

On Supporting Mobility and Multihoming in Recursive Internet Architectures

Vatche Ishakian[☆], Joseph Akinwumi, Flavio Esposito, Ibrahim Matta

*Computer Science Department
Boston University
111 Cummington Street
Boston, MA, 02215, USA*

Abstract

As the Internet has evolved and grown, an increasing number of nodes (hosts or autonomous systems) have become multihomed, *i.e.*, a node is connected to more than one network. Mobility can be viewed as a special case of multihoming—as a node moves, it unsubscribes from one network and subscribes to another, which is akin to one interface becoming inactive and another active. The current Internet architecture has been facing significant challenges in effectively dealing with multihoming (and consequently mobility), which has led to the emergence of several custom point-solutions. The Recursive InterNetwork Architecture (RINA) was recently proposed as a clean-slate solution to the current problems of the Internet. In this paper, we present a specification of the process of ROuting in Recursive Architectures (RORA). We also perform an average-case cost analysis to compare the multihoming / mobility support of RINA, against that of other approaches such as LISP and Mobile-IP. Extensive experimental results confirm the premise that the RINA architecture and its RORA routing approach are inherently better suited for supporting mobility and multihoming.

Keywords:

Routing, mobility support, multihoming, loc/id split, recursive architecture, performance analysis, simulation.

1. Introduction

Support for multihoming and mobility was not a primary goal in the original design of the Internet. As a result, the Internet's naming and addressing architecture is incomplete. Specifically, the address of a multihomed host specifies a particular interface (connection), rather than the node itself. Because routing is done based on this interface, *i.e.*, Internet Protocol (IP) address, if this active interface goes down, it is costly to switch to another operational interface.

There have been several attempts to fix this addressing problem, including the Location ID Separation Protocol (LISP)—currently being tested at Cisco [1, 2]—and Mobile-IP [3]. The basic idea behind LISP is to assign the multihomed node a provider-independent (location-independent) identifier (ID). A border router maps a destination ID to the node's location, which is the address of another border router that is known to have a path to the node. Routing is then done from the source's border router to the destination's border router. If the latter (node's location) changes due to path failure or mobility, it becomes costly to propagate that change over the whole Internet (to all possible source border routers).

Mobile-IP (MIP) allows a mobile host to seamlessly move from its home domain to a foreign location without losing connectivity. This is done by having a foreign agent (router) update the location of the mobile node at its home agent (router). Since mobility is a special (dynamic) form of multihoming, MIP can also be used to handle a change in the active interface (due to failure or re-routing) leading to a multihomed node, where a home agent directs traffic to the currently active (operational or "better") interface. However, this location update can be costly since it needs to propagate from the foreign agent to the home agent.

Note that both LISP and Mobile-IP (and combination thereof) help reduce the size of the routing tables at the core of the Internet, since several IDs can map to one location and hence be represented by one routing entry. Further elaboration on the benefits of LISP can be found in [4].

RINA [5] is a recently proposed Recursive InterNetwork Architecture. It uses the concept of Distributed Inter-process communication Facility (DIF) to divide communication processes into manageable scopes across network subsystems, which results in a reduced routing table size per DIF. RINA routes *hop-by-hop* based on the destination's node address, *not* its interface. At each hop, the next-hop node address is mapped to the (currently operational) interface to that next-hop node. This late binding of a node's address to its interface (path) allows RINA to effectively deal with interface changes due to multihoming or mobility. The cost of such late binding is relatively small since

[☆]Tel:+18572469206

Email addresses: visahak@cs.bu.edu (Vatche Ishakian[☆]),
akin@cs.bu.edu (Joseph Akinwumi), flavio@cs.bu.edu
(Flavio Esposito), matta@cs.bu.edu (Ibrahim Matta)

its scope is *local* to the routing “hop” that traverses the underlying DIF. By recursing the DIF structure to make the DIF scopes small enough, the cost of such *late* bindings (location updates) can be made arbitrarily small.

1.1. Our Contribution

We present a specification of the process of ROuting in Recursive Architectures (RORA) adopted in RINA, and highlight its inherent support for mobility and multihoming. We present a cost model to quantitatively assess the effectiveness of LISP, MIP, and RINA, in supporting multihoming / mobility. To the best of our knowledge, this paper presents a first cost comparison of these approaches. Our definition of “cost” captures both the average number of packets generated by a source node to a (multihomed or mobile) destination node, as well as the average path length from the source to the destination (as indication of delays or bandwidth usage). In our model, we compute the overall average cost for a *single* interface change experienced by the multihomed or mobile destination node. We validate our analytical model for mobility using simulation and for multihoming using trace-driven simulation based on real data from CAIDA [6].

1.2. Organization of the Paper

The rest of the paper is organized as follows: Section 2 reviews MIP, LISP, and RINA. Section 3 presents the RORA routing process. We present our general cost model in Section 4, and then we instantiate it for the various approaches. Section 5 presents numerical results for grid topologies. Section 6 evaluates the cost of supporting mobility using simulations, and Section 7 evaluates the cost of supporting multihoming using real packet traces from CAIDA [6]. Section 8 reviews related work and Section 9 concludes the paper.

2. Background

This section provides a basic background on the various architectures we study, namely MIP, LISP, and RINA—for more details, we refer the reader to references herein.

2.1. Mobile-IP

Mobile-IP (MIP) [3] has been mainly standardized to deal with the mobility of nodes. As mentioned earlier, since mobility is merely a (dynamic) form of multihoming, the MIP concept can also be used to deal with interface (path) change to a multihomed node.

In MIP, two basic mechanisms are identified: (1) a *discovery* mechanism, which allows a node to detect its new point-of-attachment, and (2) a *registration* mechanism, which allows a node to register itself with an agent that represents it at its home network.

Figure 1 shows a source node (SN) sending packets to a destination node (DN) in another Autonomous System (AS). The destination moves to a new AS and acquires a care-of-address at the Foreign Agent (FA). The FA then updates the corresponding Home Agent (HA) with DN’s new location.

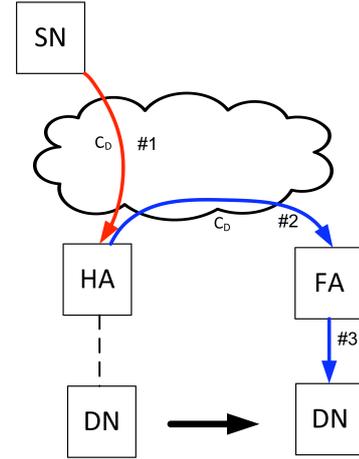


Figure 1: Mobile-IP Protocol.

The basic delivery process of data packets from a source node to a destination node is as follows (highlighted as sequence 1–3 in Figure 1):

1. The datagram is delivered to HA via standard routing.
2. The HA intercepts the datagram and tunnels it to the destination’s current location (care-of-address).
3. The FA at the current location intercepts the datagram and delivers it to the destination node.

Since its initial proposal, both for IPv4 and IPv6, different enhancements were suggested to deal with routing inefficiencies like Proxy Mobile IPv6 [7], Hierarchical MIPv6 (HMIPv6) [8], and Fast Handovers for Mobile IPv6 (FMIPv6) [9]. However these approaches incur extra communication overhead, and rely on special anchor routers with which mobile nodes register, which leads to QoS deteriorations as the anchor nodes get congested [7]. Furthermore, recent work [10] reports minor performance (packet loss, bandwidth) improvements under HMIPv6 and FMIPv6 compared to MIPv6. Thus in this work, we only consider the basic version of MIP.

2.2. LISP

The Locator/ID Separation Protocol (LISP), proposed by Farinacci *et al.* [11], separates the address space into end-systems’ identifiers (EID) for source and destination hosts, and routing locators (RLOCs) where border routers act as RLOCs for the end-systems inside their local domain. The mappings, referred to as *EID-to-RLOC* mappings, are stored in a Mapping Server (MS).

The basic delivery process of data packets from a source node (SN) to a destination node (DN) is as follows (highlighted as sequence 1–4 in Figure 2):

1. The data packet from the source is intercepted by the border router called Ingress Tunnel Router (ITR).
2. The source ITR performs a lookup query for a destination EID-to-RLOC mapping [12].
3. ITR transparently tunnels the data packet to the destination’s RLOC referred to as Egress Tunnel Router (ETR).

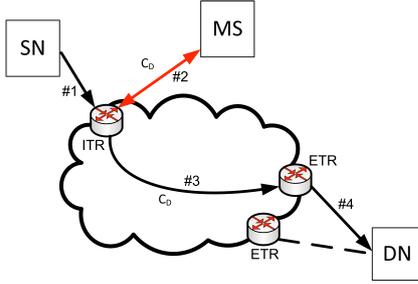


Figure 2: LISP Architecture.

4. Upon intercepting the packet, the destination’s ETR forwards the packet to the destination.

Upon failure of an active interface, a multihomed destination node sends an update to its ETR, which in turn updates the EID-to-RLOC MS. The sequence of messages is shown in Figure 3.

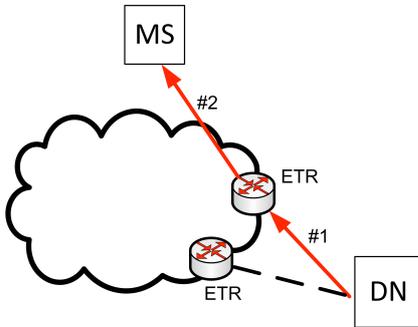


Figure 3: LISP cost of update.

Different variants of LISP only differ in how the EID-to-RLOC mapping is maintained [12]. The use of caching for lookup has also been recently explored in [13].

2.3. RINA

In RINA, application processes or services have globally unique names, and networking is viewed as distributed Inter-Process Communication (IPC). Each Distributed IPC Facility (DIF) integrates both transport and routing tasks along with other management tasks [5].

If an application process in RINA needs to communicate with another application process, it requests service from the underlying DIF. This DIF maps the destination application name to a node (process) address. A DIF in RINA can (recursively) provide transport services between source and destination application processes, using services of underlying (lower-level) DIFs.

Routing: The route to the destination node address (to which the destination application process is connected) is computed as a sequence of intermediate node addresses within the DIF. At each routing hop, the next-hop node address is in turn mapped (recursively) to a lower-level node address by the underlying DIF. This lower-level node address is viewed as the point-of-attachment of the higher-level node. Thus, RINA’s addresses are relative: a node address at a DIF level (N) is considered a

node name by a lower-level (N-1) DIF. At the (N-1)-DIF, this name needs to be mapped to a node (N-1)-address by the DIF’s directory service. Eventually, the node (process) address maps to a specific path (interface). This *late binding* to a specific interface (path) makes it easier for RINA to deal with mobility (and multihoming). If an active interface (path) to a node fails, RINA maps the (next-hop / destination) node address to another operational interface (path). The cost of such interface/location update is small because the update is only *local* to the routing hop—the next-hop / destination node address is mapped to the lower-level node address that resides within the operational lower-level DIF.

On the contrary, in the current Internet model, the interface address (*i.e.*, IP address) names both the node itself and the interface (path) to that node—this static binding makes mobility (and multihoming) difficult to manage.

RINA Example:

Figure 4 shows a source process “A” sending packets to a destination process “E” using the services of the underlying DIFs. Note that in RINA, a single system may have multiple processes which are members of different DIFs at different levels. The source process “A” and destination process “E” form a (high-level 5-node) DIF with a set of intermediate processes, which we call “intermediary” nodes (“B”, “C”, and “D”), such that the intermediary “B” can reach the destination process using two separate interfaces over two different underlying DIFs (numbered 2 and 5). Figure 4 also highlights the routing table of each of the nodes and the next hop to each of the destinations. This 5-node DIF can be thought of as an “overlay” (or private network) to which the source, destination, and intermediaries had subscribed. When a packet reaches the intermediary node (process) “B”, “B” forwards the packet based on the current best / operational path / interface (underlying DIF) leading to the destination process.

Remark: It is important to highlight the difference between how BGP [14] and RINA handle route / interface failures. In BGP, even if there is a specific path failure to a specific prefix (node), the protocol may still broadcast a path to the destination since it relies on advertising reachability to *aggregate* destination prefixes. On the other hand, RINA would handle such failures using *hop-by-hop* routing within the DIF of the destination process. For example, in Figure 4, if the communication between nodes “B” and “E” passing through node “D” and utilizing lower-level DIFs 2 and 3, goes down, then node “B” would locally adapt and start routing through node “C” utilizing lower-level DIFs 5 and 4. Thus, RINA provides finer grained control over routing to multihomed destinations.

Upon mobility, a node (process) may need to join or leave a DIF through a registration or unregistration procedure [5].

In the remaining sections we present the necessary inter and intra DIF operations, such as, registration, unregistration, node-address mapping, and routing, necessary to support mobility (and multihoming) in RINA. We then present an analytical model to compare the cost of supporting mobility/multihoming in RINA with that of other solutions, namely, MIP and LISP variants. Moreover, we validate our analysis using simulations.

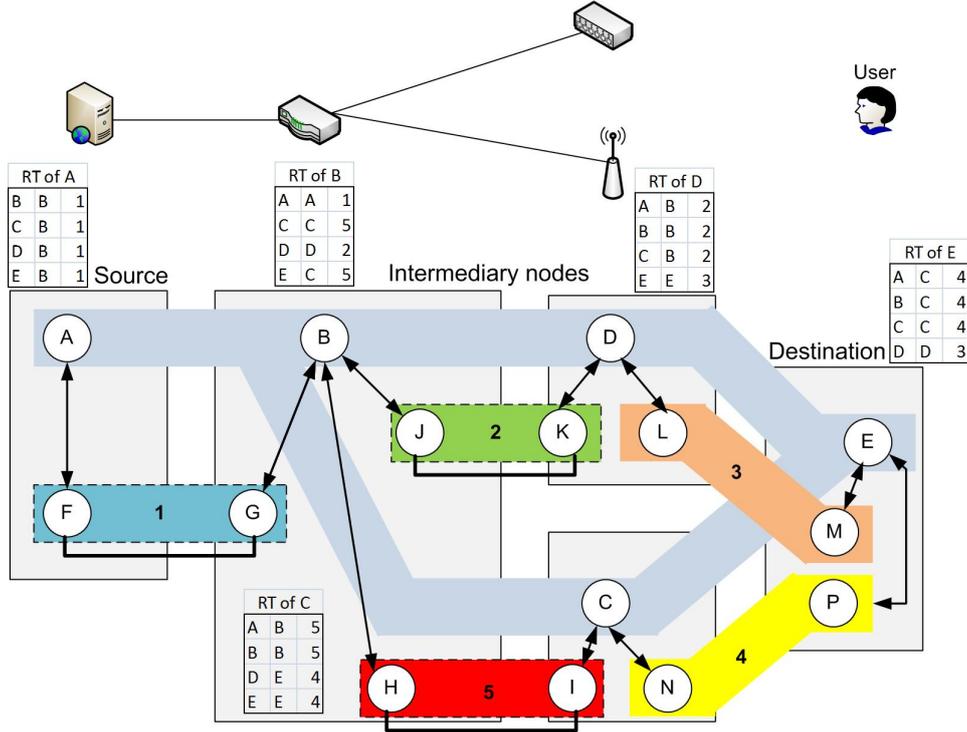


Figure 4: A RINA Network.

3. Protocol Specification

In this section we present the specifications of the process of Routing in Recursive Architectures (RORA) adopted in RINA. Naturally, the RORA functions are *recursive*, whereby each function invocation (instance), in reality, represents processing at a certain DIF level.

In our specifications, we assume the existence of a data structure, which we refer to as RIB (Resource Information Base), in each DIF. Among other information, the RIB contains a set of pairs (n, a) , where n is the application name, and a is its corresponding node address. Given RINA's relative addressing, n is viewed as the higher-level node name. For example, in Figure 4, within DIF 2, "D" is viewed as node name and "K" is its corresponding node address. Relatively speaking, "K" is also viewed as a point-of-attachment (PoA) for "D".

The RINA architecture consists of registration and unregistration phases to support the subscription and unsubscription of processes as they join and leave DIFs, respectively. RINA also requires translation / mapping functionalities and the actual recursive routing process.

3.1. Registration

In RINA, the registration process is done in a top-down fashion. As a node (process) moves from one DIF to another, it sends a registration request to a *registration* node located in that DIF (level l). After being authenticated (a mechanism outside the scope of our discussion)¹, the requester is assigned an address (line 5 in Procedure 1). The registration process is recursively propagated to the underlying DIFs to which lower-level

processes on the same machine subscribe (line 8). During the registration process, and after allocating an address to the node, the DIF updates its RIB (line 6). Once the registration process is complete, a registration response (shown in Procedure 2) is propagated upwards to each requesting DIF node (line 9).

Procedure 1 : Register_Request()

Require: n :String, l : long { n represents the node wanting to register, while l represents the DIF level}

- 1: **if** ($l == 0$) **then**
 - 2: Return {End of recursion; reached bottom DIF}
 - 3: **end if**
 - 4: **if** authenticate(n, l) **then**
 - 5: $a \leftarrow$ allocate_address(n)
 - 6: RIB.Update($n, a, \text{"ADD"}$)
 - 7: **end if**
 - 8: Register_Request($a, l-1$) {Recursively register in the lower DIF}
 - 9: Register_Response(n, a)
-

Procedure 2 : Register_Response()

Require: n :String, a :String {respond with the allocated address information to the requesting node}

Registration Example: We illustrate the registration process of RINA using the network shown in Figure 4. Assuming that the source wants to register. It starts by calling *Register_Request(source, l)* where l is the topmost level. Once it is authenticated, it will be assigned an address A , which will be registered recursively at layer $l-1$ and in turn assigned an

¹Security aspects of RINA are highlighted in [15].

Procedure 3 : UnRegister_Request():

Require: n :String, l : long
1: **if** ($l == 0$) **then**
2: **Return**
3: **end if**
4: **if** authenticate(n, l) **then**
5: $a \leftarrow$ allocate_address(n)
6: RIB_Update(n, a , "REMOVE")
7: **end if**
8: UnRegister_Request($a, l - 1$) {Recursively unregister in the lower DIF}
9: UnRegister_Response(n)

Procedure 4 : UnRegister_Response():

Require: n :String {respond with the unregistration confirmation to the requesting node}

address F .

3.2. Unregistration

Procedure 3 highlights the unregistration process in RINA which is similar to the registration process. The node n issues an *UnRegister_Request* to a node in a DIF. The node that receives the unregistration request removes n from the RIB (line 6), and subsequently issues a recursive unregistration request (line 8) to the lower DIFs. Once this process is complete, an unregistration response (shown in Procedure 4) is propagated upwards to each requesting DIF node (line 9).

Unregistration Example: We illustrate the unregistration process of RINA using the network shown in Figure 4. Assuming that the source wants to unregister. It starts by calling *UnRegister_Request(source, l)* where l is the topmost level. Once authenticated, its address will be removed and the unregistration will be processed recursively until all processes on the same machine (A and F) have unsubscribed from their respective DIF.

3.3. Mapping Functions

The mapping functions are a set of primitives that update and query a DIF's RIB data structure. They are called by a node to obtain a mapping between a name and its address.

Procedure 5 highlights the *RIB_Update* method, which is called whenever node information needs to be added to or removed from the RIB. $RIB \cup (n, a)$ and $RIB \setminus (n, a)$ should work as any database add and remove function, respectively.

Procedure 5 : RIB_Update():

Require: n :String, a :String, $type$:string
if $type ==$ "ADD" **then**
 $RIB \cup (n, a)$
else
 $RIB \setminus (n, a)$ {remove entry}
end if

The *Map_Request* function (Procedure 6) queries for node n in the RIB, and when found, its address is returned.

Procedure 6 : Map_Request()

Require: n :String
if $\exists n \in$ RIB **then**
 $a \leftarrow$ Find_address(n) {Returns the address of node n }
 Map_Response(n, a)
end if

Procedure 7 : Map_Response()

Require: n :String, a :String {replies to the requesting node with either successful connection to node address or error}

The *Map_Response(n, a)* (Procedure 7) function establishes a transport connection based on the node address, or signals an error (e.g. timeout, credential not found).

3.4. Recursive Routing

The recursive routing function (Procedure 8) is considered the core of RORA, and requires a source and a destination node, s and d . We denote by i any intermediate node, and m as the message to be delivered.

The function starts by recursively obtaining the source and next-hop/destination addresses using the *Map_Request* function (cf. Procedure 9). Based on the routing policy adopted, the next-hop (or intermediate node) to the destination is obtained by calling *getNextHop(d)*. Whenever the message reaches its destination process at the lowest DIF, the message is decapsulated and delivered to the higher level DIF directly using the function *Deliver_up* (Procedure 10). Whenever the message reaches its next-hop, it continues to be sent recursively down to its next-hop / destination. Eventually the message reaches its final destination and gets delivered to the destination application process using the function *Deliver_App*.

Routing Example: We illustrate the routing process of RINA using the network shown in Figure 4. Assuming that source A wants to send a message to E . It starts by calling *RRoute(A, E, m)*. The function finds out the next-hop node (B in this case), and sends the message down to the lower-level DIF, which maps the source and next-hop addresses to their lower-level addresses and calls *RRoute* recursively. In particular, the lower layer process F will forward the message to G , which in turn will deliver it to node (process) B . Node B will repeat the same process to send the message to D ², which in turn sends the message to node E .

Remark: It is important to note that RINA deals with scalability of routing by utilizing the following approaches: first, internal addresses in RINA are topological (location dependent), rather than today's provider/graph-based, and second, PoA addresses change more frequently in DIFs of lower rank (smaller scope). Thus upon user mobility, the scope of update is smaller than the scope of update in the current Internet. Third, communication across each hop is based on a soft-state mechanism [16]. This allows address mapping/resolution at each hop to be

²Depending on routing configurations during the forwarding stage, node B might decide to send the message to node C .

Procedure 8 : RRoute()**Require:** s :String, d :String, m :String

- 1: **if** “me” == d **then**
 - 2: Deliver_up(m) {If my address is the same as the destination, then deliver to the upper layer DIF}
 - 3: **else**
 - 4: $i = \text{getNextHop}(d)$
 - 5: $m' = (s, d) \parallel m$ {add header}
 - 6: send_down(“me”, i, m')
 - 7: **end if**
-

Procedure 9 : Send_down()**Require:** s :String, d :String, m :String

- 1: $s' = \text{Map_Request}(s)$
 - 2: $d' = \text{Map_Request}(d)$
 - 3: RRoute(s', d', m)
-

done per communication flow between a node and the next hop. So messages to the next hop are sent via the readily available established (transport) flow.

4. Cost Model

In this section we study the average (communication) cost of supporting mobility under MIP, LISP and RINA architectures. For the LISP architecture, we also analyze extended variants that employ caching for EID-to-RLOC mappings, or Mobile-IP running over basic LISP. Our cost model has similar characteristics to that of cellular systems [17, 18, 19], however, unlike cellular systems, our model captures additional characteristics and intricacies that are inherent to different proposed Internet architectures such as the multilevel hierarchy of LISP-MIP and recursive RINA DIFs.

4.1. Assumptions, Cost Definitions, and Parameters

We assume a single source-destination model where the source sends data packets according to a Poisson process. We analyze the average cost of managing a single interface (path) change to the destination due to the mobility of the destination node.

The cost of delivery of a single packet is denoted by C_D . The total cost per interface change, denoted by C_T , is a function of the *location lookup* cost (C_L), the *location update* cost (C_U), the *location inconsistency* cost (C_I), and C_D . *Location lookup* cost is defined only for LISP, to capture the cost of querying a mapping server for information about the destination’s RLOC given the destination’s EID. In computing the *location inconsistency* cost, we assume that packets delivered to the wrong location due to inconsistency of location / routing information, need to be delivered again.

We assume that for all approaches considered (LISP, MIP, and RINA), directory and routing processes have pre-established security mechanisms (*i.e.* common secret keys) that allow them to (1) encrypt and decrypt their routing/control messages for confidentiality, and (2) generate and piggyback authorization keys in routing/control messages for authentication and

Procedure 10 : Deliver_up()**Require:** m :String

- 1: **if** header(m) == ϕ **then**
 - 2: Deliver_App()
 - 3: return
 - 4: **end if**
 - 5: $(s, d) = \text{header}(m)$
 - 6: $m = m \setminus (s, d)$ {decapsulate}
 - 7: **if** “me” == d **then**
 - 8: Deliver_up(m)
 - 9: **end if**
 - 10: RRoute(s, d, m)
-

access control. Thus, our focus is on the cost of communicating such routing/control messages and its impact on the delivery performance of data messages in terms of packet delay and loss.

In our model, we assume that the inter-arrival times of data packets and the lifetime of the destination’s interface, each follows an exponential distribution, denoted by $f_p(t)$ and $f_m(t)$, respectively. We define the following two parameters:

- λ : the mean packet arrival rate, *i.e.*, $f_p(t) = \lambda e^{-\lambda t}$.
- μ : the rate at which the interface to the destination changes, or mobility rate, *i.e.*, $f_m(t) = \mu e^{-\mu t}$.

Assuming that both packet arrival and interface lifetime processes are independent, the mean number of data packets received by the destination per single interface change is given by: $\rho = \frac{\lambda}{\mu}$.

We define P to be the probability that the source has the correct (*i.e.*, consistent) location / interface information. For example, under MIP, P defines the probability that the home router contains consistent routing / location information. Under LISP, P defines the probability that the Mapping Server contains correct EID-RLOC mapping information. Under RINA, P defines the probability that the DIF contains correct routing information.

To allow for ease of analysis, we make the simplifying assumption that upon an interface change, the source does not send packets until the network state becomes consistent. We relax this assumption in the experimental section³. In steady state, P can be defined as the probability that the interface to the destination has not changed since the last packet arrival. Let t_p be the exponential random variable representing the packet inter-arrival time, and t_m be the exponential random variable representing the residual time during which the interface to the destination node does not change⁴. Thus, we have:

$$P = \text{Prob}(t_p < t_m) \quad (1)$$

³Note that our analytical model still captures the cost of being in an inconsistent state through the cost of updating routing / location information so as to make it consistent.

⁴Recall that the *residual* time of an exponentially distributed time is also exponential due to the memoryless property.

Parameters/Costs	Definitions
λ	packet sending rate of the source
μ	mobility rate of destination or rate of interface failure for multihomed destination
ρ	$\frac{\lambda}{\mu}$
C_L	Cost of lookup
C_U	Cost of location update
C_D	Cost of delivery
C_I	Cost of inconsistency

Table 1: Definitions of Parameters and Costs.

$$= \int_{t_p=0}^{\infty} f_p(t_p) \int_{t_m=t_p}^{\infty} f_m(t_m) dt_m dt_p \quad (2)$$

$$= \int_{t_p=0}^{\infty} \lambda e^{-\lambda t_p} \int_{t_m=t_p}^{\infty} \mu e^{-\mu t_m} dt_m dt_p \quad (3)$$

$$= \frac{\lambda}{\lambda + \mu} \quad (4)$$

The total cost per destination's interface change, C_T , is given by:

$$C_T = C_L + C_U + \rho(P \times C_D + (1 - P) \times C_I) \quad (5)$$

where the cost of inconsistency C_I is defined as $(C_D + C_D^{\text{OLD}})$, and C_D^{OLD} is the cost of packet delivery to the old location / interface. The term $(1 - P) \times C_I$ defines the cost of delivering packets that were in flight when the destination's interface changed. Henceforth, we take $C_D^{\text{OLD}} = C_D$, assuming that packets delivered to the wrong location need to be re-delivered to the correct location at the same cost. C_U represents the cost of messages propagated to restore the consistency of location/routing information after an interface change. For example, in the case of LISP, C_U represents the cost of updating the mapping server.

Table 1 summarizes our parameters.

4.2. MIP Cost Analysis

For MIP, we define the cost terms in Equation (5) as follows:

- $C_D = C_{\text{SN-HR}} + C_{\text{HR-DN}}$, where the cost of delivery of a single packet, C_D , is the sum of $C_{\text{SN-HR}}$, representing the cost of delivering a packet from the source node (SN) to the home router (HR), and $C_{\text{HR-DN}}$, representing the cost of delivering the packet from HR to the destination node (DN).
- $C_U = C_{\text{DN-FR}} + C_{\text{FR-HR}}$, where the cost of updating the destination's interface / location is the sum of $C_{\text{DN-FR}}$, which represents the cost of updating the foreign router, and $C_{\text{FR-HR}}$, which represents the cost of updating the home router.

Note that in MIP, $C_L = 0$, since the home router readily maintains the location of the destination node, and does not look up any mapping service.

4.3. LISP Cost Analysis

Under LISP, we define the cost terms in Equation (5) as follows:

- $C_D = C_L + C_{\text{SN-DN}}$, where the lookup cost, C_L , represents the cost of querying the EID-RLOC Mapping Server (MS) to identify the location of the destination Tunnel Router (TR). This lookup cost is incorporated in the delivery cost of every single data packet.
- $C_U = C_{\text{DN-TR}} + C_{\text{TR-MS}}$, where C_U , the cost of updating the MS, is the sum of $C_{\text{DN-TR}}$, which represents the cost of location update from the destination node to its TR, and $C_{\text{TR-MS}}$, which represents the cost of updating the MS.

4.4. RINA Cost Analysis

Support for mobility is inherent in the RINA architecture [5]. As described earlier, a data packet is delivered *hop-by-hop* to the destination across limited-scope Distributed Inter-process communication Facilities (DIFs). If the destination's interface changes, then the mapping from the destination node's address to the new interface is locally propagated. This local update involves unsubscription / withdrawal from/of the old interface (underlying DIF), and subscription / registration to/of the new interface (underlying DIF), which in turn results in updating the routing information to map to the new interface.

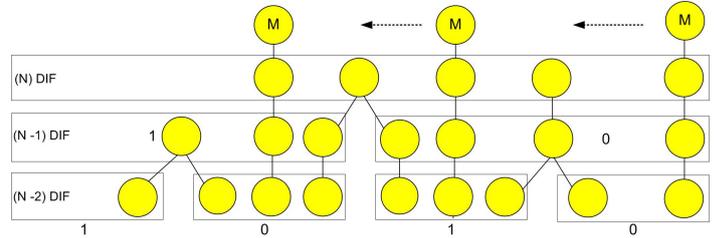


Figure 5: RINA DIF Structure.

As described in Section 3, in RINA, a node registers at a higher level DIF and gets an address assigned to it, which in turn serves as the node name for the lower level DIF. Thus, a communication request for that destination name can be readily resolved at the lower level DIF to a node address at that level. This process is repeated recursively over all RINA DIFs.

For analytical simplicity, as illustrated in Figure 5, we define the DIF structure of RINA as a binary tree where the scope (in terms of propagation delay) of each DIF is recursively partitioned into two lower-level scopes of half the size each.

To model the cost of propagation of routing change, we assign a topological based address prefix of zero or one to each part of the tree. Each node gets assigned an address whose prefix is derived from the DIF assigned values. For example, a node that resides in the rightmost lowest level DIF gets allocated an address whose prefix is 00.

When a destination node moves from one lowest level DIF to another, routing along the tree gets updated to point to its current location. The cost of update is determined by the scope of the DIF(s) that will need to be updated as a result of a user's mobility. We define l as the level (height) of routing propagations up the tree, which is obtained by taking the exclusive-or (XOR) of the destination's current address prefix and its previous address prefix, and computing l as the position of the most significant (leftmost) bit being set to one (assuming the position of the least significant bit is 1).

The total cost for routing updates is equal to:

$$\sum_{j=0}^{l-1} 2 \times \frac{D}{2^{h-j}}$$

where D is the diameter of the network, and h is the height of the tree.

Example: Referring to Figure 5, assume that a node with address prefix 00 moves to the nearby lowest level DIF to the right, then the node address prefix changes to 01. In this case, $00 \text{ XOR } 01 = 01$, so $l = 1$, and the total update cost is equal to $2 \frac{D}{2^2} = 2 \frac{D}{4}$ (given the height of the tree $h = 2$). This is the case since the parent node (with address prefix 0) needs to update its routing to point toward the new lowest level DIF instead of the old DIF. This requires the propagation of routing update across two lowest level DIFs, each of which spans a delay equal to fourth the diameter delay across the whole network.

Since our analysis deals with average costs, our goal is to compute the average value of l over possible mobility between different lowest level DIFs. To this end, we define an event β_i such that given m bit prefix addresses, bit i is flipped and bit $i + 1$ to m remain unchanged—in other words, β_i represents the probability of movement of a node that requires route updates to propagate up i levels, given a certain node mobility model. We also define the probability of bit i flipping as α_i . Thus, the probability of event $\beta_i = \alpha_i \prod_{j=i+1}^m (1 - \alpha_j)$. The expected value of the level of route update propagations l is given by $E[l] = \sum_{i=1}^m i \beta_i$.

Thus under RINA, we define the cost terms in Equation (5) as follows:

- $C_D = C_{\text{SN-DN}}$, since RINA strives to maintain a “direct” route to the destination.
- $C_U = \sum_{j=0}^{E[l]-1} 2 \times \frac{D}{2^{h-j}}$, which is the cost of routing updates upon mobility of the destination node.

As in MIP, $C_L = 0$ since each node (process) readily maintains the next-hop (routing) information to the destination node, and does not look up any mapping service.

4.5. LISP-MIP Cost Analysis

Farinacci *et al.* [11] propose the use of MIP as a means to managing fast mobility in LISP. This LISP-MIP variant can be generally used to deal with a change of destination's interface

whether because of mobility or re-routing to a multihomed destination.

Figure 6 highlights the cost of message delivery under the LISP-MIP architecture. The source is sending a packet to the destination node that has already moved to another domain and got a new care-of-address and updated its home agent, following the MIP protocol. Once the home agent intercepts the message, it tunnels it to the new location. An additional MS lookup is needed to obtain the address of the current destination Tunnel Router (TR).

Thus under LISP-MIP, assuming no caching of location information, we define the cost terms in Equation (5) as follows:

- $C_D = C_{\text{SN-L}} + C_{\text{SN-HR}} + C_{\text{HR-L}} + C_{\text{HR-DN}}$, where $C_{\text{SN-L}}$ and $C_{\text{HR-L}}$ represents the cost of querying the EID-RLOC mapping server at the source's TR, and at the destination's home TR, respectively.

The cost of update C_U in LISP-MIP is the same as that of MIP.

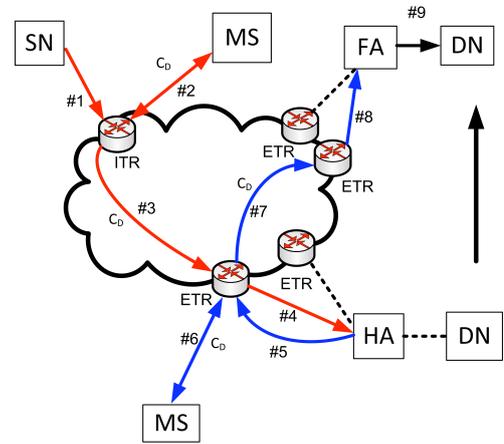


Figure 6: LISP-MIP cost of packet delivery.

4.6. LISP-Cache

Iannone *et al.* [13] studied the use of caching at the source Tunnel Router (TR) under LISP. Naturally, caching would decrease the per-packet cost of looking up the EID-RLOC mapping information, as long as the cached location information is accurate. The packet delivery process is still the same as that of Figure 2 with the only difference being that the lookup is only done once per cache entry lifetime (which, we assume, based on an infinite cache size, corresponds to the expected inter-failure time of the destination's interface). Thus we define the cost terms in Equation (5) as follows:

- $C_L > 0$, which represents the cost of querying an EID-RLOC mapping server to identify the location of the destination TR. This lookup is done *once* whenever the destination's interface changes and then cached for subsequent data packets.
- $C_D = C_{\text{SN-DN}}$, where we assume that looking up the cache for the location information is negligible, and thus does not contribute to the cost of delivery of every single data packet.

- $C_U = C_{DN-TR} + C_{TR-SN_{cache}}$,
where C_{DN-TR} represents the cost of location update from the DN to its TR, and $C_{TR-SN_{cache}}$ represents the cost of invalidating the source TR's cache due to the change in the destination's interface.

4.7. LISP-MIP-Cache

As a last LISP variant, we augment the LISP-MIP model described above with caching to reduce the cost of looking up location information. The delivery process still follows the same pattern shown in Figure 6, the only difference is that the lookup is only done once per cache entry lifetime (which, we assume, based on an infinite cache size, corresponds to the expected inter-failure time of the destination's interface). We define the cost terms in Equation (5) as follows:

- $C_L = (C_{SN-L} + C_{HR-L}) > 0$,
which represents the costs of querying a mapping server at the source's TR and the destination's home TR, respectively. We note that these lookup costs are only incurred once whenever the destination's interface changes. The location information is then cached for future use. Thus these lookup costs do not contribute to the delivery cost of every single data packet.
- $C_D = C_{SN-HR} + C_{HR-DN}$,
which defines the cost of delivery of a single data packet. The cost of looking up the cached location information is assumed to be negligible.
- $C_U = C_{DN-FR} + C_{FR-HR}$,
which defines the cost of updating the destination's location at its home router.

A summary of the costs under all schemes is shown in Table 2.

5. Numerical Results

We present numerical results using the cost equations defined in Section 4, assuming grid topologies. As mentioned earlier, we define costs in terms of average path lengths between communicating entities, *e.g.*, between a source's TR and a mapping server in LISP.

For an $N \times N$ grid topology, the average distance between any two nodes is given by $1.333(N/2)$ hops. We use this average distance as the cost of communication between two nodes that are *not* on the same network. On the other hand, if the communicating nodes are on the same network, the cost is relatively smaller (and independent of the size of the topology) — we take the cost to be two hops between a node and its TR, and one hop otherwise. For RINA we model a binary DIF tree on top of the grid topology such that each leaf (lowest level) DIF contains two network nodes.

Figure 7 presents results for an 8×8 grid for the various schemes as ρ takes on different values. The height of our RINA binary tree is 5.

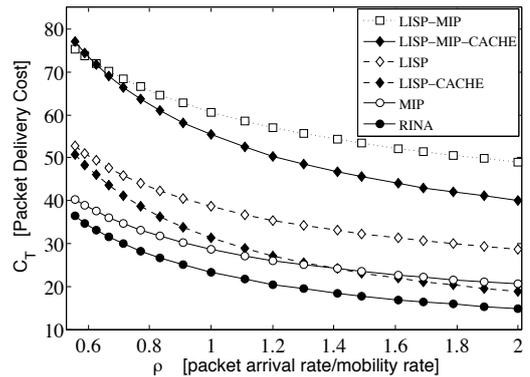


Figure 7: Numerical results for an 8×8 grid.

We assume a skewed probability distribution for the movement of nodes between (lowest level) DIFs such that the probability of moving from the leftmost DIF to the rightmost DIF is minimum — the probability of address bit i being flipped is $1/2^i$ (cf. Section 4.4). This is a reasonable assumption since non-local movements would not be practical in reality. Given the above mobility distribution, $E[l] \approx 3$.

As ρ increases, the total cost for all schemes decreases (as expected). RINA has the lowest total cost, while LISP has the worst cost. It is worthwhile to mention that even when the total cost of location update in RINA is higher than that of MIP, due to RINA's use of the "direct" path to the destination, RINA's total cost of packet delivery is lower.

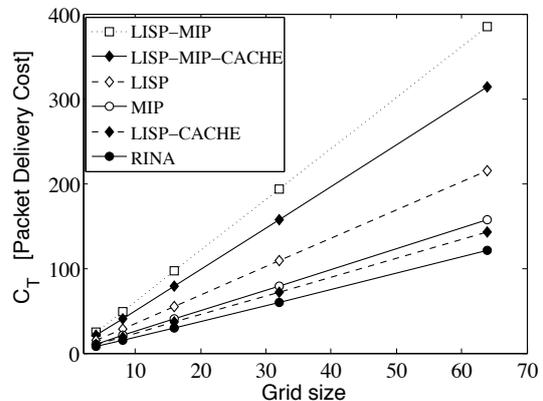


Figure 8: Numerical results for varying grid sizes.

Figure 8 shows the total costs of the various schemes for varying grid sizes N for $\rho = 2$. As N increases, the total cost for all schemes increases, with RINA incurring the lowest cost at a sublinear increase rate.

6. Mobility Simulation Results

We validate our cost model using simulation. In our simulations, "cost" is represented by average packet delay, packet drop ratio, or inverse of packet delivery ratio. To obtain an

Costs	Mobile IP	RINA	LISP	LISP-Cache	LISP-MIP	LISP-MIP-Cache
C_D	$C_{SN-HR} + C_{HR-DN}$	C_{SN-DN}	$C_L + C_{SN-DN}$	C_{SN-DN}	$C_{SN-L} + C_{SN-HR} + C_{HR-L} + C_{HR-DN}$	$C_{SN-HR} + C_{HR-DN}$
C_U	$C_{DN-FR} + C_{FR-HR}$	$C_{DIF-UNREG} + C_{DIF-REG}$	$C_{DN-TR} + C_{TR-MS}$	$C_{DN-TR} + C_{TR-SN_{CACHE}}$	$C_{DN-FR} + C_{FR-HR}$	$C_{DN-FR} + C_{FR-HR}$
C_I	$C_D + C_D^{OLD}$	$C_D + C_D^{OLD}$	$C_D + C_D^{OLD}$	$C_D + C_D^{OLD}$	$C_D + C_D^{OLD}$	$C_D + C_D^{OLD}$
C_L	0	0	$C_{TR-MS} + C_{MS-TR}$	$C_{TR-MS} + C_{MS-TR}$	$2(C_{TR-MS} + C_{MS-TR})$	$2(C_{TR-MS} + C_{MS-TR})$

Table 2: Components of total cost in response to a single interface change.

internet-like topology, we use the BRITE topology generator [20] to generate a network of autonomous systems (ASes) and their router topologies. We use the top-down generation model of BRITE which is based on two phases. In the first phase, an AS topology is initially generated using the Barabasi-Albert model with incremental growth type and preferential connectivity. In the second phase, a router-level topology is generated for each AS, where router nodes are placed randomly on the 2D-plane and connected using the Waxman model. The average path length between nodes in the generated topologies is 14 hops, consistent with Internet measurement studies [21].

We simulate a single source-destination pair where the source sends packets at a rate λ while μ defines the rate at which the destination interface changes as a result of node mobility. We adopt a random walk mobility model where the destination node moves within a specified hop radius from its current location. For MIP, we assume that the cost of update is the round-trip propagation delay between the mobile destination node’s current location and its designated home router. For LISP, we assume that updating the EID-RLOC mapping server takes an exponentially distributed time with a mean value that corresponds to the average path length, upper bounded by the network diameter. We also assume an infinite cache size for LISP-Cache and LISP-MIP-Cache.

For RINA, we assume a two-level hierarchy where at the AS level, border routers form the higher level DIF, whereas internal routers of each AS constitute the lower layer DIFs. We simulate *hop-by-hop* routing in RINA, and at the higher level DIF, whenever the destination’s interface changes due to mobility, we calculate the cost of updating the “intermediary” leading to the destination to be the round-trip propagation delay between them. If there is no path to the destination from the “intermediary”, we assume the source needs to be updated to route to a new “intermediary” leading to the destination. The cost of updating the source is calculated as the round-trip propagation delay between the source and the destination.

Figures 9 and 10 show the packet delivery ratio and the average packet delivery time under the various approaches. All results are presented with 90 percent confidence intervals. The results are consistent with our analytical results. RINA yields the lowest cost in terms of packet drop ratio, delivering packets

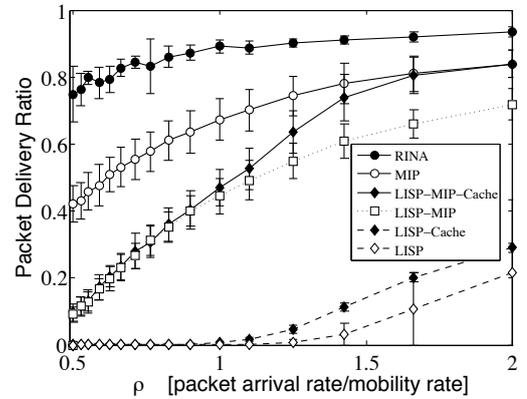


Figure 9: Packet Delivery Ratio.

at the lowest possible delay due to its local routing adaptation within the scope of the lower level DIFs connecting the intermediary and destination. LISP-MIP has higher packet delivery ratio compared to LISP, but higher average packet delivery delay.

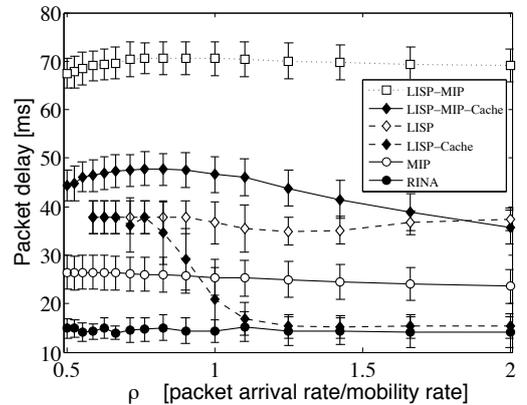


Figure 10: Average Packet Delivery Time.

As the mobility rate decreases, approaches that utilize caching like caching over LISP and caching over LISP-MIP

DataSet	Chicago	San Jose
Unique ASes	66	97
Packets	74123	74123
Nodes (routers)	2178	2425
Edges	4488	5041

Table 3: Topology Properties.

gain a significant advantage over non cached approaches. In Figure 9, and contrary to our analytical results, MIP and LISP-MIP perform better than LISP in terms of packet delivery ratio. This is due to the fact that the communication between the source node and the home router does not suffer any losses, which leads to better packet delivery ratio. However, MIP and LISP-MIP do incur a higher packet delay, which is consistent with our cost model.

7. Multihoming Trace-driven Simulation

In this section we validate our analytical results using trace-driven simulation based on CAIDA’s anonymized packet traces [6]. This simulation considers only multihoming, so we do not include experimental results for Mobile-IP. We select two datasets from two Equinix locations: Chicago and San Jose (dated 20090219-045912 and 20090219-060100, respectively). The traces consist of anonymized tcpdump packets from different source-destination pairs. Each trace file contains more than a million records. The traces provide only anonymized source-destination pairs and packet arrival times. Due to the unavailability of complete real Internet topologies, and the difficulty of mapping the packets to any real topology since they are anonymized, we use the BRITE topology generator to generate an underlying AS and router network topology in the same way described in Section 6.

To keep the simulation and generated topologies manageable, we only consider the first 74123 packets from each packet trace. We assume that all IP addresses which have a common 16-bit prefix belong to the same AS. Table 3 highlights properties of our two simulated topologies.

We utilize the packet timestamp as the packet arrival time, thus we relax the Poisson arrival assumption made in the analytical model. Furthermore, the time between link failures follows an exponential distribution. To allow for exact measurement of the effect of interface failures, we simulate the failure of only a single link (interface) at a time. We also make sure that interface failures occur only on destinations that are multihomed. We also assume an infinite cache size for LISP-Cache. Similar to our mobility simulations (cf. Section 6), we assume that for LISP, updating the EID-RLOC mapping server takes an exponentially distributed time with a mean value that corresponds to the average path length, upper bounded by the network diameter. Simulation of routing in RINA models a similar approach to the one described in Section 6.

7.1. Results

In this section, we present the results of our simulations. We measure packet drop ratio, and packet delivery delay. We experimented using topologies generated using BRITE’s top-down approach, where in the initial phase, the AS topology is generated using either the Barabasi-Albert (BA) model or Waxman’s model [22].

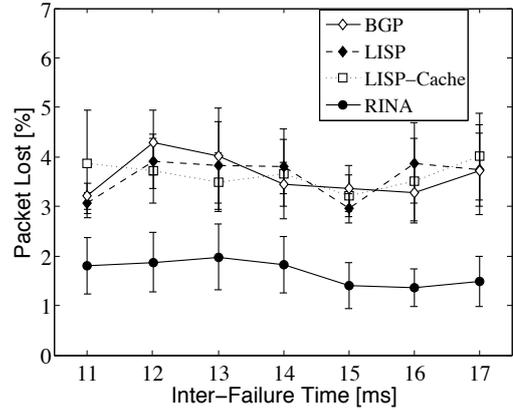


Figure 11: Packet Drop Ratio (Chicago dataset, Waxman AS-topology).

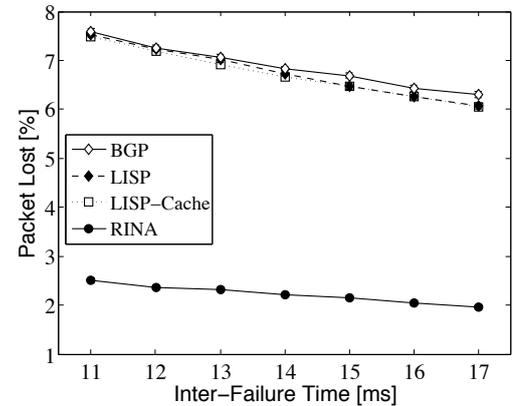


Figure 12: Packet Drop Ratio (San Jose dataset, Waxman AS-topology).

Figures 11 and 12 show the simulation results of packet drop ratio using Waxman’s AS topology based on the two datasets (Chicago and San Jose), the Figures 13 and 14 show results of packet drop ratio for BA AS topology. The results confirm our analytical model. As the link inter-failure time increases the percentage of packets dropped decreases. Based on the Waxman’s AS topology (Figures 11 and 12), RINA drops around 2% and 2.5% of the packets, respectively, while BGP, LISP, and LISP with caching, drop around 4% and 8% of the packets, respectively. BA AS topology results, are consistent with the Waxman’s AS topology (Figures 13 and 14) with RINA dropping around 4.5% and 5.5% of the packets, respectively, while BGP, LISP, and LISP with caching, drop around 6% and 10% of the packets, respectively.

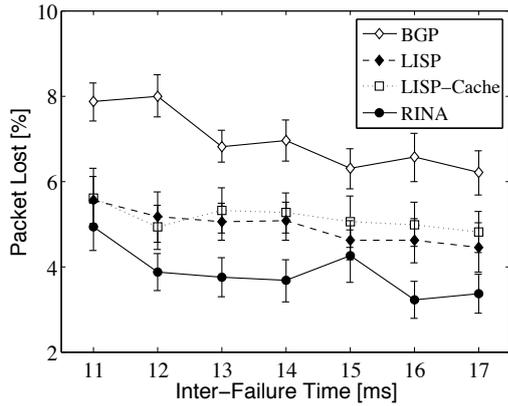


Figure 13: Packet Drop Ratio (Chicago dataset, BA AS-topology).

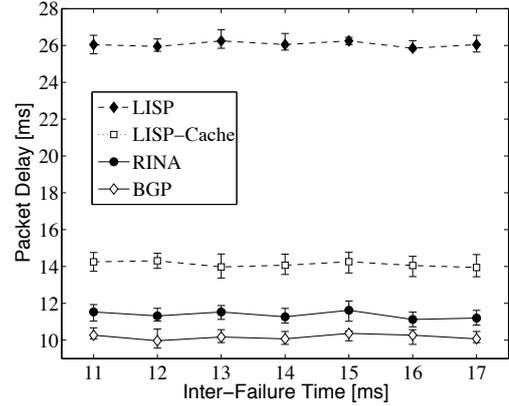


Figure 15: Average Packet Delivery Time (Chicago dataset, Waxman AS-topology).

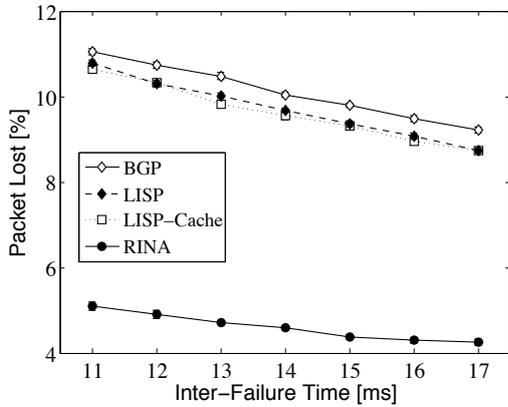


Figure 14: Packet Drop Ratio (San Jose dataset, BA AS-topology).

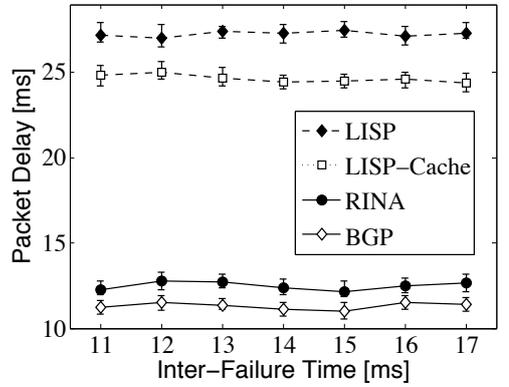


Figure 16: Average Packet Delivery Time (San Jose dataset, Waxman AS-topology).

Figures 15 and 16 show the average packet delivery time for packets that reach the destination using Waxman’s AS topology, while Figures 17 and 18 show the average packet delivery time results for BA AS topology. The delivery time of RINA and BGP is smaller due to the fact that there is no need to contact a mapping server. The benefit of caching for LISP is highlighted by a smaller average packet delivery time.

RINA yields the lowest cost in terms of packet drop ratio, delivering packets at the lowest possible delay due to its local routing adaptation within the scope of the overlay involving the source, destination, and “intermediary” node.

Note that BGP’s delay is slightly lower than that of RINA, since BGP’s lack of fine-grained routing control makes it incapable of adapting to a failure of the shortest path to a specific destination node, however, for those packets that get delivered when the shortest path is up, their delivery delay is smallest. On the other hand, RINA enables the construction of an “overlay” network between the source node, destination node, and an intermediate node(s) (intermediary) that are capable of re-routing around failed paths (interfaces). Thus, under RINA, more packets are successfully delivered, but those packets taking alternate paths when the primary paths are down, experience slightly higher delay.

We also observe that under LISP, the delay is almost double that of RINA and BGP, since LISP requires a mapping lookup which adds extra delay that is in the order of the average path length of around 14 hops (msec) in our topologies.

8. Related Work

8.1. Architectural Changes

File transfer and email were the main applications when, in the 70s the Internet protocols were designed. The number of connected hosts have grown from less than 200 in 1980, to 570 millions in 2008 [23]. Experience and technological progress that would make a redesign of the Internet nowadays substantially different, together with the deficiencies of the current architecture are motivating research efforts in how the Internet architecture should be. Such research efforts could be classified along two main dimensions:

Approach: Purist versus pluralist. The former supports flexibility to meet current and future application requirements, and the latter envisions parallel protocol stacks able to cope with multiple service requirements. The approach envisioned in

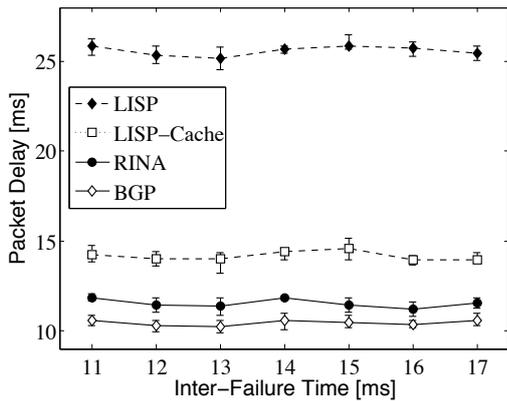


Figure 17: Average Packet Delivery Time (Chicago dataset, BA AS-topology).

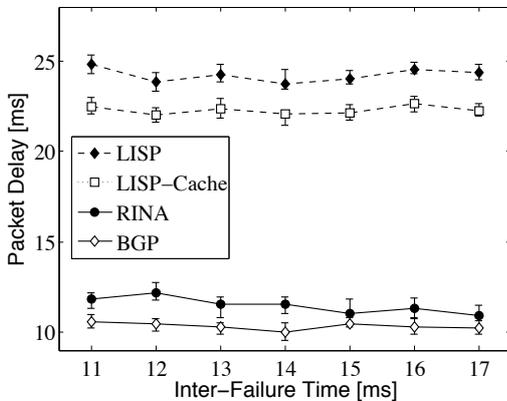


Figure 18: Average Packet Delivery Time (San Jose dataset, BA AS-topology).

RINA is classifiable as both pluralist and purist. This is because we can flexibly compose a diverse set of end-to-end services across several underlying DIFs, where each DIF may run a different set of policies to meet certain local service requirements.

Design: Evolutionary versus clean-slate [24]. For many years, extensive research have been conducted to overcome the Internet impasse, with improvements and patches that would coexist within the design constraints of the current architectures, see *e.g.*, all the efforts on overlays mostly inspired by [25] and [26]. On the other hand, clean-slate approaches ignore such constraints to exploit potential benefits [27, 28, 29].

Unlike evolutionary approaches, our RINA architecture is a clean-slate design based on the inter-process communication (IPC) principle. Quoting Robert Metcalfe: “*Networking is inter-process communication and only inter-process communication.*” (1972) In this view, all network services, *e.g.* transport and internetworking tasks, together constitute a Distributed IPC Facility (DIF) to application processes. RINA applies this concept recursively, whereby the processes that make up a DIF can themselves act as application processes to request services from lower level DIFs.

Recursion has been recently promoted in network architec-

tures, but to the best of our knowledge, this has been limited to tentative proposals of repeated functions of existing layers, and how one may either reduce duplication or create a meta-function (*e.g.*, error and flow control) that could be re-used in many layers, *e.g.*, Touch *et al* [30]. Independently, we have pursued a general theory to identify patterns in network architecture [31, 5], which the RINA architecture embodies.

8.2. Multihoming and Mobility

We are certainly not the first to advocate the need for new support for mobility [32]. In the current Internet, a system is identified by its IP / Internet address. As a result, when a system changes its point-of-attachment, its IP address changes. This makes reaching mobile systems difficult.

Multiple efforts have attempted to address this naming / addressing problem, proposing and deploying new mobility and multihoming protocols and architectures, including Mobile IP [3], and its hierarchical extension HMIPv6 [8], Mobcast [33], a system based on a *Proxy IP Anycast Service* (PIAS), Internet Indirection Infrastructure [34], Host Identity Protocol [35], and others [36, 28].

Those attempts may be classified as network or routing-oriented (*e.g.* [11, 36, 28]), host-centric ([35, 3, 37]), and hybrid edge-based solutions ([38]).

An example of network-oriented solution is the so-called LISP: location / identifier split [11]. LISP uses the locator not to locate the destination node (*i.e.*, where it is), rather a path that leads to it (*i.e.*, how to get there).

Host-centric solutions instead do routing at the end points of the network (hosts). IP addresses are given both ID and locator semantics. Route selection is done at the host where the ID is mapped to a particular locator, rather than letting routers select the best path. Shim6 [37] is an example of a host-centric solution that addresses multihoming.

The HAIR [38] architecture proposes to separate locators from identifiers in a hierarchical fashion. In fact, it can be seen as a hierarchical version of LISP. If a host moves across adjacent domains at the same hierarchical level, then routing updates do not necessarily have to propagate to the core of the Internet. Although HAIR addresses scalability by restricting the visibility of routing updates as in RINA, it suffers from the same drawbacks of network-based solutions, that is, routing based on interface (IP) addresses rather than node addresses. All these solutions in fact, bind IP/anycast host names to addresses, making it hard to utilize alternate paths in case the corresponding interface goes down.

By adopting and extending Saltzer’s naming and addressing schema [39] in a recursive fashion, RINA names applications/nodes rather than interfaces (point-of-attachment). This late binding of a node’s address to its interface allows RINA to effectively deal with interface changes due to multihoming or mobility.

9. Conclusion

We highlighted the benefits of ROuting in Recursive Architectures (RORA) as a model for the future internet. We de-

veloped a cost model to evaluate the mobility / multihoming support of RINA, LISP, and MIP. RINA incurs the lowest cost, while LISP incurs the highest cost. We also validated our model for both mobility and multihoming using simulation on an Internet-like topology and based on real packet traces from CAIDA. We are currently investigating dynamic DIF formation that optimizes routing in the RINA architecture in the presence of arbitrary node/link failures and mobility. We are also prototyping RINA's recursive routing.

Acknowledgements

The authors would like to thank members of the RINA team, in particular John Day, and Karim Mattar for their support and valuable feedback. This work has been partially supported by National Science Foundation awards: CISE / CNS #0963974, CISE / CCF #0820138, and CISE / CSR #0720604.

References

- [1] L. Iannone, D. Saucez, O. Bonaventure, OpenLISP: An Open Source Implementation of the Locator/ID Separation Protocol, in: IEEE INFOCOM, Demo paper, 2009.
- [2] D. Meyer, The Locator/Identifier Separation Protocol (LISP) (2008).
- [3] C. Perkins, IP Mobility Support for IPv4, RFC 3344, Internet Engineering Task Force (Aug. 2002).
URL <http://www.ietf.org/rfc/rfc3344.txt>
- [4] B. Quoitin, L. Iannone, C. de Launois, O. Bonaventure, Evaluating the Benefits of the Locator/Identifier Separation, in: MobiArch '07: Proceedings of 2nd ACM/IEEE International Workshop on Mobility in the Evolving Internet Architecture, ACM, New York, NY, USA, 2007, pp. 1–6. doi:<http://doi.acm.org/10.1145/1366919.1366926>.
- [5] J. Day, I. Matta, K. Mattar, "Networking is IPC": A Guiding Principle to a Better Internet, in: Proceedings of ReArch, Madrid, SPAIN, 2008.
- [6] C. Walsworth, E. Aben, k. Claffy, D. Andersen, The CAIDA Anonymized 2009 Internet Traces, www.caida.org/data/passive/passive_2009_dataset.xml (2009).
- [7] J. Lei, X. Fu, Evaluating the Benefits of Introducing PMIPv6 for Localized Mobility Management, in: Wireless Communications and Mobile Computing Conference, 2008. IWCMC '08. International, 2008, pp. 74–80. doi:[10.1109/IWCMC.2008.14](http://doi.acm.org/10.1109/IWCMC.2008.14).
- [8] H. Soliman, C. Castelluccia, K. E. Malki, L. Bellier, Hierarchical MIPv6, Internet RFC 4140 (August 2005).
- [9] E. R. Koodli, Fast Handovers for Mobile IPv6, Internet RFC 4068 (July 2005).
- [10] X. P. Costa, M. Torrent-Moreno, H. Hartenstein, A Performance Comparison of Mobile IPv6, Hierarchical Mobile IPv6, fast handovers for Mobile IPv6 and their Combination, Mobile Computing and Communications Review 7 (4) (2003) 5–19.
- [11] D. Farinacci, V. Fuller, D. Meyer, D. Lewis, Locator/ID Separation Protocol (LISP), <http://tools.ietf.org/html/draft-farinacci-lisp-12> (March 2009).
- [12] D. Farinacci, V. Fuller, LISP Map Server, online, <http://tools.ietf.org/html/draft-fuller-lisp-ms-00> (March 2009).
- [13] L. Iannone, O. Bonaventure, On the Cost of Caching Locator/ID Mappings, in: Proceedings of the ACM CoNEXT '07 conference, ACM, New York, NY, USA, 2007, pp. 1–12. doi:<http://doi.acm.org/10.1145/1364654.1364663>.
- [14] A Border Gateway Protocol 4 (BGP-4), author=Rekhter, Y. and Li, T. and Hares, S. and others, year=1995, publisher=RFC 1771, March 1995.
- [15] G. Boddapati, J. Day, I. Matta, L. Chitkusev, Assessing the Security of a Clean-Slate Internet Architecture, Tech. Rep. BUCS-TR-2008-021, CS Dept, Boston U. (June 22 2009).
- [16] G. Gursun, I. Matta, K. Mattar, Revisiting a Soft-State Approach to Managing Reliable Transport Connections, in: Proceedings of the Eighth International Workshop on Protocols for Future, Large-Scale and Diverse Network Transports (PFLDNeT), 2010.
- [17] K. Ratnam, I. Matta, S. Rangarajan, Analysis of Caching-Based Location Management in Personal Communication Networks, in: ICNP '99: Proceedings of the Seventh Annual International Conference on Network Protocols, IEEE Computer Society, Washington, DC, USA, 1999, p. 293.
- [18] K. Ratnam, S. Rangarajan, A. Dabhura, An Efficient Fault-Tolerant Location Management Protocol for Personal Communication Networks, Vehicular Technology, IEEE Transactions on 49 (6) (2000) 2359–2369.
- [19] K. Kyamakya, K. Jobmann, Location Management in Cellular Networks: Classification of the Most Important Paradigms, Realistic Simulation Framework, and Relative Performance Analysis, Vehicular Technology, IEEE Transactions on 54 (2) (2005) 687–708.
- [20] A. Medina, A. Lakhina, I. Matta, J. Byers, BRITE: An Approach to Universal Topology Generation, in: Proceedings of MASCOTS, 2001.
- [21] V. Paxson, End-to-end Routing Behavior in the Internet, SIGCOMM Comput. Commun. Rev. 36 (5) (2006) 41–56. doi:<http://doi.acm.org/10.1145/1163593.1163602>.
- [22] B. Waxman, Routing of Multipoint Connections, Selected Areas in Communications, IEEE Journal on 6 (9) (1988) 1617–1622. doi:[10.1109/49.12889](http://doi.acm.org/10.1109/49.12889).
- [23] J. Roberts, The Clean-slate Approach to Future Internet Design: a Survey of Research Initiatives., Annals of Telecommunications 64 (2009) 271–276.
- [24] J. Rexford, C. Dovrolis, Future Internet Architecture: Clean-slate versus Evolutionary Research (2010). doi:<http://doi.acm.org/10.1145/1810891.1810906>.
- [25] D. G. Andersen, H. Balakrishnan, F. M. Kaashoek, R. Morris, Resilient Overlay Networks, in: Symposium on Operating Systems Principles, 2001, pp. 131–145.
- [26] J. Touch, S. Hotz, The X-Bone, in: Proceedings of Global Internet Mini-Conference Globecom, 1998.
- [27] D. D. Clark, J. Wroclawski, K. R. Sollins, R. Braden, Tussle in Cyberspace: Defining Tomorrow's Internet, in: Proceeding of ACM SIGCOMM, 2002, pp. 347–356.
- [28] X. Xu, D. Guo, R. Jain, J. Pan, S. Paul, Internet 3.0: Next Generation Internet (2008).
- [29] C. Esteve, F. L. Verdi, M. F. Magalhães, Towards a new Generation of Information-oriented Internetworking Architectures, in: Proceedings ACM CoNEXT Conference, ACM, New York, NY, USA, 2008, pp. 1–6. doi:<http://doi.acm.org/10.1145/1544012.1544077>.
- [30] J. Touch, Y.-S. Wang, V. Pingali, A Recursive Network Architecture, Tech. rep., USC/ISI (October 2006).
- [31] J. Day, Patterns in Network Architecture: A Return to Fundamentals, Prentice Hall, 2008.
- [32] D. Le, X. Fu, D. Hogrefe, A Review of Mobility Support Paradigms for the Internet, IEEE Communications Surveys and Tutorials 8 (2006) 2–15.
- [33] C. Lee, K. Attrey, C. Caballero, N. Feamster, M. Mihail, J. Copeland, MobCast: Overlay Architecture for Seamless IP Mobility using Scalable Anycast Proxies, in: IEEE Wireless Communications and Networking Conference (WCNC), 2007, pp. 3872–3876.
- [34] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, S. Surana, Internet Indirection Infrastructure, in: In Proceedings of ACM SIGCOMM, 2002, pp. 73–86.
- [35] R. Moskowitz, P. Nikander, Host Identity Protocol Architecture, in: Internet Draft, draft-ietf-hip-arch-03, 2005, p. 24.
- [36] H. Balakrishnan, K. Lakshminarayanan, S. Ratnasamy, S. Shenker, I. Stoica, M. Walfish, A Layered Naming Architecture for the Internet, in: SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications, ACM, New York, NY, USA, 2004, pp. 343–352. doi:<http://doi.acm.org/10.1145/1015467.1015505>.
- [37] E. Nordmark, M. Bagnulo, Shim6: Level 3 Multihoming Shim Protocol for IPv6, RFC 5533, Internet Engineering Task Force (Jun. 2009).
- [38] A. Feldmann, L. Cittadini, W. Mühlbauer, R. Bush, O. Maennel, HAIR: Hierarchical Architecture for Internet Routing, in: ReArch '09: Proceedings of the 2009 workshop on Re-architecting the internet, ACM, New York, NY, USA, 2009, pp. 43–48.
- [39] J. Saltzer, Naming and Binding of Objects, in: R. Bayer (Ed.), Operating Systems, Lecture notes in Computer Science, Vol. 60, Springer-Verlag, New York, 1978.