A lifecycle experience of PolKA: From prototyping to deployment at Géant Lab with RARE/FreeRtr

Everson S. Borges^{1,2}, Edgard Pontes¹, Cristina Dominicini², Marcos Schwarz³, Csaba Mate⁴, Frederic Loui⁵, Rafael Guimarães², Magnos Martinello¹, Rodolfo Villaça¹, Moisés R. N. Ribeiro¹

¹Federal University of Espírito Santo (UFES) Av. Fernando Ferrari, 514, Goiabeiras – 29.075-910 – Vitória – ES – Brazil

²Federal Institute of Espírito Santo (IFES)

³RNP-National Research Network

⁴RARE - Router for Academia Research & Education

⁵GÉANT - Pan-European Research and Education Networks

{everson.borges,edgard.pontes}@edu.ufes.br

Abstract. In this paper, we take the position of developers that need to deploy emerging programmable protocols (or services) with specific network requirements and want to know how to benefit from an open router platform and testbed infrastructure. Our focus is to exploit the PolKA lifecycle experience as a use case to figure out a balance between integrating and reusing legacy protocols with new protocols.

Resumo. Neste artigo, assumimos a posição de desenvolvedores que precisam implantar protocolos (ou serviços) programáveis emergentes com requisitos de rede específicos e querem saber como se beneficiar da plataforma de roteador aberta e da infraestrutura de teste. Nosso foco é explorar o equilíbrio na integração e reutilização de protocolos legados com novos protocolos.

1. Introduction

Source routing (SR) is a prominent alternative to table-based routing for reducing the number of network states. In this approach, a source adds a route label in the packet header to define the paths along with the network. Segment Routing and PolKA [Dominicini et al. 2020] are two new approaches to the source routing paradigm that allow a flow to be routed to a specific topological path while maintaining a per-flow state only at the ingress node to the SR domain.

However, a significant challenge in developing a new protocol is designing and implementing it with its encapsulation representing an encoding abstraction that needs to be efficiently processed by the switches. For example, implement an integer modulo operation that is not natively supported in commodity or even programmable network hardware. Therefore, a network developer needs to use software switches [Martinello et al. 2014], or depend on synthesizing integer division to ASICs or NetFPGAs[Liberato et al. 2018].

In this paper, we take the position of developers that need to deploy emerging programmable protocols (e.g., PolKA[Dominicini et al. 2020]) with specific network requirements (e.g., mod operation). In addition, we want to know how to benefit from the open router platform and testbed infrastructure. Our focus is to exploit the balance of integrating and reusing legacy protocols with new protocols.

2. PolKA and RARE/freeRtr in a nutshell

2.1. PolKA- Polynomial Key-Based Architecture

Let us assume that a packet should be routed via a selected path, represented by N core nodes and their respective output ports. Let $S = \{s_1(t), s_2(t), \ldots, s_N(t)\}$ be the multiset of the polynomials representing the *nodeID*s of the nodes in this path. The set S must be composed of pairwise co-prime polynomials, and satisfy the condition $degree(s_i(t)) \ge$ $\log_2(nports)$, where *nports* denotes the number of ports in the node. For simplicity, we assume that $s_i(t)$ are irreducible polynomials. Let $O = \{o_1(t), o_2(t), \ldots, o_N(t)\}$ be the multiset of N polynomials, where $o_i(t)$ represents the output port for the packet at the core node $s_i(t)$, for $i = 1, 2, \ldots N$, satisfying the condition that $degree(s_i) > degree(o_i)$. For instance, if the output port polynomial is $o_i(t) = 1 \cdot t^2 + 1 \cdot t$, it maps the port 110 and the packet is forwarded to port label 6 at node $s_i(t)$. Based on the definition of the path represented by S and O, the Controller calculates the *routeID* using the polynomial CRT [Shoup 2009, Bajard 2007] as the polynomial R(t) that satisfies:

$$R(t) \equiv o_i(t)s_i(t), \quad for \quad i = 1, 2, \dots N$$
(1)

The *routeID* is embedded in the packet by the edge, and the forwarding operation in each core node calculates the output port as the remainder of the euclidean division of the *routeID* in the packet by its *nodeID*: $o_i(t) = \langle R(t) \rangle_{s_i(t)}$

Figure 1 shows an example of a path composed of three core nodes, which received their *nodeIDs* from the Controller in a network configuration phase: $s_1(t) = t+1 = 11$, $s_2(t) = t^2 + t + 1 = 111$, $s_3(t) = t^3 + t + 1 = 1011$. Considering the path $s_1 \rightarrow s_2 \rightarrow s_3$, the set O is composed by: $o_1(t) = 1$, $o_2(t) = t = 10$, $o_3(t) = t^2 + t = 110$. For this example, the *routeID*, calculated according to the polynomial CRT, is R(t) = 10000. The Controller may proactively compute this R(t) or calculate it when the first packet of a flow arrives. To configure the path, the Controller installs flow entries in the edges, which embed the *routeID* 10000 into the packets. Then, each node can calculate its *portID* by dividing the *routeID* of the packet by its own *nodeID*. For example, the remainder of R(t) = 10000 divided by $s_2(t) = 111$ is $o_2(t) = 10$ (port label 2).

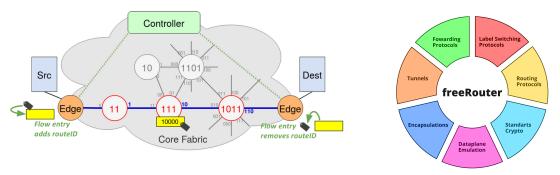




Figure 2. freeRtr.

2.2. RARE/freeRtr

RARE/freeRtr is a control plane platform shown in **Fig. 2**, which uses UNIX socket to forward packets and this is a key feature that is leveraged to connect the control plane **RARE/freeRtr** to any type of protocol plane legacy, such as OSPF, IPv4, IPV6, ATM, ethernet, etc. In **RARE/freeRtr** everything is in a VRF so there is no global VRF, this design choice has positive consequences like. There are no VRF recognition questions,

having multiple BGP processes for the same **RARE/freeRtr** instance linked to a different VRF. in addition, **RARE/freeRtr** uses dual stack, so the entire feature set supports both IPv4 and IPv6 without compromise.

In order to support legacy applications, **RARE/freeRtr** brings along a densily populated control plane able to handle: i) diverse encapsulation; ii) cryptography; iii) both tablebased and also lable-based at L2 and L3 forwarding; iii) diverse tunneling; iv) both IGP and EGP routing protocols; and v) Policy-Based Routing services, NAT, QoS, and other services, highlighted by **Fig. 2**. A more detailed list of supported protocols can be found in the RARE project website¹.

3. PolKA lifecycle experience

This section describes the PolKA lifecycle experience from its proposal, prototyping, validation, deployment, and integration in production-grade testbeds.

3.1. Initial proposal: Prototyping PolKA in an emulated environment

Previous works explored the integration of RNS-based SR [Martinello et al. 2014, Liberato et al. 2018]. However, these works relied on integer RNS, and the integer modulo operation cannot be implemented in current commodity network hardware. Therefore, they either used software switches implementations [Martinello et al. 2014], or depended on synthesizing integer division to ASICs or NetFPGAs [Liberato et al. 2018].

To solve this problem, PolKA was introduced [Dominicini et al. 2020], as a RNS-based SR scheme that replaces the integer arithmetic by polynomial arithmetic (Galois field (GF) of order 2 or GF(2)). The immediate benefit was to enable the reuse of commodity network functions based on polynomial arithmetic. In fact, PolKA explores the Cyclic Redundancy Check (CRC) hardware to execute the *modulo*, and evaluated it in an emulated environment with Mininet and P4-16 language. A hardware prototype with Netronome SmartNICs was used for measuring forwarding latency, but, as these SmartNICs only support fixed CRC polynomials, it restricts the use of its CRC hardware for PolKA.

3.2. PolKA as use case of the RARE testbed: Prototyping for Tofino targets

Then, PolKA was selected as an use case in RARE/GÉANT P4 Lab². The partnership enabled an implementation of PolKA with the high-performance switching ASIC Tofino, and the first hardware-based comparison of PolKA with traditional approaches. As a result, we deployed a PoC [Dominicini et al. 2021], which comprised four Intel/Barefoot Tofino WEDGE100BF32X switches that were geographically distributed in Europe and connected at 10Gbps links. The experimental results showed that PolKA matches the data plane performance of traditional L2 table-based forwarding and list-based source routing.

In these experiments, the focus was on the data plane evaluation, and the control plane was implemented as a set of scripts that automated the essential node configuration tasks. In addition, our P4-based data plane was individually deployed in each switch.

3.3. Integration of PolKA in RARE/freeRtr

The next step was the integration of **PolKA** in the RARE project. Besides being the first non-standard protocol to be available for any experimenter in RARE, we inherited various

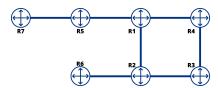
¹http://rare.freertr.org

²https://wiki.geant.org/pages/viewpage.action?pageId=148085131

features, such as diverse encapsulations, access control lists, and policy-based routing. Moreover, **RARE/freeRtr** allows the emulation of any topology with the same configuration files that can be tested and, then, deployed in a real P4-enabled infrastructure [Loui et al. 2022]. In addition, **PolKA** can now be evaluated with baseline approaches, like Segment Routing, which are also supported by **RARE/freeRtr**.

The control plane in the RARE project is based on **RARE/freeRtr**. Thus, to develop PolKA's control plane we have to options: (i) to develop a centralized controller that interacts with the **RARE/freeRtr** API, or (ii) to reuse standard distributed protocols already running in RARE and **RARE/freeRtr**. As a first integration effort, we chose the second option, by integrating the PolKA control plane logic to the **RARE/freeRtr** workflow³ to: (i) retrieve available topology information from link-state routing protocols to calculate the routeIDs when creating **PolKA** tunnels, and (ii) reuse the Segment Routing indexes of the nodes to create a static table that maps these indexes to nodeID polynomials.

The DPDK framework and the P4 behavioural language are used to describe the packet processing behaviour of RARE data plane. The supported P4 target platforms are the BMv2 software switch with the V1Model architecture, and the Tofino high performance network processor with the PSA (Portable Switch Architecture) architecture. For the integration of **PolKA** data plane, we ported our previous code for both the DPDK and the Tofino data planes⁴. In this implementation, our fixed-length PolKA header can use different encapsulations (e.g., MPLS, Ethernet, VLAN, and PPP), and the routeID has 128 bits with CRC 16, which allows the maximum of 8 hops.



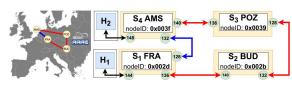


Figure 3. Edge-Core Experiment

Figure 4. RARE/GEANT testbed

4. Deployment and validation of PolKA with RARE/freeRtr

As proof of concept, it was implemented a ring topology: R1, R2, R3 and R4 are core nodes, while R5 and R6 are edge nodes, and R7 is the source, as shown in the **Fig. 3**. This topology is composed by a shortest and a longest path chosen to be equivalent to the topology of RARE/Géant P4 Lab testbed (**Fig. 4**).

4.1. PolKA Experiment in RARE/freeRtr

Retrieving topology information and configuring nodeIDs: For these functionalities, our design choice was to reuse legacy protocols rather than re-implementing these features in a centralized controller. In particular, for this experiment, the OSPF routing protocol was enabled on all ethernet interfaces to gather the topology information. Besides, we map the already assigned unique Segment Routing indexes of each router to an index of a static table of irreducible polynomials to set a unique **PolKA**'s nodeID for each router.

Creating PolKA Tunnels: Two tunnels were created, one to the shortest path (2 hops) **Fig. 5** and a second one to the longest path (4 hops) with an example of settings in

```
<sup>3</sup>https://docs.freertr.org/guides/reference/
```

```
<sup>4</sup>https://bitbucket.software.geant.org/projects/RARE/repos/rare/
browse/p4src/include
```

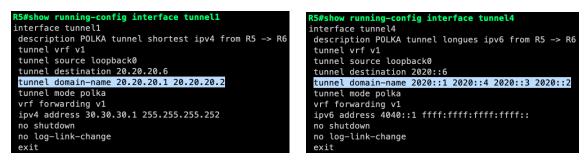


Figure 5. IPv4 Tunnel settings

Figure 6. IPv6 Tunnel settings

IPv6 Fig. 6. An important parameter is tunnel domain-name, that defines the path, containing the address of all the routers, and enables the *routeID* creation at the edge.

Visualizing tunnels in the edges: PolKA routeID is automatically computed when creating PolKA tunnels. **Fig. 7** and **Fig. 8** allow us to observe the configurations on the edge router R5. This tunnel inserts the routeID in the **PolKA**'s packet header. Then each router along the path extracts the routeID to forward the packet to the next hop, by using a modulo operation with routeID and its nodeID.

R5#sh	ow polka ro	uteid t	unnel1					
mode	routeid							
hex	00 00 00 0	0 00 00	00 00	00 00 26	d9	4d b	0 6b	6b
poly	1001101101	1001010	0110110	110000001	1010	1101	1010	11
index	coeff	poly	crc	equal				
0	00010000	27499	27499	true				
1	00010001	2	2	true				
2	00010003	6	6	true				
3	00010005	23389	23389	true				
4	00010009	56209	56209	true				
5	0001000f	33006	33006	true				
6	00010011	0	0	true				
7	0001001b	55909	55909	true				
8	0001001d	33248	33248	true				
9	0001002b	8069	8069	true				

		_								
	ow polka ro	uteid t	unnel4							
mode	routeid		_							
hex	00 00 00 0	0 00 00	69 76	5d c5 d7 be 75 93 96	9a					
poly	1101001011	1011001	.0111013	1100010111010111101111	10					
01110101100100111001011010011010										
index	coeff	poly	crc	equal						
0	00010000	38554	38554	true						
1	00010001	4	4	true						
2	00010003	6	6	true						
3	00010005	2	2	true						
4	00010009	3	3	true						
5	0001000f	3401	3401	true						
6	00010011	0	0	true						
7	0001001b	25646	25646	true						
8	0001001d	21719	21719	true						
9	0001002b	64376	64376	true						

Figure 7. RouteID on Tunnel 1

Figure 8. RouteID on Tunnel 4

Repository and dissector: We provide a github repository⁵ to allow the reproducibility of the **PolKA** experiments on **RARE/freeRtr**, including installation, testing, and debugging instructions. In addition, a Wireshark dissector⁶ allows the analysis of the PolKA header.

4.2. Agile Path Reconfiguration with Policy Based Routing

In order to demonstrate an agile path reconfiguration, we make use of a policy based routing (PBR) for traffic classification at the edges, according to **Figs. 9 and 10**. So, firstly, the traffic is classified through an access list, and then it is forwarded through one of the available tunnels. This allows us to quickly steer and balance the traffic through the **PolKA** tunnels, avoiding congestion due to an unequal use of the network.

In **Fig. 9**, we show a traffic in green sent between edge router R5, passing through core routers R1 and R2, to reach at edge router R6 over **PolKA** tunnel 1. **Fig. 10** shows a **PolKA** tunnel in IPv6 to transport traffic over the longest path crossing the core routers R1, R4, R3, and R2 until reaching the edge router R6.

5. Conclusion

This work describes the PolKA architecture and protocol implementation within the freeRtr open-source router platform. We carried out experiments deploying PolKA in

⁵https://github.com/eversonscherrer/wpeif2022

⁶https://github.com/eversonscherrer/dissector-polka

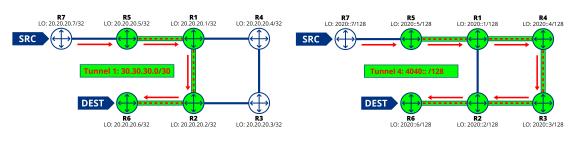


Figure 9. Shortest Path

Figure 10. Longest Path

a topology composed of edge and core nodes. The required topological information for PolKA was enabled by reusing the OSPF legacy protocol. For validation purposes, tunnels were set both with IPv4 and IPv6 to reach distinct destinations through different paths demonstrating that the connectivity service is functioning correctly. As lessons learned, the lifecycle experience from prototyping to deployment is a loop-intensive development process. The right balance between implementing everything from scratch instead of reusing legacy protocols is key in this process. Although open-source router platforms such as freeRtr bring along legacy protocols, it still needs to improve design modularity and minimal documentation set for developers willing to integrate and test their protocols.

6. Acknowledgments

This work received the 2021 Google Research Scholar Award, and the 2022 Intel Connectivity Research Grant (Fast Forward Initiative). It is also supported in part by the following Brazilian agencies: CNPq, RNP, CAPES (Finance Code 001), FAPESP/M-CTI/CGI.br (PORVIR-5G #20/05182-3, SAWI #20/05174-0, and SFI, #18/23097-3), FAPES (#94/2017, #281/2019, #515/2021, and #284/2021). CNPq fellows Dr. Martinello #306225/2020-4 and Dr. Ribeiro #315463/2020-1.

References

- Bajard, J. C. (2007). A residue approach of the finite fields arithmetics. In 2007 Conference Record of the Forty-First Asilomar Conference on Signals, Systems and Computers, pages 358–362.
- Dominicini, C. et al. (2020). Polka: Polynomial key-based architecture for source routing in network fabrics. In 2020 6th IEEE Conference on Network Softwarization (NetSoft), pages 326–334. IEEE.
- Dominicini, C. et al. (2021). Deploying polka source routing in p4 switches. In 2021 *International Conference on Optical Network Design and Modeling (ONDM)*, pages 1–3. IEEE.
- Liberato, A. et al. (2018). RDNA: Residue-Defined Networking Architecture Enabling Ultra-Reliable Low-Latency Datacenters. *IEEE TNSM*.
- Loui, F., Mate, C., Gall, A., and Wisslé, M. (2022). Project overview: Router for academia research & education (rare).
- Martinello, M. et al. (2014). KeyFlow: a prototype for evolving SDN toward core network fabrics. *IEEE Network*, 28(2):12–19.
- Shoup, V. (2009). A computational introduction to number theory and algebra. Cambridge university press.