

Assessing the Security of a Clean-Slate Internet Architecture

Gowtham Boddapati[‡] John Day[†] Ibrahim Matta[†] Lou Chitkushev[‡]

[‡]Metropolitan College [†]College of Arts & Science
Computer Science, Boston University
{gowtham, day, matta, ltc}@bu.edu

Paper #26 (6 pages)

Abstract—The TCP/IP architecture was originally designed without taking security measures into consideration. Over the years, it has been subjected to many attacks, which has led to many patches to counter them. Our investigations into the fundamental principles of networking have shown that carefully following an abstract model of Interprocess Communication (IPC) addresses many problems [1]. Guided by this IPC principle, we designed a clean-slate Recursive INternet Architecture (RINA) [2]. In this paper, we show how, without the aid of cryptographic techniques, the bare-bones architecture of RINA can resist most of the security attacks faced by TCP/IP and of course is only more secure if cryptographic techniques are employed. Furthermore, the RINA model indicates specifically where those security measures reside. We also show how hard it is for an intruder to compromise RINA. Then, we show how RINA inherently supports security policies in a more manageable, on-demand basis, in contrast to the monolithic one-size-fits-all approach of TCP/IP.

I. INTRODUCTION

The TCP/IP architecture has shown signs of weakness as the Internet has grown and evolved. These problems are partly due to changing requirements—including mobility, quality-of-service, and security—but partly because of the architecture’s rigid one-size-fits-all structure. In this paper, we focus on the security properties that are inherent in an internet architecture.

As is often lamented, the TCP/IP architecture was originally designed without taking security considerations into account. Over the years, many vulnerabilities have been discovered and led to many patches to counter them. Given its rigid structure, security mechanisms have mostly been inserted into TCP/IP as “shim” sublayers, lacking a comprehensive approach to security and its reliance on overloading functions.

Most recently, there have been attempts to design *clean-slate* internet architectures. Our own investigations into the fundamental principles of communication led to a rather simple, elegant model based on a generalization of Inter-Process Communication (IPC). However, this model, referred to as RINA (Recursive INternet Architecture) [2], was developed from IPC considerations alone, without explicitly considering security. Hence, it seemed wise to investigate its security properties at the outset.

Space does not allow us to consider all aspects of the security of RINA in this paper. (We hope to cover other aspects in subsequent papers.) Here, after a very brief overview of the

pertinent aspects of RINA, we consider four types of vulnerabilities that have been found in TCP/IP: port-scanning attacks, connection-opening attacks and data-transfer attacks. What we find is that unlike the TCP/IP architecture, without the aid of cryptographic techniques, the bare-bones architecture of RINA is more secure and resistant to these attacks, even if we assume that a RINA network has been fundamentally compromised. Though further analysis is to be conducted, this might suggest that good design is as important to good security as explicit consideration of security.

We also show how RINA’s model organizes cryptographic techniques as building blocks that can be recursively applied on-demand, rather than the piecemeal approach of TCP/IP.

The rest of the paper is organized as follows. Section II reviews elements of TCP/IP and RINA that are most relevant to the security aspects discussed in this paper, specifically access control, addressing, and connection management. Section III compares the resiliency of TCP/IP and RINA to transport attacks, namely port-scanning, connection-opening and data-transfer. Section IV compares the two architectures in terms of their organizational support for diverse security policies. Section V concludes the paper.

II. BACKGROUND: TCP/IP VS. RINA

Figure 1 illustrates the TCP/IP architecture. In [2], we identified the shortcomings of this architecture and attributed them to: (1) exposing addresses to applications, (2) artificially isolating functions of the same scope¹, and (3) artificially limiting the number of layers (levels).

Figure 2 illustrates our RINA architecture [2], which leverages the inter-process communication (IPC) concept.² In an operating system, to allow two processes to communicate, IPC requires certain functions such as locating processes, determining permission, passing information, scheduling, and managing memory. Similarly, two applications on different endhosts should communicate by utilizing the services of

¹Transport and routing/relaying are split into two layers: Data Link and Physical layers over the same domain/link, and Transport and Network layers internet-wide.

²We use IPC in its long lost original sense of passing data messages between processes, rather than the broader current sense used today that encompasses this as well as all synchronization techniques.




Fig. 1. TCP/IP Architecture

a distributed IPC facility (DIF). A DIF is an organizing structure—what we generally refer to as a layer. What functions constitute this layer, however, is fundamentally different. A DIF is a collection of IPC processes (nodes). Each IPC process executes routing, transport and management functions. IPC processes communicate and share state information. How a DIF is managed, including addressing, is hidden from the applications. To understand why we organize layers this way, see [1].

The goal of a DIF is to provide a distributed service that allows application processes to communicate, one use might be as a private network or overlay. Two novel aspects of a DIF is that it *repeats* and is *relative*. Each repetition addresses a different range of operation and/or scope. As shown in Figure 2, two IPC processes A and B in an N-level DIF communicate by utilizing the services of an (N-1)-level DIF. Thus, while the specific function of IPC processes is to do IPC, they are also application processes requesting IPC from a lower layer. Our IPC-based architecture can be found in [2]. In this section, we only highlight key aspects of this architecture that have a fundamental impact on security.




Fig. 2. RINA Architecture

A. Access Control

Unlike TCP/IP, RINA requires explicit enrollment for an IPC process within a system to either join an existing DIF, or create a new DIF.

a) *Adding a New Member to an (N)-DIF*: Suppose that DIF A consists of a number of IPC processes on a set of systems. Suppose that DIF B wants to join DIF A, and that DIF B represents a single IPC process. The IPC process, b , in B knows the application (service) name of an IPC process, a , in A, not its address — B has no way of knowing the address of

any process in A. A and B are connected by an underlying (N-1)-DIF³. Using the underlying (N-1)-DIF, b requests that the (N-1)-DIF establish an IPC channel (connection) with a using the application name of a . In RINA, application processes incorporate a common protocol for establishing application connections that includes a plug-in module for authentication.

The (N-1)-DIF determines whether a exists and whether b has access to a . After the connection has been established, a authenticates b and determines whether it can be a member of A. This authentication can be as strong or as weak as required by the DIF. If the result is positive, a assigns an (N)-address to b . Note that the address is taken from the name space for DIF A, i.e., DIFs have their own name (address) space. b uses the (N)-address to identify itself to other members of DIF A. Other initialization parameters associated with DIF A are exchanged with b . The IPC process, b , is now a member of DIF A.

b) *Creating a New DIF*: Creating a new DIF is a simple matter. A management or similar application with the appropriate permissions causes an IPC process to be created and initialized, including pointing it to one or more (N-1)-DIFs. As part of its initialization, the IPC process is given the means to recognize allowable members of the DIF (e.g., a list of application process names, a digital signature, and so on). It might be directed to initiate enrollment with them or to simply wait for them to find this initial IPC process. When this has been achieved, adding more members to the DIF proceeds as described earlier.

B. Addresses and their Binding

The TCP/IP architecture has a global addressing space, which allows any system to freely connect to any other system. On the contrary, in RINA, the addresses are *internal* to a DIF. For two application processes to communicate, they have to have access to a DIF in common. If there is no common DIF, then one must be created either by joining an existing DIF or creating a new one. This provides the opportunity to restrict access based on the security policy of the DIF.

In the TCP/IP protocol suite, TCP overloads the port-id to be both port-id and connection-endpoint-id. And by overloading it again by giving it application semantics as a well-known port forces the receiver to rely on the sender's id information rather than ids it generated, which makes it easier the intruders to guess the port and thwarts any consistency checking by the receiver. Unlike TCP/IP, in RINA, applications do not listen to a well-known port. Rather an application process requests service using the destination application-name.⁴ The local IPC process returns a port-id with only local significance to the user to use as an opaque handle. The request is translated into a set of policies for an EFCP (Error and Flow Control Protocol) flow and instantiates this end of the flow by creating an EFCP-instance, identified by a different local identifier, referred to as a connection-endpoint-id (CEP-id). The local IPC process

³Ultimately the lowest level DIF is the physical medium.

⁴Application (service) names are global, but node (process) addresses are internal to the DIF.

then issues a create-request to find the destination application and allocate the flow.

When the IPC process at the destination gets the create-request, it determines if it can accept the request. The degree of access control is a matter of policy — it could be quite elaborate, or weak like the current Internet. If the request is accepted, the destination IPC process instantiates an EFCP-instance with its own local CEP-id, and the result is returned to the requesting application. The source and destination CEP-ids are concatenated for use as a connection or flow id. If the create-request returns with a negative response, it is determined whether the cause is fatal or not. If not fatal, the source IPC process may modify the request and try again. If the create-request returns with a positive response, the CEP-id is bound to the port-id. Figure 3 illustrates RINA's management of data-transfer connections.




Fig. 3. RINA connections

C. Data Transfer

RINA uses a soft-state data transfer protocol, built around Watson's Delta-t protocol [3]. This is in contrast to the hybrid hard-state/soft-state approach of TCP. In Delta-t, unless refreshed by data/ACK packet arrivals, a flow state is deleted after $2 \times MPL$ (Maximum Packet Lifetime) at the receiver (Figure 4), and $3 \times MPL$ at the sender (Figure 5). TCP, on the other hand, requires explicit control messages to establish and close connections. This makes TCP more vulnerable to attacks that fabricate its connection-management messages, or cause them to be dropped [4].




Fig. 4. RINA Receiver

III. TRANSPORT ATTACKS ON TCP/IP VS. RINA

A. Port-Scanning Attacks

Port scanning is often viewed as a first step for an attack, wherein the attacker explores “open” ports to which processes




Fig. 5. RINA Sender

on a system are listening. In RINA, a service is accessed by its application-name—the requesting applications never see addresses nor CEP-ids. In fact they are not privy to any data transfer identifiers. This is in contrast to TCP/IP in which a destination application process is assumed to listen to a well-known port. RINA also supports local access control domains that restrict which applications are visible to the DIF that the requestor belongs to. As described earlier, source and destination port-ids then get assigned locally on-demand. Ports are also dynamically mapped to separate data-transfer (connection) endpoints, and contrary to TCP/IP, ports are not part of the flow/connection id. This makes traditional port-scanning attacks not possible in RINA.

In RINA, however, the attacker might try to guess application names. But this is more difficult because application names are of variable length String, a far larger name space. Furthermore, the malicious user has to be a member of the same DIF to be able to address other members in the DIF. Joining a DIF requires that the new IPC process be authenticated, providing further barriers to compromise RINA.

B. Connection-Opening Attacks

In this type of attack [5], the intruder attempts to establish a connection with the server, impersonating a trusted user A.

In TCP/IP, this attack exploits the explicit three-way handshake of TCP in which the client and server exchange (synchronize) their Initial Sequence Numbers (ISN) prior to data transfer. A malicious handshake sequence with server S, intruder X, and spoofed client A, may look like:

```

X  →  S : SYN(ISNx), SRC= A
S  →  A : SYN(ISNs), ACK(ISNx)
X  →  S : ACK(ISNs), SRC = A
X  →  S : ACK(ISNs), SRC = A , malicious-data

```

In this attack, we assume that the attacker X already knows the destination port and IP address, as well as the source IP address. The destination port and IP address are easy to obtain, as they are generally published, as well-known ports. The source IP address is also generally easy, as this is simply the client that is being spoofed. As this is a connection establishment phase, the intruder can use any one of the ports available as source port-id. This attack also assumes that the acknowledgment (ACK) sent by the server and destined to the spoofed system A, is lost or delayed, either because A itself

was down or slow (possibly through a separate attack) or the ACK is intercepted and dropped by the intruder X.

The difficult part of launching this attack is determining the ISN of the server. This could be more easily obtained if the intruder is in the middle and observes the (unencrypted) traffic between A and S. Otherwise, the intruder has to guess ISNs, which given 32-bit sequence numbers and random selection of ISNs, involves 2^{32} possibilities.

In TCP/IP the packet following the three way handshake can be any arbitrary packet. Which can lead to attacks such as connection opening attack. On the other hand, in RINA the IPC process involved in communication uses the ACSE protocol for establishing and releasing application connections. By using this protocol the receiver expects the authentication packets to be followed after the connection establishment phase but not some arbitrary packets which greatly reduces the risk of connection opening attacks. In RINA, the intruder also has to predict the server's CEP-id, which gets dynamically allocated by a resource-allocation protocol.

```

X  → S : create-request(service-name, A, S,
                        source CEP-id, QoS, ...)
S  → A : creat-response(OK, destination CEP-id, ...)
X  → S : destination CEP-id, ISNc, ...
S  → A : Challenge (ACSE), ...
X  → S : response, ...
X  → S : data

```

If we assume that the intruder X has somehow thwarted the enrollment authentication described earlier as well as the authentication after connection establishment and is a member of the DIF as are A and S, but note that these are the hurdles that a TCP intruder does not need to overcome. In this case, X is able to know the addresses of A and S, *i.e.*, X is launching an insider attack. As this is a connection establishment phase, the intruder can use any source CEP-id. Also, in RINA, since there is no need for synchronizing sequence numbers [3], the sender can use any initial sequence number. Assuming X does not observe the reply with the destination CEP-id, it has to guess this CEP-id. Furthermore, since RINA uses a soft-state approach, CEP-ids are only allocated for $2 \times MPL$. Thus, intruder X has to successfully guess the destination CEP-id within a limited time window of $2 \times MPL$. Given 16-bit CEP-ids, this involves 2^{16} possibilities. This makes this type of attack equivalent to port-scanning attacks, in which an intruder may be attempting an unallocated destination CEP-id. Such attacks raise more suspicion (and hence, are easier to detect) than TCP attacks that guess ISN.

Should the intruder be successful in guessing an acceptable CEP-id, the application expects the common application protocol to be the data, followed by the application's authentication procedure. The intruder will have difficulty inserting its malicious data.

C. Data-Transfer Attacks

Data-transfer attacks, known as *blind in-window attacks* [6], are those where the attacker does not have access to the data packets of the victim connection but still attempts to inject packets that seem legitimate. Forming a legitimate packet requires guessing various fields in the packet's header.

In TCP/IP, the goal of this type of attack might be to abort an ongoing connection by injecting a "reset" TCP control packet [6]. The damage depends on the application running above the TCP connection. One such application is BGP, where a connection abort would result in entries of the routing table being flushed. In this attack, we assume that the attacker knows the destination port and IP address, as well as the source IP address. The destination port and IP address are easy to obtain, as they are published. The source IP address is also generally easy, as this is simply the spoofed client. The intruder has to guess the source port as well as the sequence number that has to lie within the window of the receiver.

To guess the source port-id, given 16-bit port numbers, we have at most 2^{16} possibilities. Furthermore, for each possible source port-id, given 32-bit sequence numbers and say 64KB window size⁵, we have $\frac{2^{32}}{2^{19}} = 2^{13}$ possibilities for selecting a sequence number that lies within the current receiver's window. Thus, there is a total of $2^{16+13} = 2^{29}$ possibilities. Note that for larger window sizes⁶, typical of higher bandwidth-delay-product networks, the attack will be easier to launch.

In the case of RINA, the intruder can launch an attack during two different phases of a connection: (1) after the resource-allocation request is complete and before the data transfer phase starts, or (2) during the data transfer phase. Again here we assume that the intruder is in the same DIF, so the attacker knows the addresses of the source and destination IPC processes. In RINA Model even though the attacker is successful in exploiting the Data transfer phase, he would be getting hold of the pipe which runs through the connection endpoints of the two IPC process but not the pipe which runs from application to application as in the case of TCP. In TCP the connection runs from application to application. Figure 6 and Figure 7 illustrates RINA and TCP management of data-transfer connections.




Fig. 6. RINA Connection

⁵64KB is the default TCP maximum window size, without window scaling options.

⁶Larger window sizes are possible using window scaling options.




Fig. 7. TCP Connection

In the first case, the attacker has to guess the source CEP-id and the destination CEP-id. The attacker also has to guess other agreed-upon parameters of the connection, such as the QoS-id, though as a member of the DIF, he/she knows the legal range of QoS-ids. Since the data transfer phase has not started, the attacker can use any ISN. Given 16-bit port numbers and 8-bit QoS-id, the attacker has $2^{16+16+8} = 2^{40}$ possibilities for guessing the CEP-ids and QoS-id for the victim connection.

In the second case, in addition to the CEP-ids and QoS-id, the attacker has to guess the sequence number which falls within the window of the server. This guessing involves $2^{40+13} = 2^{53}$ possibilities, assuming 64KB window size. This type of attack is made even harder because of RINA's use of a Data-Run-Flag (DRF) during data transfer [3]. If the DRF bit is set, this implies that the sender has no data left to be acknowledged or it is starting a new data run. Thus, the DRF bit periodically synchronizes the sender and receiver, and so setting it incorrectly in the attack packet would raise suspicion.

For example, if the DRF bit is not set and the receiver's connection state had timed out (because it has not been refreshed by new data from the sender), the packet is simply dropped by the receiver. Let's then assume that the attacker always sets the DRF bit, along with an arbitrary sequence number, in its attack packet. This attack packet is accepted only if the receiver had no state for this connection. Otherwise, the receiver can verify whether the setting of the DRF bit makes sense, which is the case only if the receiver has indeed acknowledged all prior packets.

Finally, this type of attack is not possible or harder to launch in RINA for two reasons: (1) RINA uses a soft-state approach in managing connections, thus it does not use explicit connection "reset" messages, which precludes "reset" attacks,⁷ and (2) RINA supports the dynamic assignment of CEP-ids during the lifetime of a connection, binding them to the same port-ids that are only locally-visible. This would make it very hard for an attacker to guess the source and destination CEP-ids.

D. Blind TCP data injection through fragmented IP traffic

Zalewski described a possible attack that can be performed on TCP/IP that doesn't require the attacker to guess or know

the aforementioned TCP connection parameter and could therefore be successfully exploited in some scenarios with less effort than that required to exploit the more traditional data-injection attacks.

The attack is performed when one system is transferring information to a remote peer by means of TCP, and the resulting packet gets fragmented. In that case, the first IP fragment will usually contain the entire TCP header, including port numbers, sequence number, and other information that may be relatively difficult for a third party (the attacker) to guess otherwise. The other fragments carry the remaining section of the TCP/IP payload, and would be put back together with the headers and previous sections of the packet once received at the receiving side. The attacker may spoof the second IP fragment, instead of attempting to determine the sequence number, and insert data into the TCP payload.

In order to successfully exploit this attack, the attacker would be facing with two problems, the first part is figuring out the IP identification values. Usually a minor inconvenience, since a majority of systems use sequential numbers, and so it is possible to guess the next value with no effort and the other problem is sending a fragment that would, after reassembly, still validate against TCP/IP checksum in the headers. If the attacker knew the data being transferred over the target connection, then the only real unknown is the sequence number in there - the remainder can be usually either predicted to a degree, or simply overwritten with overlapping fragments.

This problem arises because in the TCP/IP architecture fragmentation/reassembly is done by both IP and TCP. In RINA, because the Relaying/Multiplexing function and Error and Flow Control functions are in the same layer fragmentation/reassembly occurs only once.

Table I summarizes our comparison of RINA against TCP/IP under transport-level attacks. We assume 32-bit sequence numbers, 16-bit port-ids/CEP-ids, 64KB window size, and 3-bit QoS-id. To be able to make a direct comparison, we had to assume that a RINA network had been compromised and a rogue member had been allowed to join—a hurdle that is not present in TCP/IP networks.

IV. SECURITY POLICIES IN TCP/IP VS. RINA

RINA decouples the various security functions of authentication and confidentiality/integrity. The former is done at the top of the DIF where applications of the DIF authenticate each other. The latter is done at the bottom of the DIF where the IPC processes encrypt their traffic if they do not trust the lower DIFs. These security functions are applied recursively, so IPC processes themselves would authenticate each other when communicating through lower-level DIFs. Policies of the DIF determine the levels of authentication and encryption. Figure 8 illustrates this functional organization.

In contrast, TCP/IP implements security functions piecemeal, for example, using SSL under the application layer and IPsec below the network layer. The TCP/IP organizing structure is rigid and can only accommodate security functions

⁷In a soft-state approach, the connection's state at the receiver is automatically reset after $2 \times MPL$ if not refreshed by the sender [3], thus there is no need for an explicit "reset" message.

TABLE I

COMPARISON OF TCP/IP AND RINA UNDER TRANSPORT ATTACKS. TO BE ABLE TO MAKE A DIRECT COMPARISON, WE HAD TO ASSUME THAT A RINA NETWORK HAD BEEN COMPROMISED AND A ROGUE MEMBER HAD BEEN ALLOWED TO JOIN—A HURDLE THAT IS NOT PRESENT IN TCP/IP NETWORKS.

Vulnerability	TCP/IP	RINA
Port-scanning	possible due to well-known ports	not possible with unknown CEP-ids
Connection-opening	2^{32} possibilities to guess ISN	2^{16} possibilities to guess destination CEP-id within $2 \times MPL$
Data-transfer (right after conn. open)	2^{29} possibilities to guess source port-id and valid SN	2^{40} possibilities to guess source and destination CEP-ids and agreed-upon QoS-id
Data-transfer (after transfer started)	2^{29} possibilities to guess source port-id and valid SN	2^{53} possibilities to guess source and destination CEP-ids, agreed-upon QoS-id, and valid SN
DoS	possible due to well-known ports and no access control	thwarted with access control and dynamic construction of new DIFs

as “shim” sublayers, rather RINA accommodates them as an integral part of (recursive) inter-process communication.




Fig. 8. Security policies applied recursively

Figure 9 illustrates a “middlebox” solution to enable the support of “private” domains in TCP/IP. Such a middlebox is known as Network Address Translator (NAT) since it aggregates private addresses of systems inside the private domain (such as system “B” in the figure) into the NAT public address. Communication across the private domain and the public (Internet) domain, say between systems “B” and “A”, is done through the NAT, which translates between its public NAT address and port number, which identifies “B” externally, and B’s actual private address and port number. Furthermore, the NAT acts as a firewall, preventing attacks on private addresses and ports. However, it is clear that this kind of hand-crafted arrangement makes it hard to coordinate communication across domains when we want to.




Fig. 9. Security through NATs in TCP/IP

Figure 10 illustrates the procedure in RINA, where communication is established between application processes to

join the same DIF. First, process “B” joins DIF z, which initially only contains process “C” (Figure 10(a)). As mentioned earlier, this explicit enrollment procedure happens using a common underlying DIF (DIF y, in this example), and involves authenticating that B is a valid member of DIF z, initializing it with current DIF information, and assigning B an internal address for use in coordinating communication within DIF z. Then, similarly, process “A” joins DIF z (Figure 10(b)). Thus, in RINA, there are no “middleboxes” per se, but rather systems join and leave DIFs as determined by management (security) policies. Furthermore, such enrollment procedures can be repeated horizontally to create concurrent DIFs, or vertically to create stacked DIFs.




Fig. 10. (a) Process “A” is about to join DIF z, (b) Process “A” after joining DIF z.

V. CONCLUSION

In this paper, we compare a clean-slate internet architecture, RINA, that is based on fundamental IPC principles, to TCP/IP in terms of architectural support for security. We specifically compare the resiliency of RINA to security vulnerabilities found in the TCP/IP architecture. In some cases, to make a

fair comparison, we had to assume that a RINA network had been compromised and a rogue member had been allowed to join. (A hurdle that is not present in TCP/IP networks.) Even so, we found RINA to be more secure and resistant to these attacks.

We focused on access control, addresses and their binding, and data transfer. We contrast the open access of TCP/IP to the controlled access of RINA, which requires an *explicit* enrollment phase to join a network of IPC processes (DIF). Unlike TCP/IP, in RINA, node addresses (of IPC processes) are internally assigned by a DIF, and are not exposed to application processes. Furthermore, data connections are *dynamically* assigned connection endpoint ids (CEP-id), which are bound to dynamically assigned ports. This late (dynamic) binding of addresses / ids provides levels of indirection that make RINA inherently more secure than TCP/IP, which exposes static addresses and port numbers to applications.

We compare the resiliency of RINA and TCP/IP to transport-level attacks. We show how the static assignment of addresses and ports, as well as the hard-state approach of TCP/IP to managing connections, makes TCP/IP quite vulnerable to port-scanning, connection-opening, and data-transfer attacks. On the other hand, the dynamic assignment of addresses and ports, the decoupling of port numbers from CEP-ids, and the soft-state approach to managing connections, makes RINA quite resilient to such attacks. We believe that this is an interesting result, given that no more consideration of security was present in the development of RINA than in the development of the TCP/IP architecture. One might be led to conclude that strong design is as important to good security as explicit consideration of security. In other words, TCP does not suffer as much from a lack of foresight as a weak design.

We also show that RINA's support for dynamic construction of DIFs, with their own security and resource-allocation policies, enables RINA to effectively thwart DoS attacks. Furthermore, Distributed DoS (DDoS) attacks are even harder to mount, since a large-scale DDoS would require the intruder to join (and be authenticated by) many DIFs—many more hurdles to overcome!

Finally, we argue that the recursive nature of RINA organizes the security policies in a clean way, decoupling authentication from integrity and confidentiality.

REFERENCES

- [1] J. Day, *Patterns in Network Architecture: A Return to Fundamentals*. Prentice Hall, 2008.
- [2] J. Day, I. Matta, and K. Mattar, ““Networking is IPC”: A Guiding Principle to a Better Internet,” in *Proceedings of ReArch'08 - Re-Architecting the Internet*. Madrid, SPAIN: Co-located with ACM CoNEXT 2008, December 2008.
- [3] R. Watson, “Timer-Based Mechanisms in Reliable Transport Protocol Connection Management,” *Computer Networks*, vol. 5, pp. 47–56, 1981.
- [4] G. Gursun, I. Matta, and K. Mattar, “On the Performance and Robustness of Managing Reliable Transport Connections,” CS Department, Boston University, Tech. Rep. BUCS-TR-2009-014, April 17 2009. [Online]. Available: <http://www.cs.bu.edu/techreports/pdf/2009-014-reliable-conn-mgmt.pdf>
- [5] S. M. Bellovin, “Security Problems in the TCP/IP Protocol Suite,” *Computer Communication Review*, vol. 19, no. 2, pp. 32–48, 1989.

- [6] P. Watson, “Slipping in the Window: TCP Reset attacks,” Presentation at 2004 CanSecWest, 2004, <http://cansecwest.com/csw04archive.html>.