PART IV

Research and Experimentation PROJECTS RECENTLY Funded

19

Recursive InterNetwork Architecture (ARCFIRE, Large-scale RINA benchmark on FIRE)

Sven van der Meer[1], Eduard Grasa[2], Leonardo Bergesio[2], Miquel Tarzan[2], Diego Lopez[3], Dimitri Staessens[4], Sander Vrijders[4], Vincenzo Maffione[5] and John Day[6]

[1]LMI
[2]Fundacio 12CAT, Spain
[3]TID
[4]University of Ghent, Belgium
[5]Nextworks, Italy
[6]Boston University, USA

19.1 Introduction

The main goal of ARCFIRE is to bring RINA from the labs into the real-world. RINA, the Recursive InterNetwork Architecture, is an innovative "back-to-basics" network architecture that solves current limitations and facilitates full integration between distributed computing and networking. RINA addresses the challenges that drive the communications industry in moving from dedicated hardware to almost completely virtualised infrastructure.

New technologies such as 5G will change the communication industry even more significantly before 2020. Here, ARCFIRE contributes by providing experimental evidence of RINA's benefits, at large scale, in compelling and realistic business cases. This will motivate RINA adoption. ARCFIRE demonstrates – experimentally – RINA's key benefits integrating current European investment in advanced networks (IRATI, PRISTINE) and Future Internet testbeds (FIRE+) focusing on five goals:

1. Facilitate comparison of converged operator networks using RINA against operators' current network designs;
2. Produce a robust RINA software suite ready for Europe to engage in large-scale deployments and long-living experiments;
3. Provide relevant experimental evidence of RINA benefits to network operators, their equipment vendors, application developers and end-users;
4. Build on the current EU Future Internet community and raise the number of organisations involved in RINA development and deployment;
5. Enhance the FIRE+ infrastructure with ready to use RINA software.

ARCFIRE will have long-term sustainable impact on how we build infrastructure for the Networked Society. ARCFIRE's deployed software suite will enable equipment vendors to shorten their innovation life cycle, network operators to run advanced networks addressing their needs in a future-proof fashion, European SME's to find and exploit specialised markets, and application developers to explore unseen opportunities.

19.2 Problem Statement

The leitmotiv of ARCFIRE is to *experimentally demonstrate at large scale they key benefits of RINA, leveraging former EC investments in Future Internet testbeds (FIRE+) and in the development of the basic RINA technology (IRATI, PRISTINE).*
ARCFIRE's contribution is

1. showcase the benefits and viability of RINA via large-scale experimental deployments;
2. quantify those benefits by comparing RINA with current Internet technologies using different Key Performance Indicators (KPIs) and
3. motivate the academic and industrial computer networking research communities to engage in RINA research, development and innovation activities.

ARCFIRE addresses the following specific objectives:

1. *Compare the design of converged operator networks using RINA to state-of-the art operator network designs* – ARCFIRE analyses the design of current state of the art converged operator networks, carry out an equivalent design using RINA and compare both approaches using a set of KPIs.

2. *Produce a robust RINA software suite; mature enough for large-scale deployments and long-lived experiments* – IRATI, the most ambitious RINA implementation to date, is today mature enough to support short-lived experiments that allow only minor traffic variations in the range of a few hours (2–3) with a relatively small number of systems (up to 20), supporting only a couple of DIF levels. ARCFIRE improves the open source IRATI software suite so that it is possible to make large-scale experimental deployments with up to 100 nodes (physical or virtual), supporting tens to hundreds of DIFs, up to 5 levels deep, running experiments for up to a week. These metrics will allow the IRATI implementation to be used both for rich experimental research activities and for internal trial deployments by network operators.
3. *Provide relevant experimental evidence of RINA benefits for network operators, application developers and end-users* – ARCFIRE, via WP4, performs 4 extensive experiments with the goal of experimentally evaluating different aspects of converged RINA operator networks:
   - T4.2 looks at the benefits of RINA when managing multiple layers over multiple access technologies;
   - T4.3 assess how RINA improves the operation of resilient, virtualised services over heterogeneous physical media;
   - T4.4 analyses end-to-end service provisioning across multiple RINA network providers and
   - T4.5 studies the effectiveness of RINA against Distributed Denial of Service Attacks (DDoSs).

   All experiments will target large-scale deployments and run for relatively long periods of time (as defined in Objective 2).
4. *Raise the number of organisations involved in RINA research, development and innovation activities* – ARCFIRE implements a set of actions in order to raise the acceptance of RINA by the computer networking research community. These actions, refined as part of T5.1 activities, are designed to overcome two of the main reasons for the current low number of researchers involved in RINA R&D: facilitate the understanding of the RINA concepts and mechanics and disseminate experimental results that prove the benefits of RINA in high-impact scientific publications and conferences.
5. *Enhance FIRE+ as a platform for large-scale experimentation with RINA* – Facilitate experiments with the IRATI RINA implementation on the FIRE+ facilities by documenting all the experiments carried out

by the consortium using the FIRE+ infrastructure, publishing of all the configurations (and Virtual Machine) templates used in those experiments and adapting or extending the generic FIRE+ experiment provisioning, control and monitoring tools. ARCFIRE also provides feedback on these tools with respect to join FIRE+-GENI experiments.

19.3 Background and State of the Art

RINA, the Recursive InterNetwork Architecture is a "back to basics" approach learning from the experience with TCP/IP and other networking technologies in the past. To better understand the implications that this approach has uncovered and to explore its consequences, it is necessary to build systems that adhere to that architecture. Research results to date have found that many long-standing network problems are inherently solved by the structure of the resulting RINA theory of networking. Hence, additional mechanisms or, more commonly, the series of hacks and patches found with current technologies are not required.

Our back-to-basics approach reminded us that from its inception, networking was viewed as InterProcess Communication (IPC). Hence, RINA starts from the premise that networking is IPC and only IPC. Networking provides the means by which application processes on separate systems communicate, generalising the model of local IPC. Figure 19.1 shows a diagram of the RINA architectural model. In contrast to the fixed, five-layer model of the Internet, where each layer provides a different function, RINA is based on a single type of layer, which is repeated as many times as required by the network designer. The layer is called a Distributed IPC Facility (DIF), which is a distributed application that provides IPC services over a given scope[1] to the distributed applications above (which can be other DIFs or regular applications. These IPC services are defined by the DIF API, which provides operations to

1. allocate flows to other applications by specifying an application name and a set of characteristics for the flow such as delay, loss, capacity;
2. read and write data from the flows; and
3. de-allocate the flows and free the resources associated to them.

[1]Scope is the locus of distributed shared state that forms a layer, examples of these layers with different scopes may occur in point-to-point links, networks, networks of networks, virtual private networks, etc.

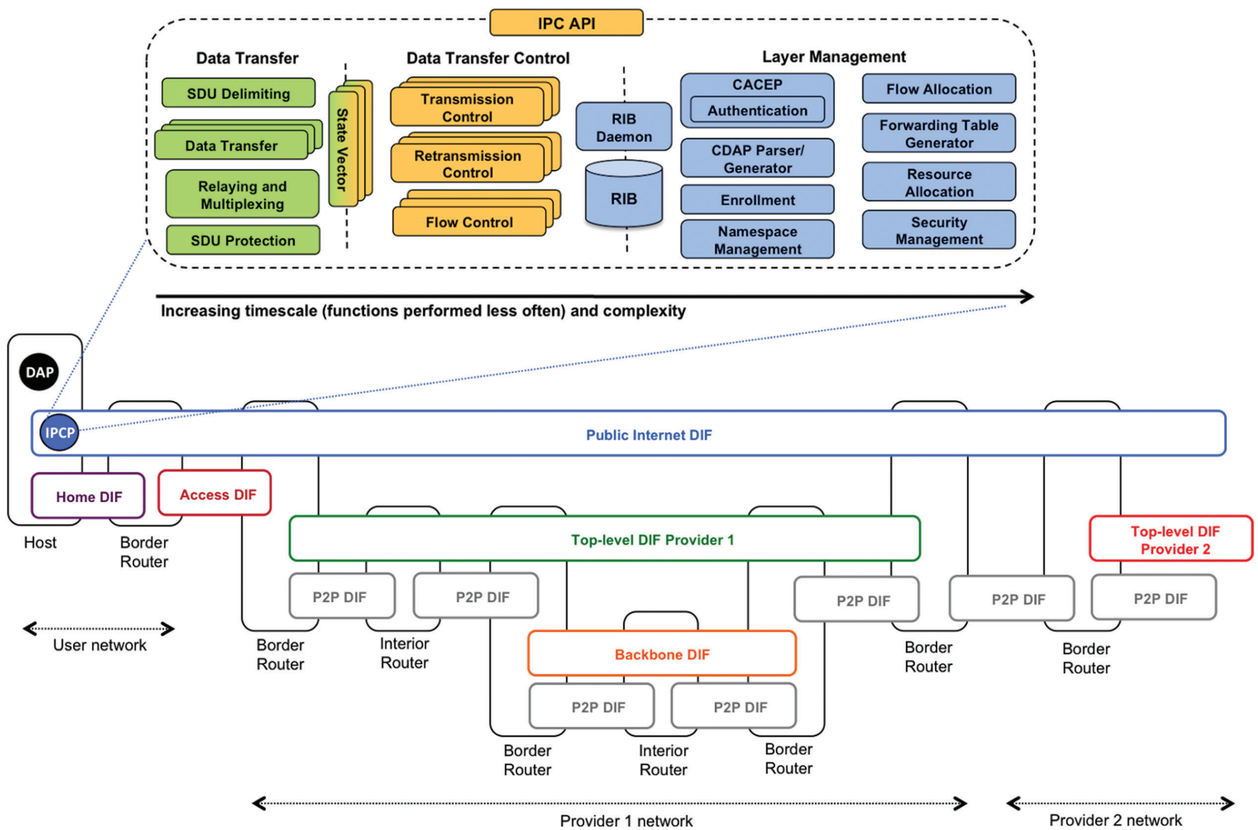

Figure 19.1 An example of the RINA architecture, with the same type of layer (DIF) repeated as required over different scopes. Different sets of policies customise each layer to its operational range.

All DIFs offer the same services through their API and have the same components and structure. However, not all the DIFs operate over the same scope and environment nor do they have to provide the same level of service. In RINA, invariant parts (mechanisms) and variant parts (policies) are separated in different components of the architecture. This makes it possible to customise the behaviour of a DIF to optimally operate in a certain environment with sets of policies for that environment instead of the traditional "one size fits all" approach or having to re-implement mechanisms over and over again.

The principles behind RINA, were first presented by John Day in his book "Patterns in Network Architecture: A return to Fundamentals". Since the book was published in 2008, several organisations have stated their interest in further researching RINA, as well as into turning the theory into practice by deploying RINA in the real world. The http://pouzinsociety.orgPouzin Society (PSOC) was formed in 2009 to coordinate all the international activities around RINA research and development. Three research initiatives have been previously funded by the European Commission:

- the FP7 IRATI project, which succeeded in developing the first RINA implementation over Ethernet and showing the benefits of the RINA structure;
- FP7 PRISTINE, which is building on IRATI's results to further improve this nascent technology and start demonstrating the benefits of RINA in specific areas such as congestion control, resource allocation, routing, security or network management; and
- the G3+ open call winner IRINA, which researched RINA as a potential alternative for the future architecture of GEANT and National Research and Education Networks (NRENs).

Both IRATI and PRISTINE have been/are very development-intensive projects, with a strong focus on implementation activities, simulation and low-scale experimentation. IRINA was focused on studying the benefits of RINA for NRENs and performing a small lab experiment with the IRATI stack. The results obtained by ARCFIRE will provide a definitive answer to the question of why should the different computer networking community stakeholders (academia, industry, funding bodies, etc) invest on RINA.

19.4 Approach

The ARCFIRE technical approach is based on three research and development activities, each of them with their specific inputs and outputs:

1. **Design of converged operator networks with RINA – WP2**
   This is the use case scenario of the project. ARCFIRE has chosen this scenario because it is the one that can best illustrate the benefits of RINA and allow for a greater diversity of experiments. ARCFIRE will take as an input the latest release of the RINA specification from the Pouzin Society, as well as the current state of the art in the design of converged networks (mainly contributed by ERICSSON and TID). The outputs of this design activity will be a set of new and enhanced RINA specifications (contributed back to the Pouzin Society). The specifications form the theoretical framework that will allow the design of the experiments and selection of KPIs and the set of new features in the RINA software suite that will need to be implemented by ARCFIRE.
2. **RINA software suite adaptation, enhancement and robustification – WP3**
   ARCFIRE will build on the results of FP7 IRATI, FP7 IRINA and FP7 PRISTINE and enhance the open source IRATI RINA software suite in order to make it compliant with the ARCFIRE requirements. This work will involve the design and development of new features required by the experimentation activities, but a good share of ARCFIRE's development activities will go into maturing the RINA stack, making it more stable and automated to enable large-scale deployments and long-running experiments. The resulting RINA software suite will be contributed back to the IRATI open source project, to make it available to the individuals and organisations interested in experimenting with RINA.
3. **Large-scale experimentation with the IRATI RINA implementation on FIRE+ facilities – WP4**
   The core activities of the project will exploit the large catalogue of Future Internet experimental facilities available in FIRE+, as well as GENI at the US. ARCFIRE will look at some of the key aspects in the operation of RINA-based converged operator networks to setup experiments that analyse and quantify some of the key benefits of RINA. These aspects include the management of multi-layer networks, provisioning of reliable services over heterogeneous physical media, end-to-end service provisioning across multiple network operators and effectiveness against DDoS attacks. Experimental activities will produce the key result of ARCFIRE: experimentally verification and quantification of RINA benefits over the current state of the art Internet architecture. A secondary but important output, which will be contributed back to FIRE+, are procedures and tools to facilitate large-scale RINA experimentation on FIRE+ experimental facilities.

19.5 Technical Work

Currently, operator networks suffer from a set of limitations and inefficiencies due to the design errors in the current Internet protocol suite and the series of patches that have been introduced to solve the problems caused by these errors. The greatest problem for the design and operation of service provider networks today is complexity. Today operators have to deploy separate, independently managed networks to support the public Internet and applications with strong quality guarantees, such as separate networks to carry voice to allow the usage of IP as a replacement of the SS7/TDM system in the Public Switch Telephone Network. The lack of a well-defined security architecture also fills the network with all sorts of equipment such as Network Address Translation boxes (NATs) or firewalls. The limited number of layers in the TCP/IP architecture, tied to the rigidness of their functions requires the introduction of pseudo-layers such as MPLS, VLANs, Q-in-Q, MAC-in-MAC, L2 and L3 tunnels and lately different network overlay/virtualisation technologies such as VXLAN or NGVRE; all of those requiring dedicated equipment and/or software to provide the new functionality. The poor support for mobility and multi-homing in IP networks requires a completely separate architecture for the mobile access network, with their own set of standards, protocols and equipment (GSM, UMTS or LTE as examples). Moreover, the lack of flexibility and adaptability in the Internet's transport layer makes it hard for operator networks to provide an optimal performance over heterogeneous physical media.

In contrast, the simple structure exhibited by RINA can be leveraged for designing simpler, more performant and predictable operator networks. Figure 19.2 shows a simplified example of an operator network (center), connected to a customer network (left) and to another operator network (right). If we focus at the border between the customer network and the operator 1 network, we can see that the customer gets access to one or more DIFs via the operator's top level DIF (called Operator 1 Metropolitan DIF in the Figure). These DIFs may be general-purpose DIFs with a large number of users (such as a public Internet DIF), or community or event application-specific DIFs with more specialised policies and tighter security. If the customer wishes to do so these set of DIFs can span to the customer's network (floating on the customer's internal layers with its own private addressing not visible from outside of the customer's network) and be exposed to individual applications running at the customer's premises: for instance, a "banking DIF" could be made available to a client application of an online banking tool, but hidden from other applications such as the web browser or the email client.

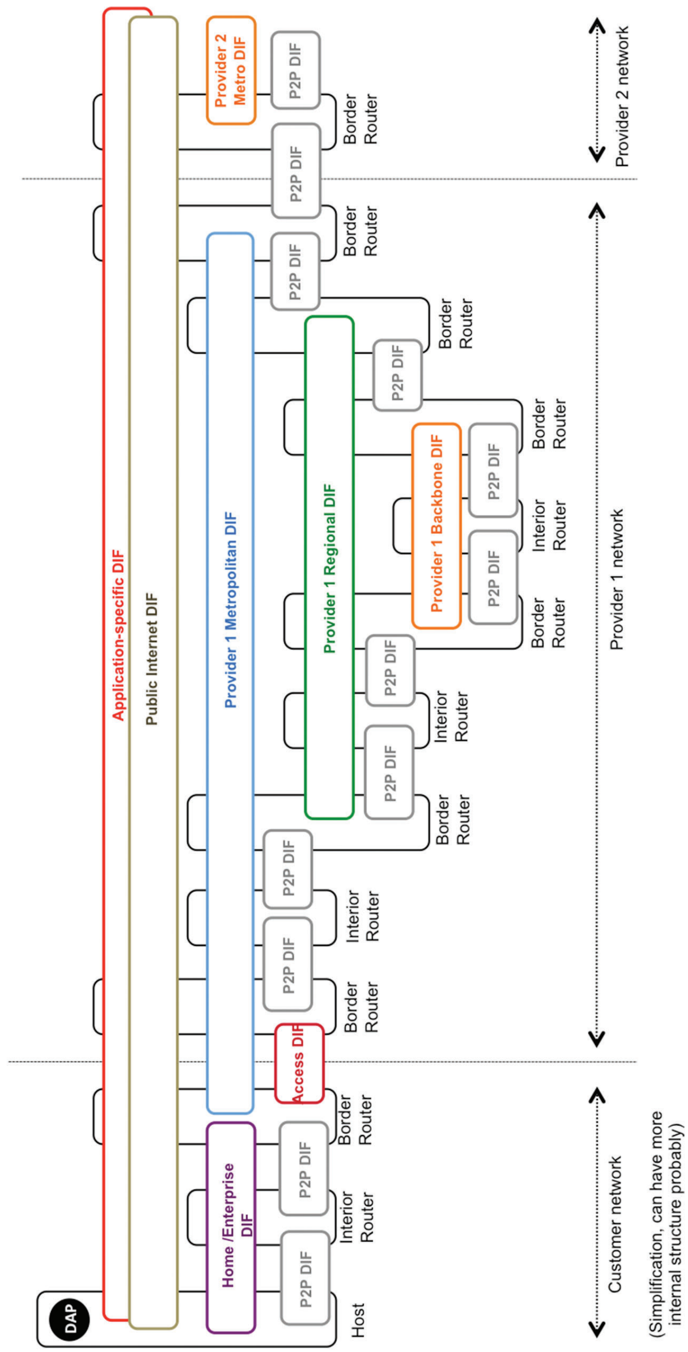

Figure 19.2 An example of a RINA operator network.

A crucial difference with the operator networks of today is the fact that RINA provider network only has two kinds of systems: interior routers and border routers (which is where recursion happens, going a layer up or down as shown in Figure 19.2). There is no need for middleboxes such as NATs, firewalls, tunnel termination devices: all the required functions are contained in interior routers or in border routers. The simplification potential in terms of management and operations is huge: networks built this way would be much simpler and cheaper to design, build and operate, which would make them easier to automate, more secure and predictable.

The designer of a RINA network can use as many layers (DIFs) as necessary to accomplish the design goals. The network designer can bound the size of the DIF, and just break a DIF into smaller DIFs if the first one becomes unwieldy. With RINA network designers have a structural tool to scale the network up: the DIF. There is no need to support layers than have to grow indefinitely such as in the current Internet. Also, interactions between DIFs are much more predictable than interactions between layers in the Internet, since all DIFs have the same interface and follow the same architectural patterns, albeit with different policies.

With regard to existing RINA software, ARCFIRE will use the IRATI open source RINA stack implementation, the Network Manager developed by the PRISTINE project and the open source RINA traffic generator (rina-tgen), initially developed by the IRINA project and now also part of the open source IRATI initiative. All these software components will be improved during the project lifetime. WP3 will be the work package responsible for adapting, improving and maintaining the RINA software suite for the purposes of the ARCFIRE project.

With regard to large scale experimentations, ARCFIRE will not focus only on a single FIRE+ facility; it will choose the most appropriate facility for the requirements of each individual experiment. This approach is today a realistic option due to the work the FIRE+ community has devoted to developing common tools for accessing the FIRE+ experimental facilities, deploying and controlling experiments and obtaining the experiments' data. By using this federated approach researchers don't have to master different toolsets when changing from one facility to another. FED4FIRE is leading the development of the FIRE+ federation of facilities, with 21 individual testbeds involved as of today, many of which offer open access programs to experimenters. ARCFIRE's preferred method to access FIRE+ facility will be via the FE4FIRE federation of testbeds. However, if a facility that is not member of FED4FIRE provides an interesting environment for the ARCFIRE experiments, the project will also consider its use.

19.6 Conclusion

ARCFIRE is guided by excellence in science and industrial leadership in integrated infrastructure with smart networking, focused on experimentally validating and benchmarking a breakthrough technology (RINA), thus bringing it closer to the market. In line with the H2020 objectives and the EU industrial policy goals, the impacts of ARCFIRE's results will contribute to the increase of European competitiveness and thought leadership in the area of networking and distributed computing, helping to create jobs and supporting growth. SMEs make up some $\frac{2}{3}$ of European industry's employment and a large share of EU industry's growth and jobs potential is to be found in its lively and dynamic SMEs. As such, the involvement of an SME in ARCFIRE NEXTWORKS – complemented by the participation of two industry players like Ericsson (the project coordinator) and Telefonica I+D is crucial for maximising the expected impact of ARCFIRE's actions. Last but not least, corporations Interoute – and SMEs in the External Advisory Board – TRIA Network Systems, Martin Geddes Consulting will provide a secondary exploitation path to ARCFIRE's outcomes.