On Balancing Load to Quickly Detect and Stop Attack Traffic

Nabeel Akhtar* Marzieh Babaeianjelodar[x] Ibrahim Matta* Yaoqing Liu[x]
*Boston University [x]Clarkson University

Abstract

- Resources are elastically managed using Software Defined Networking (SDN), Network Function Virtualization (NFV) and Control theory
- Control theory (PI controller), along with SDN, is used to balance load across different VNF instances running Snort-IDS
- RINA [1] management architecture is used to monitor VNF instances over the GENI testbed
- A generalized framework using **Attack Analyzer** is used to analyze different types of attacks
- Attacks are detected faster with load balancer

What is RINA? [1]

- RINA: Recursive InterNetwork Architecture
- A clean-slate network architecture
- Based on the fundamental principle that *networking is Inter-Process Communication (IPC) and only IPC*
- Distributed IPC Facility (DIF): a collection of distributed IPC processes with shared states
- Distributed Application Facility (DAF): a set of application processes cooperating to perform a certain function
- Two design principles: (i) divide and conquer (recursion), and (ii) separation of mechanisms and policies

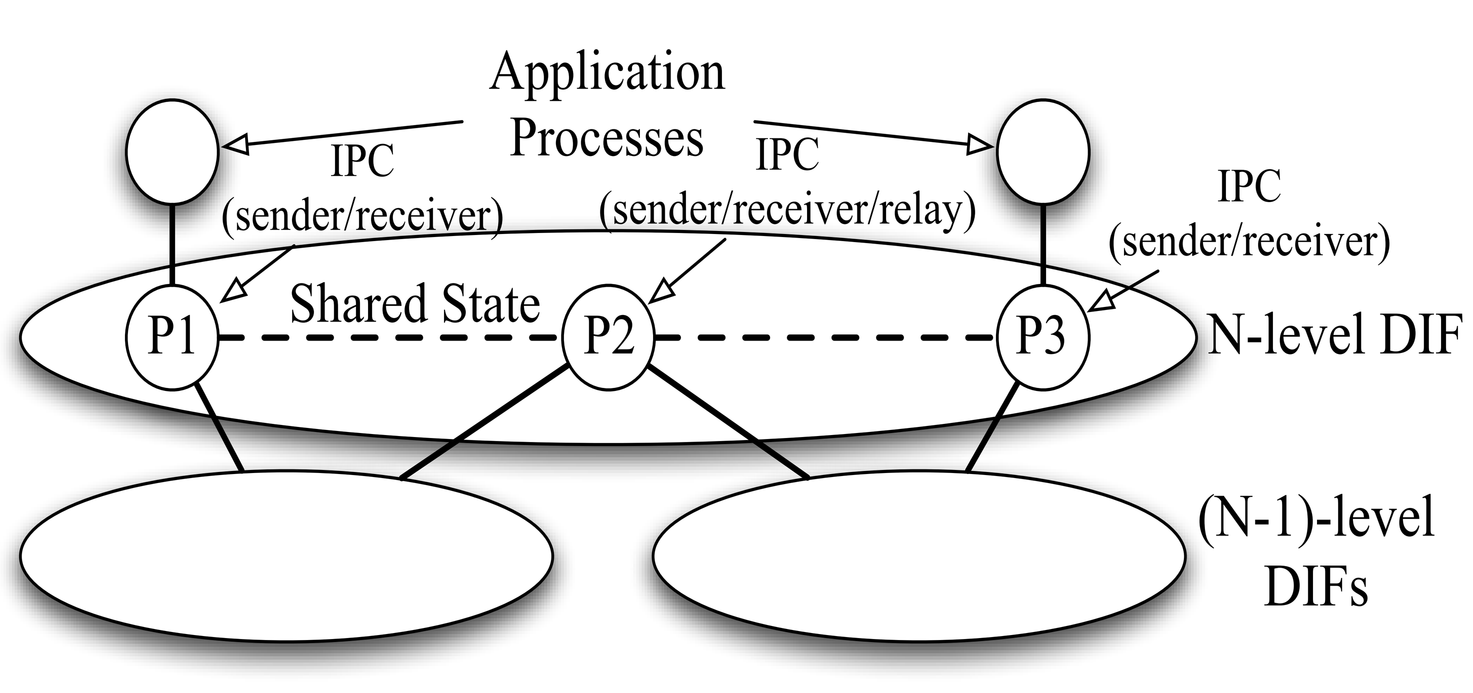

Fig. 1: RINA overview

References

[1] Boston University RINA Lab. http://csr.bu.edu/rina

Experiment over GENI

- Experimental setup (Fig. 2):
  - two VNF instances running Snort IDS (*VNF1* and *VNF2*)
  - one OVS switch and one open-flow controller
  - two sources (*S1* and *S2*) and one destination (*destination*)
- Traffic is sent to Snort-IDS running on *VNF1* or *VNF2*
- RINA management architecture is used to send load and Snort-IDS alerts of VNF instances to *Controller*
- **Load Balancer** determines the fraction of traffic to divert from VNF1 to VNF2 and updates the OVS controller
- **Attack Analyzer** processes Snort-IDS alerts and updates the attacker list for the OVS controller
- OVS controller updates OpenFlow rules on the OVS switch based on Load Balancer and Attack Analyzer input

Control Theoretic Load Balancer

- PI controller (Fig. 3):

$$x(t) = \max[0, \min[1, x(t-1) + K(\tfrac{L(t)}{T} - 1)]]$$

$x(t)$: ratio of traffic diverted to VNF2 at time t

$L(t)$: load on VNF1

T: target load on VNF1

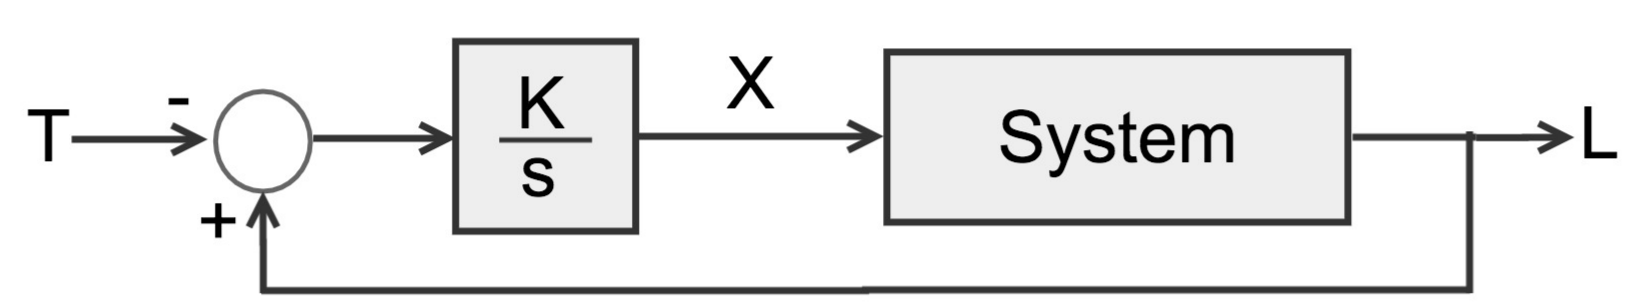

Fig. 3: System load $L(t)$ and target load T of VNF1 is used to compute $x(t)$, i.e. ratio of traffic diverted to VNF2

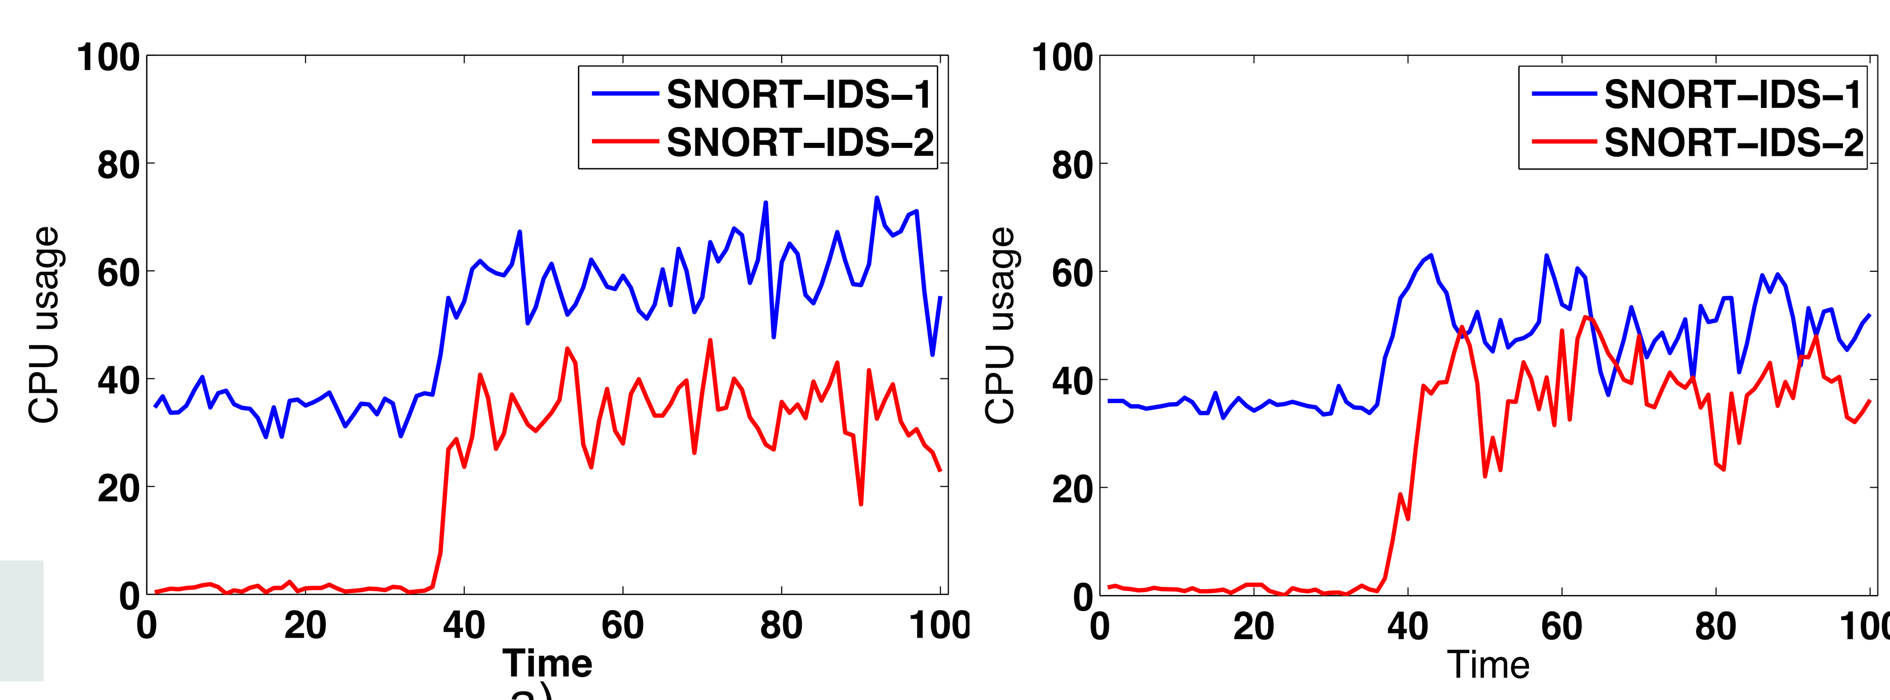

Fig. 4: (a) Simple Round Robin load balancing; (b) Load balancing based on PI control (T = 50%)

Management Architecture

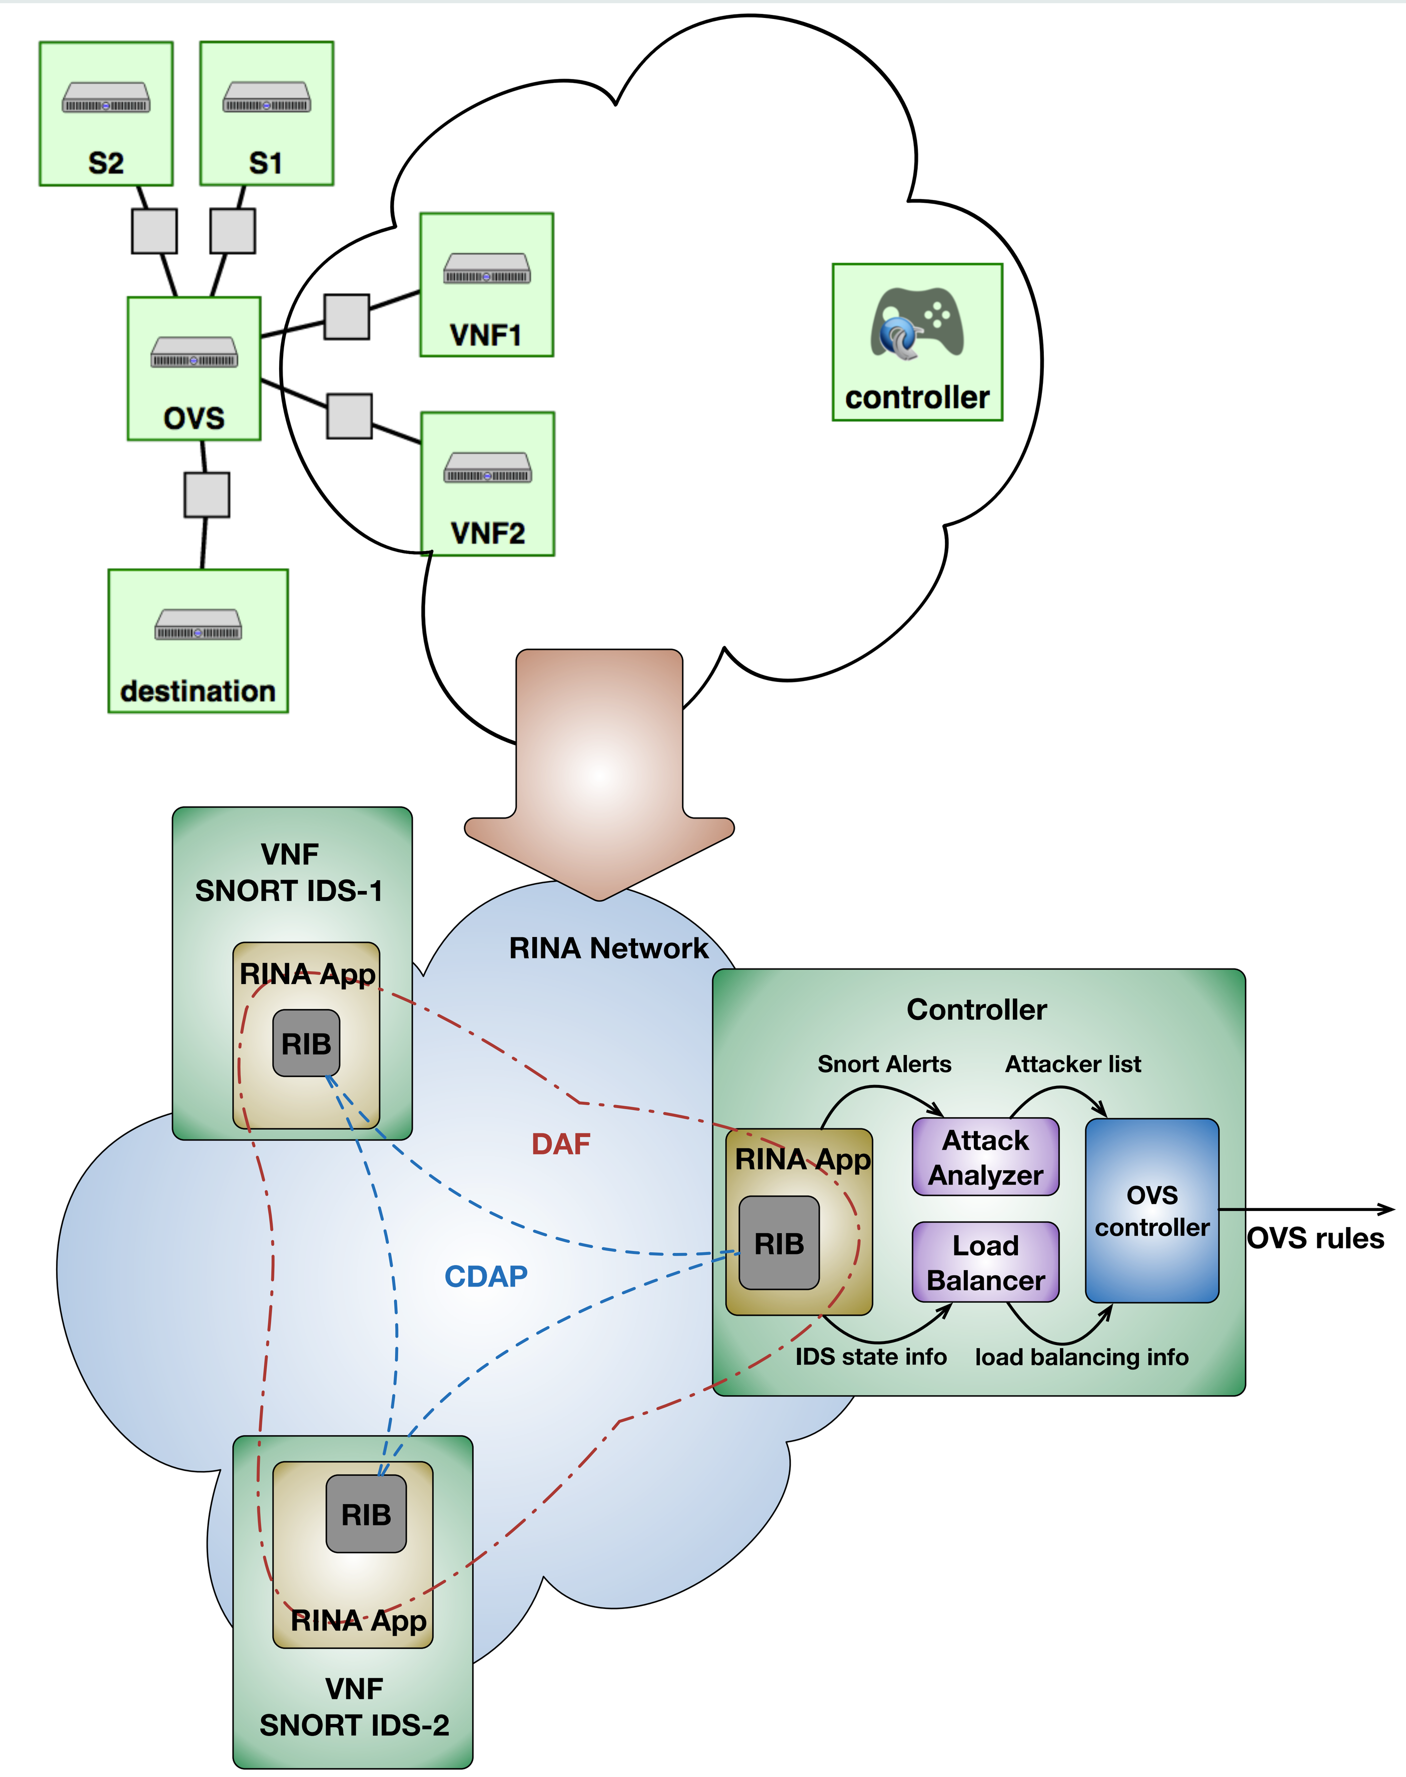

Fig. 2: RINA management architecture used for communication between VNF instances and Controller

Results

- On average, attacks are stopped significantly faster under load balancer
- Attacks sometimes go undetected without load balancer

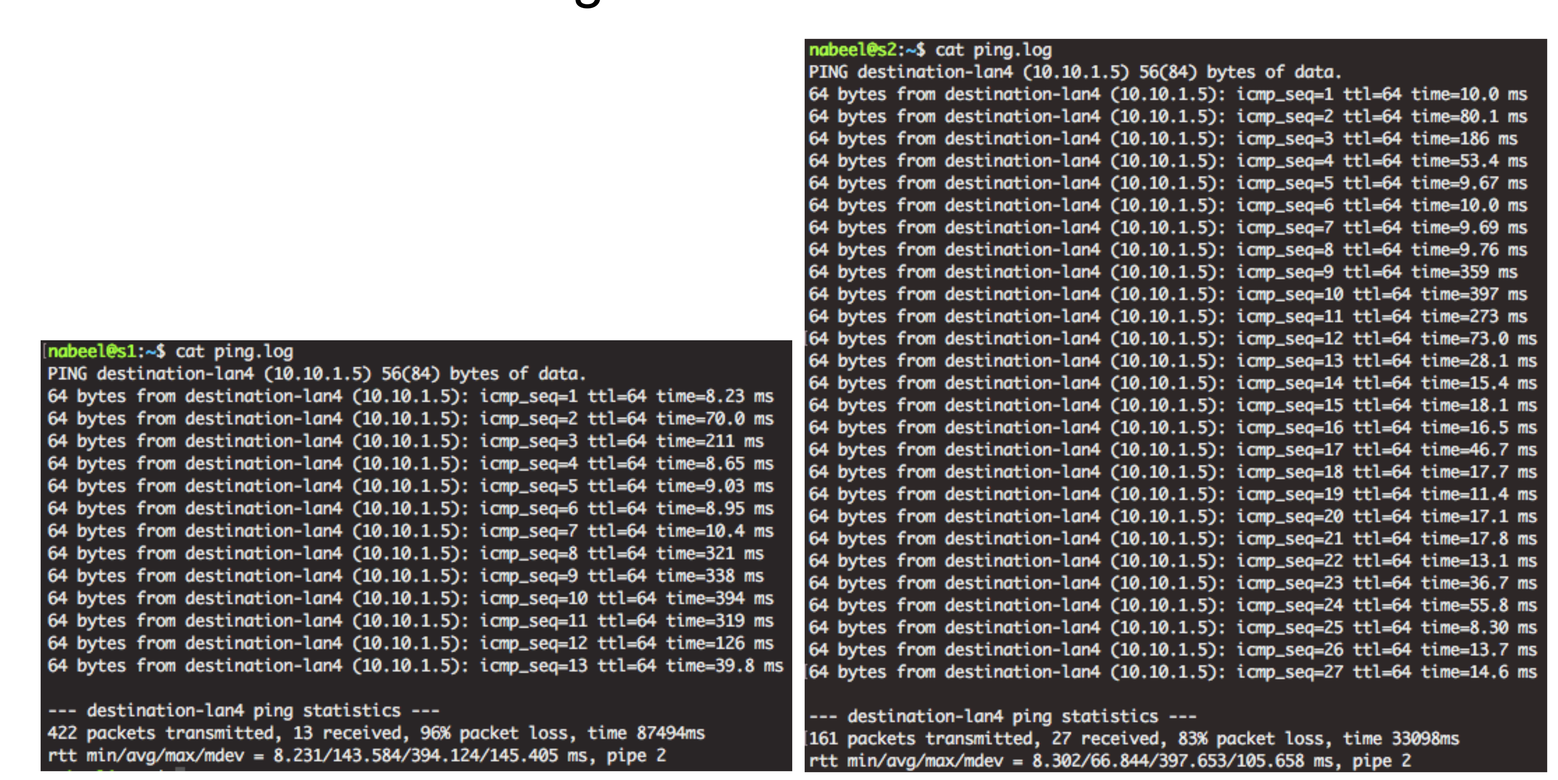

a) b)

Fig. 5: Port Scanning Attack (a) 2.6 seconds with load-balancer (b) 5.4 seconds without load-balancer