applications
- ❑ We hacked in “middleboxes”
- ❑ ...

5


Ex1: Bad Addressing & Routing



- ❑ Naming “interfaces” – i.e., (early) binding (of) objects to their attributes (Point-of-Attachment addresses) – makes it hard to deal with multihoming and mobility
 - Mobility is a dynamic form of multihoming
- ❑ Destination application process identified by a **well-known (static)** port number

6

Ex2: Ad hoc Scalability & Security



- ❑ Network Address Translator aggregates **private** addresses
- ❑ NAT acts as **firewall**
 - preventing attacks on **private** addresses & ports
- ❑ But, **hard to coordinate communication** across domains when we want to


7

Our Solution: divide-and-conquer

- ❑ Application processes communicate over a Distributed IPC Facility (DIF)
- ❑ DIF management is hidden from applications
 - ➔ **better security**
- ❑ IPC processes are application processes of lower IPC facilities
- ❑ **Recurse** as needed
 - ➔ **better management & scalability**
- ❑ Well-defined service interfaces
 - ➔ **predictable service quality**
- Applications ask for a *location-independent service*
- The underlying IPC layer maps it to a *location-dependent node name, i.e. address*

8


Recursive Architecture based on IPC



DIF = Distributed IPC Facility (locus of shared state=scope)
 Policies are tailored to scope of DIF

9


What Goes into a DIF?



- All what is needed to manage a "private" (overlay) network
 - A DIF integrates routing, transport and management
 - In TCP/IP, we artificially isolated functions of same IPC / scope

10


What Goes into a DIF?



- Processing at 3 timescales, decoupled by either a **State Vector** or a **Resource Information Base**
 - **IPC Transfer** actually moves the data (*tightly coupled* mechanisms)
 - **IPC Control** (optional) for error, flow control, etc. (*loosely coupled*)
 - Good we split TCP, but **we split it in the wrong direction!**
 - **IPC Management** for routing, resource allocation, locating applications, access control, monitoring lower layer, etc.
 - We need **only one "stateless" Common Application Protocol** to access objects: CREATE, DELETE, UPDATE, ...

11


RINA allows scoping of services



- The **DIF is the building block and can be composed**
- Good we split TCP, but **we split TCP in the wrong direction!**
- E2E (end-to-end principle) is not relevant
 - Each DIF layer provides (transport) service / QoS over its scope
- IPv6 is/was a waste of time! A single ubiquitous address space is unnecessary
 - Have many levels without too many addresses within a DIF layer

12


RINA: Good Addressing – private mgmt



- Destination application is identified by “name”
- Each IPC Layer (DIF) is privately managed
 - It assigns **private** node addresses to IPC processes
 - It internally maps **app/service name** to **node address**
 - Need a global namespace, but not address space
 - Destination application process is assigned a port number **dynamically**

13

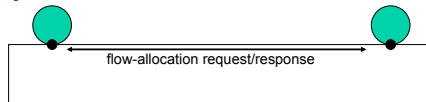
RINA: Good Addressing - late binding




- Addressing is **relative**: node address is name for lower IPC Layer, and point-of-attachment (PoA) for higher IPC Layer
- Late binding of **node name** to **PoA** address
- A machine subscribes to different DIF layers

14

Only one Data Transfer Protocol




- In RINA, service is accessed by its application name
- Port allocation and access control **decoupled from** data transfer
- At each end, port and conn ID are allocated **dynamically** and bound to each other by management, in a **hard-state** fashion




15

Only one Data Transfer Protocol (2)




- Once allocated, Data Transfer can start following **Delta-t** [Watson' 81], a **soft-state** protocol
 - Flows without data transfer control are UDP-like. Flows without reliability requirement do not ACK. Different policies support different requirements
 - If there is a long idle period, conn state is discarded, but ports remain
 - Conn IDs can be changed during data transfer and bound to same ports



16


Where security goes ...

- Authentication and encryption are applied recursively – no “shim” sublayers



17

RINA: Better Scalability & Security – secure containers



- Nothing more than applications establishing communication
 - Authenticating that A is a valid member of the DIF
 - Initializing it with current DIF information
 - Assigning it an internal address for use in coordinating IPC
 - This is **enrollment**, i.e. explicit negotiation to join DIF (access control)
 - RINA **decouples authentication from connection management and integrity/confidentiality**

18


Port Scanning Attacks

- Goal: first step for an attack, explore “open” ports
- In RINA, requesting applications never see addresses nor conn IDs
 - No well-known ports
 - Ports, dynamically allocated, are not part of conn IDs
 - Service requested by application name
- Traditional port scanning attacks not possible
- Scanning application names is much more difficult
- Attacker has to join the DIF too
 - For the sake of comparison, we assume the attacker overcame this hurdle!

19

Connection Opening Attacks: TCP/IP

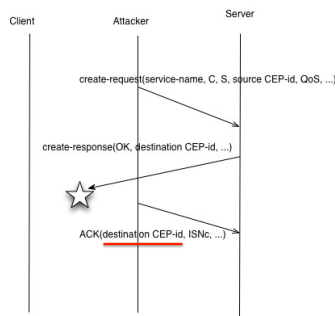
- Attacker has to guess server's Initial Sequence Number (ISN)
- Given 32-bit sequence number, **2³² possibilities**



20

Connection Opening Attacks: RINA

- Attacker has to guess destination CEP-id
- Given 16-bit CEP-ids, **2¹⁶ possibilities**
- Akin to port-scanning attacks, which raise more suspicion
- Client can use any ISN



21

Data Transfer Attacks

TCP/IP

- Goal is to inject a legitimate packet, e.g. TCP "reset"
- Attacker has to guess source port and SN within transmission window
- Given 16-bit port numbers and 16-bit max window, $2^{16} * 2^{(32-16)-16} = 2^{32}$ guesses

RINA

- **Right before data transfer starts**
- Attacker has to guess conn IDs and QoS ID
- Given 8-bit QoS ID, $2^{(16+16+8)} = 2^{40}$ guesses
- **During data transfer**
- Attacker has to also guess SN, so $2^{(40+16)} = 2^{56}$ guesses
- **Note: RINA can change conn IDs on the fly**

22

Attacking the reassembly of TCP segment

- Attack by inserting malicious data into IP fragment carrying part of TCP payload
- **Not possible in RINA**
- Transport and relaying are integrated in each DIF layer
- Fragmentation/reassembly is done once as data enters/leaves the DIF layer

23


Good Design leads to Better Security

- In RINA, requesting apps never see addresses nor conn IDs
 - traditional port scanning attacks not possible
- Underlying IPC processes must be authenticated to join DIF
 - only "insider" attacks possible
 - a hurdle that is not present in TCP/IP networks

24


Good Design leads to Better Security (2)

- ❑ Conn IDs are allocated dynamically, so they are hard to guess
- ❑ State of data transfer is soft, so there aren't explicit control messages to fabricate
 - **Note:** Delta-t was developed in the 80's with NO consideration for security!



25


RINA: Good Routing



- ❑ Back to naming-addressing basics [Saltzer '82]
 - Service name (location-independent) →
 - node name (location-dependent) →
 - PoA address (path-dependent) → path
- ❑ We clearly distinguish the last 2 mappings
- ❑ **Route:** sequence of node names (addresses)
- ❑ **Late binding** of next-hop's node name to PoA at lower DIF level

26

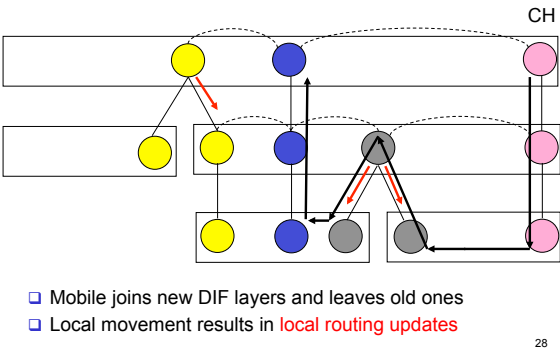
Mobility is Inherent



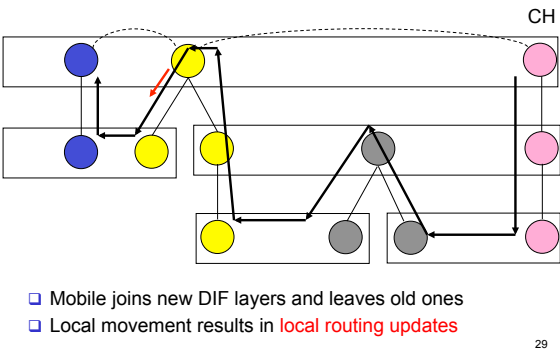
- ❑ Mobile joins new DIF layers and leaves old ones
- ❑ Local movement results in **local routing updates**

27

Mobility is Inherent

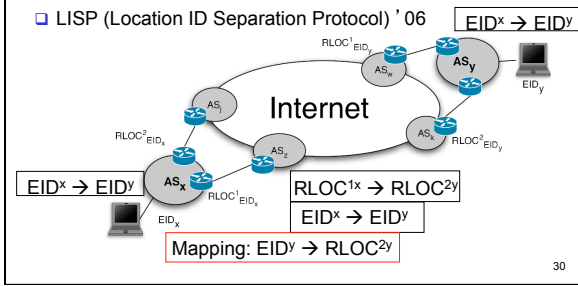


Mobility is Inherent




Compare to loc/id split (1)

- Basis of solutions to the multihoming issue
- Claim:** the IP address semantics are overloaded as both location and identifier
- LISP (Location ID Separation Protocol) '06



Compare to loc/id split (2)

- Ingress Border Router maps ID to loc, which is the location of destination Egress BR
- Problem:** loc is path-dependent, does not name the ultimate destination




31

LISP vs. RINA vs. ...

- Total Cost per loc / interface change = $\text{Cost of Loc / Routing Update} + \rho [P_{\text{cons}} * \text{DeliveryCost} + (1 - P_{\text{cons}}) * \text{InconsistencyCost}]$
- ρ : expected packets per loc change
- P_{cons} : probability of no loc change since last pkt delivery
- RINA's routing modeled over a binary tree of IPC Layers: update at top level involves route propagation over the whole network diameter D ; update at leaf involves route propagation over $D/2^h$, h is tree height

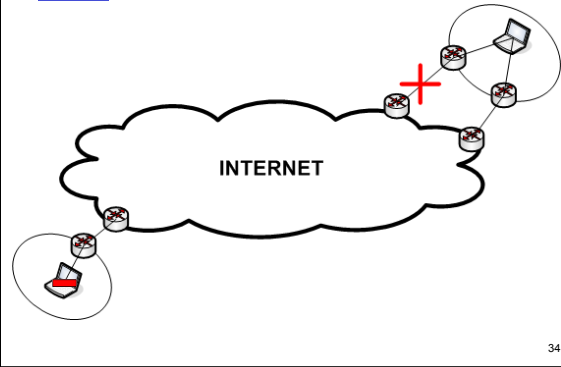
32

LISP



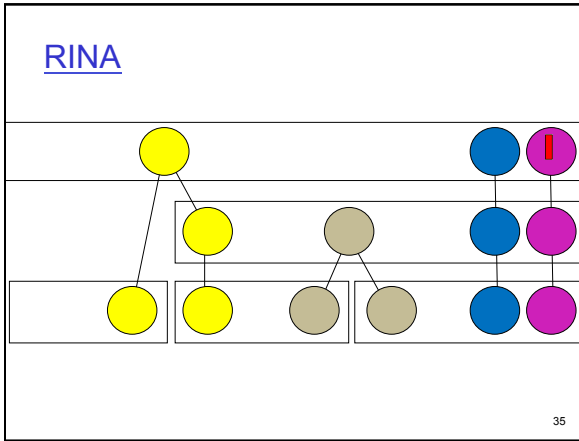
33

LISP




34

RINA




35

RINA




36

RINA




MobileIP




LISP vs. RINA vs. ...

8x8 Grid Topology
 RINA uses 5 IPC levels; on average, 3 levels get affected per move



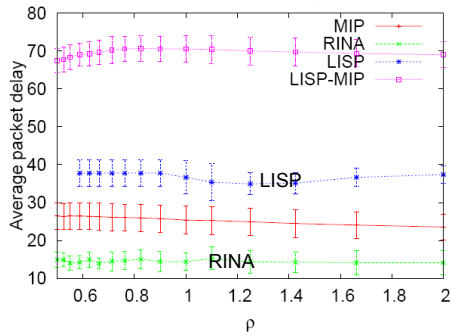
Simulation: Packet Delivery Ratio

- BRITE generated 2-level topology
- Average path length 14 hops
- Random walk mobility model
- Download BRITE from www.cs.bu.edu/brite



40

Simulation: Packet Delay



41

Bottom Line: RINA is less costly

- RINA inherently limits the scope of location update & inconsistency
- RINA uses “direct” routing to destination node

42

Adoptability

- ❑ ISPs get into the IPC business and compete with host providers
 - Provide transport services everywhere
- ❑ A user joins any IPC network she chooses
- ❑ All IPC networks are **private**
- ❑ We could still have a public network with weak security properties, i.e., the current Internet
- ❑ Many IPC providers can join forces and compete with other groups

43



More @
<http://csr.bu.edu/rina>

44
