ComputerScience

Boston University

# Application-Driven Network Management using ProtoRINA

Yuefeng Wang  Ibrahim Matta  Nabeel Akhtar

## What is RINA? [1][2]

- RINA: Recursive InterNetwork Architecture
- A clean-slate network architecture that overcomes inherent weaknesses of the current Internet, e.g., security, mobility support
- Based on the fundamental principle that *networking is Inter-Process Communication (IPC) and only IPC*
- Distributed Application Facility (DAF): a set of application processes cooperating to perform a certain function, such as communication service, management service, *etc*.
- Distributed IPC Facility (DIF): a collection of distributed IPC processes with shared states. They provide communication service to application processes over a certain scope (i.e., range of operation)
- DIF is a special DAF whose job is to provide communication services

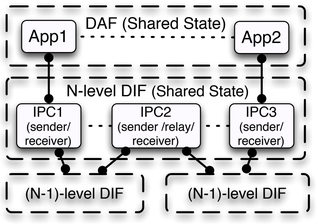


Fig 1: RINA overview

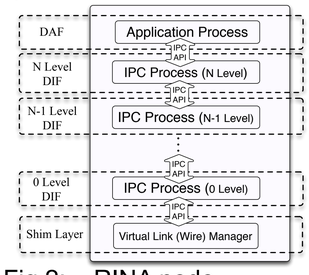


Fig 2: RINA node

## ProtoRINA: A RINA Prototype [3]

- ProtoRINA is Boston University's user-space prototype of RINA
- Experimental tool for developing (non-IP based) user and management applications, and teaching tool for networking and distributed systems classes
- A *RINA node* is a host (or machine) where application processes and IPC processes reside, and each RINA node has a *DIF allocator* which manages the use of various existing DIFs

## Traditional Network Management

- Tied to the Internet architecture, and inevitably inherits various problems, such as:
  (1) static management and one-size-fits-all structure
  (2) ad-hoc mechanisms with no common management framework
- Software-Defined Networking (SDN) aims to provide better and more flexible management, but still suffers from aforementioned problems due to it being tied to the TCP/IP architecture [4]
- Some work (such as FlowVisor and ADVisor) has been done to support network virtualization based on application requirements, but their virtual network is limited to routing and not for transport purpose, and they do not support dynamic formation of virtual networks

## Application-Driven Network Management

- Definition: given the physical topology of the network, virtual networks can be built on the fly to satisfy application-specific demands and achieve better network performance
- With the development of new networking service models (such as Private Cloud as as Service or Software as a Service), as well as the demand for different SLAs (Service-Level Agreements), we believe application-driven network management is necessary and will become the norm
- In RINA, the DIF is such a virtual network, *i.e.,* a secure transport container providing inter-process communication. A DIF can be dynamically formed, where each DIF has its own scope, and they all use the same RINA mechanisms but can have different policies

## RINA's Management Architecture

- A common management framework: the DAF-based architecture
- Application processes providing management functionalities form different management DAFs, and the same DAF-based management structure repeats over different scopes
- For a single DIF, the *Management Application Entity* of each IPC process forms the management DAF, and this DAF manages a single DIF
- For a network made up of multiple DIFs, the *DIF allocator* of each RINA node forms the management DAF, and this DAF manages the whole network
- The DIF allocator is able to form new DIFs dynamically to meet various application requirements

## Network Management for Video Multicast

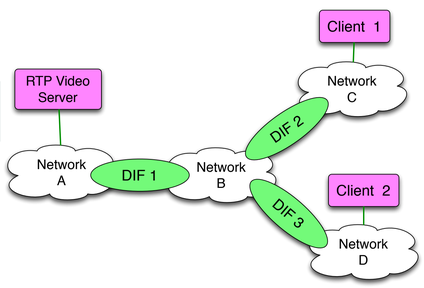


Fig 3: The whole network is made up of four enterprise networks (networks A, B, C, D). RTP video server provides a live video streaming service, and multiple clients, e.g., clients 1 and 2, want to access this service

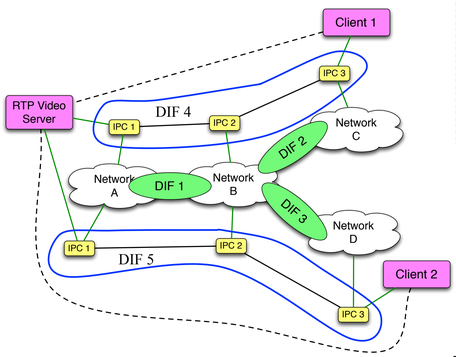


Fig 4: Video streaming can be done through unicast connections supported by two level-1 DIFs (DIFs 4 and 5), which consume unnecessary bandwidth over DIF 1

## Two Multicast Solutions

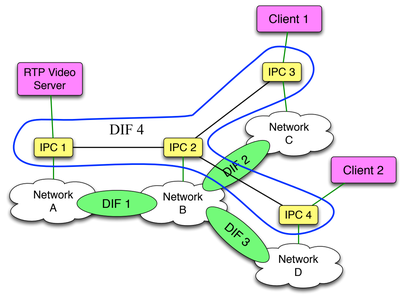


Fig 5: Video multicast through multicast service provided by level-1 DIF 4

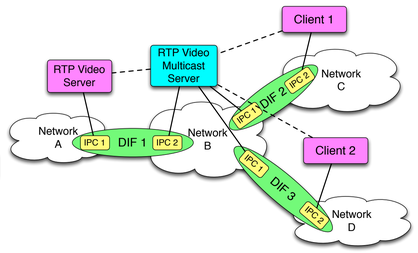


Fig 6: Video multicast through an RTP video multicast server which can placed closer to clients

## Experiments over GENI

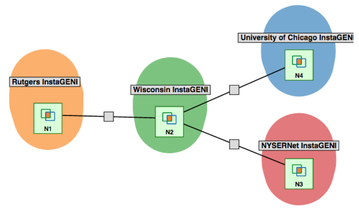


Fig 7: VMs from four InstaGENI aggregates connected through stitched VLANs

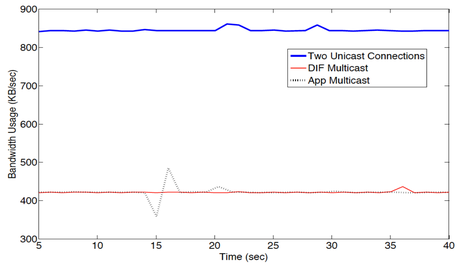


Fig 8: Comparison of bandwidth usage over DIF 1: unicast vs. multicast

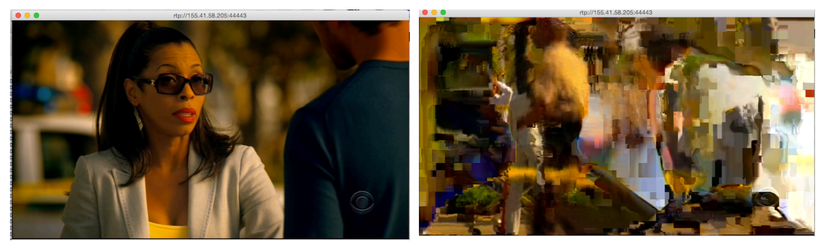

Fig 9: Video seen by clients when RTP video multicast server (cf. Fig 6) is placed at two different aggregates

## References

[1] John Day, Ibrahim Matta and Karim Mattar. "Networking is IPC: A Guiding Principle to a Better Internet". In ReArch 2008.
[2] Boston University RINA Lab. http://csr.bu.edu/rina.
[3] Yuefeng Wang, Ibrahim Matta, Flavio Esposito and John Day. "Introducing ProtoRINA: A Prototype for Programming Recursive-Networking Policies." In ACM SIGCOMM CCR, July, 2014.
[4] Yuefeng Wang and Ibrahim Matta. "SDN Management Layer: Design Requirements and Future Direction". In CoolSDN 2014.