

# On the Cost of Supporting Mobility and Multihoming

VATCHE ISHAKIAN   IBRAHIM MATTA   JOSEPH AKINWUMI

COMPUTER SCIENCE DEPARTMENT, BOSTON UNIVERSITY

{visahak, matta, akin}@cs.bu.edu

**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. Multihoming can be viewed as a special case of mobility—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 mobility (and consequently multihoming). The Recursive InterNetwork Architecture (RINA) [1] was recently proposed as a clean-slate solution to the current problems of the Internet. In this paper, we perform an average-case cost analysis to compare the mobility / multihoming support of RINA, against that of other approaches such as LISP and Mobile-IP. We also validate our analysis using simulation.

## I. 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 (mobile) host specifies a particular interface (connection), rather than the node itself. Because routing is done based on this interface (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 [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 update the location of the mobile node at its home agent. 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 [1] is a recently proposed Recursive InterNetwork Architecture. It uses the concept of Distributed IPC 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 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.

### *Our Contributions:*

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 (mobile or multihomed) 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 mobile or multihomed destination node. We also validate our analytical model using simulation.

## II. BACKGROUND

### A. Mobile-IP

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.

The basic delivery process of data packets from a source node (SN) to a destination node (DN) is as follows: (1) The SN sends the packet to the DN via standard routing. (2) The Home Agent (HA) intercepts the datagram and tunnels it to the

destination’s current location (care-of-address) which the DN acquires from the Foreign Agent (FA). (3) The FA intercepts the datagram and delivers it to the DN.

### B. LISP

The Locator/ID Separation Protocol (LISP), proposed by Farinacci *et al.* [2], 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 Map 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 1(a)): (1) The SN forwards the packet to its border router called Ingress Tunnel Router (ITR). (2) The source ITR performs a lookup query for the destination-to-RLOC mapping [5]. (3) The ITR tunnels the data packets to the destination’s RLOC referred to as Egress Tunnel Router (ETR). (4) Upon intercepting the packet, the destination’s ETR forwards the packet to the DN.

Upon failure of an active interface, a multihomed destination node would send an update to its ETR, which in turn would update the EID-to-RLOC MS. The sequence of messages is shown in Figure 1(b).




Fig. 1. (a) LISP Architecture, and (b) LISP cost of update.

Different variants of LISP only differ in how the EID-to-RLOC mapping is done [5]. The use of caching for lookup has also been recently explored in [6].

### C. RINA

In RINA, application processes or services have globally unique names, and networking is viewed as distributed Inter-Process Communication (IPC) [1].

If an application process in RINA needs to communicate with another application process, it requests service from the underlying Distributed IPC Facility (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.

The route to the destination node address (to which the destination application process is connected) is computed as a sequence of intermediate node addresses. 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 (PoA)

of the higher-level node. Thus, RINA’s addresses are relative. 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.




Fig. 2. RINA Architecture

**RINA Example:** Without loss of generality, Figure 2 shows a source process sending packets to a destination process using the services of the underlying DIFs.<sup>1</sup> The source and destination form a (high-level) DIF with an intermediate process, which we call “intermediary”, such that the intermediary is connected to the destination using two separate interfaces over two different underlying DIFs. This 3-node DIF can be thought of as an “overlay” to which source, destination, and intermediary had subscribed. When a packet reaches the intermediary, it forwards it based on the current best / operational interface leading to the destination.

It is important to highlight the difference between how BGP and RINA handle route / interface failures. In BGP, even if there is a specific path failure to a specific prefix (node), BGP 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. In Figure 2, if the (solid) overlay link I–D that uses the underlying DIF *B* goes down, node I would locally adapt and start using the (dotted) overlay link I–D that uses the underlying DIF *C*. Thus, RINA provides finer grained control over routing to multihomed destinations.

### III. COST MODEL

In this section we study the average (communication) cost of supporting mobility under MIP, LISP and RINA architectures.

#### A. Assumptions, Cost Definitions, and Parameters

We assume a single source-destination model where the source sends data packets at a constant rate. We analyze the

<sup>1</sup>Note that in RINA, a single system may have multiple processes which are members of different DIFs at different levels [1].

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 ( $C_U$ ), and *location inconsistency* cost ( $C_I$ ). *Location lookup* cost is defined only for LISP, to capture the cost of querying a mapping server (MS) 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.

In our model, we assume that the inter-arrival time 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:

- $\lambda$ : the mean packet arrival rate, *i.e.*,  $f_p(t) = \lambda e^{-\lambda t}$ .
- $\mu$ : 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 agent (router) contains consistent routing / location information. Under LISP,  $P$  defines the probability that the Map Server contains correct routing information. Under RINA,  $P$  defines the probability that the DIF contains correct routing information.

In steady state,  $P$  can be defined as the probability that the interface to the destination has not changed since the last packet delivery. 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 change<sup>2</sup>. Thus, we have:

$$\begin{aligned} P &= \text{Prob}(t_p < t_m) \\ &= \int_{t_p=0}^{\infty} f_p(t_p) \int_{t_m=t_p}^{\infty} f_m(t_m) dt_m dt_p \\ &= \int_{t_p=0}^{\infty} \lambda e^{-\lambda t_p} \int_{t_m=t_p}^{\infty} \mu e^{-\mu t_m} dt_m dt_p \\ &= \frac{\lambda}{\lambda + \mu} \end{aligned}$$

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

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

where  $C_I$  is defined as  $(C_D + C_D^{\text{OLD}})$ , and  $C_D^{\text{OLD}}$  is the cost of packet delivery to the old location / interface. 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.

### B. MIP Cost Analysis

For MIP, we define the cost terms in Equation (1) 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

<sup>2</sup>Recall that the *residual* time of an exponentially distributed time is also exponential due to the memoryless property.

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 agent (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.

### C. LISP Cost Analysis

Under LISP, we define the cost terms in Equation (1) as follows:

- $C_D = C_L + C_{\text{SN-DN}}$ , where the lookup cost,  $C_L$ , represents the cost of querying the Map 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.

### D. RINA Cost Analysis

Support for mobility is inherent in the RINA architecture [1]. As reviewed 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.




Fig. 3. RINA DIF Structure

Registration in RINA is done in a bottom-up fashion [1] where 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 ease of analysis we define the DIF structure of RINA as a binary tree, where a tree node represents a network node that is subscribed to a DIF of size indicated in Figure 3, as well as to all lower level DIFs in its subtree of half the size each. Thus, to route over the scope of the whole network, say of size  $n$  nodes, routing can start from the root of the tree and proceed recursively down toward the lowest level DIF to which the destination is connected.

To assign addresses to nodes, we assign to each tree edge a binary value of zero or one. Each node gets assigned an address whose prefix is derived from the tree edge values. For example, a node that resides in the leftmost 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 total number of nodes 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 given 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 3, 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. Note that we further multiply the update cost by two to account for acknowledgements.

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  $\beta_i$  such that given  $m$ -bit addresses, bit  $i$  is flipped and bit  $i + 1$  to  $m$  remain unchanged. We also define the probability of bit  $i$  flipping as  $\alpha_i$ . 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 (1) as follows:

- $C_D = C_{SN-DN}$ , 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.

### E. LISP-MIP Cost Analysis

Farinacci *et al.* [2] 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 4 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 lookup is needed to obtain the address of the current destination tunnel router.

Thus under LISP-MIP, assuming no caching of location information, we define the cost terms in Equation (1) as follows:

- $C_D = C_{SN-L} + C_{SN-HR} + C_{HR-L} + C_{HR-DN}$ , where  $C_{SN-L}$  and  $C_{HR-L}$  represent the cost of querying the map server at the source's TR, and at the destination's home TR, respectively.

The cost of update  $C_U$  in LISP-MIP is similar to the update cost of MIP.




Fig. 4. LISP-MIP cost of packet delivery

### F. Numerical Results

We present numerical results using the cost equations defined above for 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.




Fig. 5. Numerical results for an  $8 \times 8$  grid

Figure 5 presents results for an  $8 \times 8$  grid for the various schemes as  $\rho$  takes on different values. The height of our RINA binary tree is 5. For calculating the expected value  $E[l]$  we assume a skewed probability distribution for movement of nodes between (lowest level) DIFs such that the probability of moving from the leftmost DIF to the rightmost DIF is minimum—this is a reasonable assumption since such movements would not be practical in reality. We take  $E[l] = 3$ . As  $\rho$  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 the total cost of location update in RINA is higher than that of MIP, but due to the “direct” path to the destination, RINA’s total cost of packet delivery is lower.

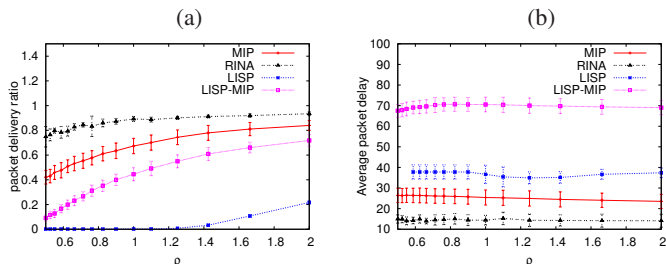


Fig. 6. (a) Packet Delivery Ratio, (b) Average Packet Delivery Time

#### IV. SIMULATION RESULTS

We validate our cost model using simulation. In simulation, “cost” is represented by average packet delay or inverse of packet delivery ratio. To generate an internet-like topology, we use the BRITE topology generator [7] 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 [8].

We simulate a single source-destination pair where the source sends packets at a rate  $\lambda$  while  $\mu$  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 map server takes an exponentially distributed time with a mean value that corresponds to the average path length, upper bounded by the network diameter. For simplicity, we assume the delay of one hop is 1 ms.

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.

Figure 6(a) and 6(b) show the packet delivery ratio and the average packet delivery time under the various approaches. The results are consistent with our analytical results. 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 lower level DIFs connecting the “intermediary” and destination. LISP-MIP has higher packet delivery ratio compared to LISP, but higher average packet delivery delay.

#### V. CONCLUSION

We developed a cost model to evaluate the mobility support of RINA, LISP, and MIP. RINA incurs the lowest cost, while LISP incurs the highest cost. We also validated our model using simulation on an Internet-like topology. 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 will also prototype RINA’s recursive routing protocol.

#### VI. ACKNOWLEDGEMENTS

This work has been partially supported by National Science Foundation awards: CISE / CNS #0963974, CISE / CCF #0820138, and CISE / CSR #0720604. Thanks to other members of the RINA team for their support and valuable feedback.

#### REFERENCES

- [1] J. Day, I. Matta, and K. Mattar, “Networking is IPC”: A Guiding Principle to a Better Internet,” in *Proc. of ReArch’08*, Madrid, SPAIN, Dec. 2008.
- [2] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis, “Locator/ID Separation Protocol (LISP),” March 2009, draft-farinacci-lisp-12.txt.
- [3] C. Perkins, “IP Mobility Support for IPv4,” Internet Engineering Task Force, RFC 3344, Aug. 2002. [Online]. Available: <http://www.ietf.org/rfc/rfc3344.txt>
- [4] B. Quoitin, L. Iannone, C. de Launois, and O. Bonaventure, “Evaluating the Benefits of the Locator/Identifier Separation,” in *MobiArch ’07*.
- [5] D. Farinacci and V. Fuller, “LISP Map Server,” March 2009, draft-fuller-lisp-ms-00.txt.
- [6] L. Iannone and O. Bonaventure, “On the Cost of Caching Locator/ID Mappings,” in *CoNEXT ’07*, 2007.
- [7] A. Medina, A. Lakhina, I. Matta, and J. Byers, “BRITE: An approach to universal topology generation,” in *Proceedings of MASCOTS*, 2001.
- [8] V. Paxson, “End-to-end Routing Behavior in the Internet,” *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 5, pp. 41–56, 2006.