FLAMINGO

European Seventh Framework Network of Excellence

http://www.fp7-flamingo.eu/

WP6 — Automated Configuration and Repair

Deliverable D6.3 — Third Year Report on Automated Configuration and Repair

University of Twente, The Netherlands (UT)
Institut National de Recherche en Informatique et Automatique, France (INRIA)
University of Zurich, Switzerland (UZH)
Jacobs University Bremen, Germany (JUB)
Universität der Bundeswehr München, Germany (UniBwM)
University Politecnicà de Catalunia, Spain (UPC)
iMinds, Belgium (iMinds)
University College London, United Kingdom (UCL)

Project funded by the European Union under the
Information and Communication Technologies FP7 Cooperation Programme
Grant Agreement number ICT-FP7 318488

Document Control

Title:	D6.3 — Third Year Report on Automated Configuration and Repair
Type:	Public
Editor(s):	Gabi Dreo Rodosek
E-mail:	gabi.dreo@unibw.de
Doc ID:	D6.3
Delivery Date:	31.10.2015
Author(s):	Anthéa Mayzaud, Anuj Sehgal, Gaëtan Hurel, Gabi Dreo, Christos Tsiaras, Anna Sperotto, Daniel Dönni, Daphne Tuncer, Marinos Charalambides, Mario Flores, Jeroen Famaey, Mario Golling, Maxim Claeys, Niels Bouten, Nikolay Melnikov, Radhika Garg, Rashid Mijumbi, Ricardo Schmidt, Frank Tietze, Rick Hofstede, Sebastian Seeber, Steven Latré, Corinna Schmitt, Abdelkader Lahmadi, Jair Santanna, Stefano Petrangeli, Peter Hillmann, Bram Naudts, Javier Rubio-Loyola, Sofie Verbrugge

For more information, please contact:

Dr. Aiko Pras
Design and Analysis of Communication Systems
University of Twente
P.O. BOX 217
7500 AE Enschede
The Netherlands
Phone: +31-53-4893778
Fax: +31-53-4894524
E-mail: <a.pras@utwente.nl>

Legal Notices

Executive Summary

Emerging trends, such as Future Internet, Internet of Things or Internet of Everything pose new challenges to the network and service management. Due to the increasing complexity, management needs to be part of the functionality of the managed objects instead of a thought after. Therefore, it is necessary to think of a new management paradigm, namely Management-by-Design. Taking into account the Future Internet Management requirements, automation of tasks and processes is inevitable. This gets amplified through the increasing amount of networked devices with a nearly unimaginable range of capabilities. The management goals need to cover a full range of technologies (starting from simple dumb sensors to smart devices) and various communication infrastructure (from isolated networks connected to the cloud to inter-clouds). WP6 addresses these aspects of automated configuration and repair.

D6.3 describes the achievements reached in WP6 during the third year of FLAMINGO. The main focus of Y3 were further steps towards the FLAMINGO integrative architecture for automated configuration and repair. Due to the two main pillars in the use cases security and content the FLAMINGO integrative architecture for automated configuration and repair is based on strong and well proven approaches. With respect to the S.M.A.R.T. (Specific, Measurable, Achievable, Relevant, Timely) objectives (Section B.1.1.5 of the Description of Work) we claim the full achievement of our third year targets for WP6.

Due to the tight collaboration of this work package with other work packages like WP1 and WP5, that has even been more intense in the third year, several PhD students are directly contributing to WP6 specific objectives. As a scientific output, we can report a total of 73 papers, where 40 are strongly related to WP6, that have been already published in Y3, and several other papers that are currently under review. This scientific output is exceeding the objective initially set in the DoW of 20 papers. We point to Deliverable D8.3 for details. The work package specific objectives center around the following three tasks, namely (i) to develop innovative architectural approaches for automated configuration and repair (Task 6.1), (ii) to identify enablers for these new architectures (Task 6.2) and (iii) to analyze the applicability of the developed approaches to selected application domains (Task 6.3). Key achievements of WP6 in the third year, as specified in the DoW and as documented in D6.3, are summarized below:

Task 6.1: Architectures Whereas in Y1 the field of attention was related to the development of an inventory of architectures and approaches in the area of automation, and in Y2 the focus was extended to the area of software defined networking (SDN), in Y3 we draw attention to the research work done in WP6. This architecture is built upon our strong pillars in the use cases of security and content delivery. Therefore, WP6 developed an approach for robust geometric forest routing with tunable load balancing to support inter-cloud capabilities. Addressing the security use case, WP6 evaluated multiple security event exchange mechanisms, e.g. IDMEF, Syslog, X-ARF, IODEF. In the area of content delivery an adaptive resource management and control framework was proposed by WP6 which showed significant gain in terms of link utilization and energy consumption.

Task 6.2: Enablers The inventory of enablers and approaches in the field of network and service management which was built in Y2 serves in Y3 primarily as a guideline for the development of new approaches and their applicability. Nonetheless, additional enablers are part of the developed solutions within Y3. In addition, WP6 investigated drawbacks of existing approaches in HTTP adaptive streaming (HAS). Based on these findings an improved approach using machine learning techniques to increase the fairness among HAS clients was developed. The success of the proposed solution was demonstrated by a comparison with existing

popular solutions (e.g. Q-Learning and FESTIVE) for HAS. The approach is based on our comprehensive analysis of enabling technologies done in Y2. Further results based on this enabler analysis have been achieved using the enabler data mining, clustering and semantic reasoning for developing an approach for security function chaining for android devices.

Task 6.3: Application Domains As reported in the past years, also in Y3 the approaches developed within WP6 are spread across the four application domains. For the area of wireless sensor networks a distributed monitoring architecture was developed to detect and mitigate anomalies in the Routing Protocol for Low power and lossy networks (RPL). The approach for outsourcing mobile security function to cloud services addresses an additional application domain. HAS as a new approach for delivering video content in a fair manner addresses the area of content-aware networks. Considering SDN technology, a framework to manage and control resources to support static and dynamic management applications was developed within WP6.

Thus, our expert knowledge gained in the first two years enabled us to implement promising approaches in Y3.

To summarize, we are convinced that all S.M.A.R.T as well as work package specific objectives in the third year have been fully achieved.

Contents

1 Introduction

WP6 is strongly devoted to the automation of management since this is a key precondition of the Future Internet. Deliverable D6.3. reports on the achievements of WP6 in the third year of FLAMINGO. Therefore, S.M.A.R.T(Specific, Measurable, Achievable, Relevant, Timely) objectives as well as WP6-specific objectives are addressed.

The first S.M.A.R.T. objective is the integration of PhD students. For a detailed list of the fully integrated PhD students, we refer to Deliverable D8.3. Many PhD collaborations within the consortium, which started during the first and second year are still ongoing. During the third year most of them started publishing their achievements. In addition, several new collaborations started in the third year. We refer to Section 3 for more details.

The second S.M.A.R.T. objective refers to the scientific output of the project. In Y3, the research work packages published in total 73 papers at major conferences and in journals, where 40 are strongly related to WP6. We report this summarized number in all research WP deliverables due to the tight research integration of WP5, WP6 and WP7 which is also manifested in the joint publications. To recall, Datasets and monitoring approaches, which are addressed in WP5 serve as input for automation approaches developed and implemented within WP6. In addition, WP6 is able to reconfigure the monitoring systems(WP5) to get more precise data sets. Furthermore, WP7 ensures that monitoring (WP5) and automated (WP6) actions are all performed within the boundaries of the economic, legal and regulative constraints. For a detailed list of the FLAMINGO published and submitted papers, we refer to Deliverable D8.3.

Deliverable D6.3 is structured as follows: Achievements related to S.M.A.R.T and WP-specific objectives are summarized in Section 2.1. Details about the progress in research related to our most relevant use cases security and content delivery with respect to the FLAMINGO automation architecture are described in Section 4.1.

Section 4.2 summarizes our selected highlights of the research conducted in WP6 during Y3. Since PhD collaborations form the basis of the research work done in FLAMINGO, Section 3 presents the PhD contributions and describes each ongoing collaboration and achievements in detail. Section 5 concludes the deliverable.

2 Objectives and Activities

This section presents an overview of the S.M.A.R.T. objectives for WP6. For each S.M.A.R.T. objective, we indicate how it was achieved in the reported year of the project. WP6-specific objectives summarize the activities that have taken place among the consortium members in Y3.

2.1 S.M.A.R.T. Objectives

To meet the S.M.A.R.T. objectives, WP6 has been active in the following aspects.

- **Integration of PhD students** – *The Description of Work (Section B.1.1.5) states that "after 9 months each research WP will have identified at least two fully integrated Ph.D. students, which means that these students will be jointly supervised and financially paid by FLAMINGO".*
  In the first two years of the project 14 PhD students have joined FLAMINGO. In the third year, three more PhD students have joined the NoE. These students, their affiliations and the co-supervising institutions are listed in D8.3. Since collaborations are a cornerstone of research within FLAMINGO. It is important that they are not only taking place between fully integrated PhD students, but also among students that are not financially paid by FLAMINGO but jointly supervised. Detailed information on the integration of PhD students can be found in Section 3.1.

- **Scientific Output** – *The Description of Work (Section B.1.1.5) states that "after 18 month at least 20 scientific papers will be submitted / published".*
  In the first two years the project had exceeded the expected number of publications. In the third year the research work packages published 73 papers at major conferences as well as in journals, and exceeds the expected number of papers. The strong collaboration between the WPs is based on the intensive PhD collaborations described in Section 3.1. There is a special intense collaboration between WP5 and WP6 since the monitoring data (WP5) builds the basis for numerous approaches addressed in WP6 and vice versa. WP6 adjusts the monitoring of WP5. This monitoring architecture developed by WP5 in Y1 is still maintained to allow WP6 delivering requirements for further data acquisition. A joint list of papers is included in Deliverable D8.3. which includes an assignment for each paper per work package. In Year 3, 40 papers are strongly related to WP6.

  Table 1 reports the result of collaborations with other European projects and institutions. Within the FLAMINGO consortium the published papers co-authored by more than one FLAMINGO member are listed in Table 2.

  Partners also targeted top conferences and journals in the network management field and high-end conferences and journals in the field of networking and measurements as suggested by the reviewers during the last evaluation. To address this, papers have been published at IEEE INFOCOM 2015, IEEE Communications and Network Security (CNS) 2015, IEEE International Symposium on Cyberspace Safety and Security (CSS 2015) and ACM Multimedia Conference (ACM MM)2015.

Table 1: FLAMINGO publications in Y3 in collaboration with other EU projects and institutions related to WP6.

Authors	**Title**	**Venue**	**EU project/ institution**
A. Lareida, T. Bocek, M. Pernebayev and B. Stiller	Automatic Network Configuration with Dynamic Churn Prediction	IFIP/IEEE International Symposium on Integrated Network Management (IM)	SmartenIT (317846)
C. Schmitt, M. Noack, W. Hu, T. Kothmayr and B. Stiller	Two-way Authentication for the Internet-of-Things	Book series on Advances in Information Sercurity, Privacy, and Ethics (AISPE) by IGI Globa	SmartenIT (317846)
C. Schmitt, and B. Stiller	Secure and Efficient Wireless Sensor Networks	ERCIM News - Special Issue: The Internet of Things and The Web of Things	SmartenIT (317846)
P. Porambage, A. Braeken, C. Schmitt, A. Gurtov, B. Ylianttila and B. Stiller	Group Key Establishment for Source Multicasting in IoT-enabled Wireless Sensor Networks	40th IEEE Conference on Local Computer Networks (LCN 2015)	European Celtic-Plus Project CONVINcE (C2013/2-1)
R. Houthooft, S. Sahhaf, W. Tavernier, F. De Turck, D. Colle and M. Pickavet	Robust Geometric Forest Routing with Tunable Load Balancing	IEEE Conference on Computer Communications (INFOCOM 2015)	European EULER project (258307) part of Future Internet Research and Experimenta-tion (FIRE)

Table 2: Publications in Y3 authored by multiple FLAMINGO partners related to WP6.

Authors	**Title**	**Venue**	**FLAMINGO partners**
A. Mayzaud, A. Sehgal, R. Badonnel, I. Chrisment, and J. Schönwälder	Mitigation of Topological Inconsistency Attacks in RPL based Low Power Lossy Networks	International Journal of Network Management	INRIA, JUB
S. Latré, M. Charalambides, J. François, C. Schmitt and B. Stiller	Intelligent Mechanisms for Network Configuration and Security	International Conference on Autonomous Infrastructure, Management, and Security (AIMS 2015)	iMinds, UCL, INRIA, UZH
R. Mijumbi, J. Serrat, J. Rubio-Loyola, N. Bouten, F. De Turck and S. Latré	Dynamic Resource Management in SDN-based Virtualized Networks	Network and Service Management (CNSM)	UPC,iMinds
N. Bouten, R. de O. Schmidt, J. Famaey, S. Latré, A. Pras and F. De Turck	Qoe-driven in-network Optimization for Adaptive Video Streaming based on Packet Sampling Measurements	Computer networks	iMinds, UT
N. Bouten, J. Famaey, R. Mijumbi, B. Naudts, J. Serrat, S.+Latré and F. De Turck	Towards NFV-based Multimedia Delivery	IEEE International Symposium on Integrated Network Management(IM)	iMinds, UPC
R. Mijumbi, J. Serrat, J. Gorricho, N. Bouten, F .De Turck and S. Davy	Design and Evaluation of Algorithms for Mapping and Scheduling of Virtual Network Functions	IEEE Conference on Network Softwarization (NetSoft)	UPC, iMinds
J. Steinberger, A. Sperotto, M. Golling and H. Baier	How to Exchange Security Events? Overview and Evaluation of Formats and Protocols	IFIP/IEEE International Symposium on Integrated Network Management (IM)	UT,UniBwM

2.2 Work Package Specific Objectives

Inside FLAMINGO each work package has its own defined objectives. This section reports on the ongoing WP6-specific objectives and the achievements during the third year.

OBJECTIVE 1 - To integrate European research in the area of automated configuration and repair: In cooperation with WP3 and WP5, WP6 was involved in the organization of the first IEEE Conference on Network Softwarization (NetSoft 2015) [1]. WP6 has also actively participated in the presentations and discussions within the 2015 EuCNC Workshop on Network Function Virtualisation (NFV) and Programmable Networks [2] and contributed with a presentation from Daphne Tuncer (UCL) and Niels Bouten (iMinds). WP6 was involved during the organization of the 9th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2015) [3]. Marinos Charalambides (UCL) was co-chairing the main track. During AIMS 2015 a tutorial on how to deploy NFV experiments on the JFed testbed was presented by the collaboration **iMinds-UPC-NFV** [4]. This tutorial was also presented during the FED4FIRE-GENI research experiment summit (FGRE 2015) [5]. For more information about interaction between industry and academia we refer to D3.4. Furthermore, a IJNM Special Issue on "Advances in Management of Multimedia Services" was organized by Marinos Charalambides (UCL). WP6 also organized the IJNM special issue "Measure, Detect and Mitigate Challenges and Trends in Network Security", a collaboration between UT, UZH, UniBwM and CAIDA. In addition, UPC has also been (and is still) engaged in the organization of conferences and workshops, for example the upcoming Conference on Communications Networks Service Management (CNSM 2015), which will take place at the UPC. UniBwM is actively involved in attracting collaborations with national research projects (e.g. SVFUA and fit4sec). Furthermore, together with ENISA UniBwM is involved in developing a solution for secure access to sensitive health data in eHealth processes. The aim of the collaborating projects reaches from key-management and key-exchange technologies to the identification of trending topics.

OBJECTIVE 2 - To create and maintain articles within Wikipedia and other online systems in this area: The ongoing research generated valuable insights that have been contributed to Wikipedia. WP6 contributed to Wikipedia on the following pages QUALITY OF EXPERIENCE [6] and ADAPTIVE BITRATE STREAMING [7]. The Wikipedia Page Policy-based management [8] was updated with information on policy refinement, which is an important part of the policy management lifecycle. In addition, WP6 has contributed to and maintained the Wikipedia page on SOFTWARE DEFINED NETWORKING [9] with relevant security mechanisms that can be implemented using the SDN paradigm. Furthermore, WP6 added a Wikipedia article related to Schengen Routing [10]. The page regarding content delivery networks [11] has been updated with information concerning telco content delivery networks (CDN). In the area of Wireless Sensor Networks WP6 edited the Wikipedia page on RPL and an additional section in the 6LoWPAN page [12]. Detailed information about Wikipedia editing done within FLAMINGO can be found in Deliverable D2.3.

[1] http://sites.ieee.org/netsoft2015/
[2] http://www.eucnc.eu/?q=node/113
[3] http://www.aims-conference.org/2015/
[4] http://www.aims-conference.org/2015/labs.html#session2
[5] http://www.fed4fire.eu/fed4fire-geni-research-experiment-summit-fgre-2015/
[6] https://en.wikipedia.org/wiki/Quality_of_experience
[7] https://en.wikipedia.org/wiki/Adaptive_bitrate_streaming
[8] https://en.wikipedia.org/wiki/Policy-based_management
[9] https://en.wikipedia.org/wiki/Software-defined_networking
[10] https://en.wikipedia.org/wiki/Schengen_Routing
[11] https://en.wikipedia.org/wiki/Content_delivery_network
[12] https://en.wikipedia.org/wiki/6LoWPAN

OBJECTIVE 3 - To develop an inventory of approaches for automated configuration and repair: **iMinds-UPC-NFV** created a survey on Network Function Virtualization (NFV) where the relationships between NFV on the one hand and SDN and cloud computing on the other hand are discussed [1]. The state-of-the-art NFV architectures and algorithms are analyzed together with the relevant research projects, standardization efforts and commercial products. The promising research directions in this area are identified and highlighted. Furthermore, **UT** contributes to this objective with an overview of formats and protocols to exchange security events, which are a basis for automated configuration and repair in the area of automated cyber defence [2] in a multi provider environment.

OBJECTIVE 4 - To specify guidelines about the applicability of approaches for automated configuration and repair to specific application domains: The collaboration **iMinds-UPC-NFV** has investigated the applicability of Network Function Virtualization for the delivery of multimedia services. To this end, a model was built taking into account the costs and benefits involved in such a scenario which gained insights in how to optimally distribute the underlying physical resources for the delivery of Multimedia services [3].

The collaboration **UCL-iMinds-Cache** has focused on the content aware routing application domain by demonstrating the benefits of applying hybrid cache management strategies in Telco-operated CDN scenarios. In addition, the collaboration also investigated the effect of the parallelization of the content placement decision-making process on the performance in terms of network and caching costs [4].

UCL has focused on the software defined networking (SDN) application domain by investigating the requirements of SDN based management and control framework to support static and dynamic resource management applications. UCL also developed algorithms to decide on the placement of distributed managers and controllers and derived guidelines for the best allocation of the relevant entities in the network [5].

OBJECTIVE 5 - To develop new architectures for automated configuration and repair approaches across administrative boundaries: The collaboration **UCL-iMinds-Cache** has extended previous work and developed new cache management approaches for a Video-on-Demand use case that can be used across administrative boundaries. This not only lead to an improvement in terms of network and caching performance, but can also support larger problem instances [4]. The collaboration between UPC and iMinds (**iMinds-UPC-NFV**) on network and function virtualization proposes new approaches for automated resource configuration in virtualized networks and functions. The results can be seen in the following publications: [6], [7], [8], [3], [9], [10]. Furthermore, the collaboration **INRIA-JUB-Distr** developed a distributed monitoring architecture which establishes the basis for automated configuration and repair actions in IoT environments. Combining existing IDS solutions and cloud based Intrusion Detection System (IDS) **UniBwM** developed an approach [11] to refine monitoring in order to increase the detection accuracy by a directly interfering packet forwarding.

OBJECTIVE 6 - To develop information models, algorithms, learning techniques and knowledge description approaches as enablers for automated configuration and repair: The collaboration **iMinds-UPC-NFV** formulated an online virtual function mapping and scheduling problem and proposed a set of algorithms for solving it. Three greedy algorithms and a tabu search-based algorithm were proposed and evaluated in terms of successful service mappings, total service processing times, revenue, cost, etc, under varying network conditions [7].

The collaboration **iMinds-UCL-Cache** developed a hybrid caching approach that combines proactive content placement and reactive cache replacement. Virtualization allows to separate the proactive and reactive caches. Periodically, the caching capacity is allocated and content is placed proactively based on history-based predictions of the future request pattern. Reactive caching capacity is uniformly split across the network to deal with unpredicted popularity fluctuations and errors in the request prediction.

In [12] UPC surveyed the importance of information modelling to the automated management, configuration, repair and orchestration of virtualized networks and functions.

iMinds developed new caching algorithms that are able to take into account the temporal relationships in video streaming as well as announcements on future requests. This allows to optimize the caching strategy in the presence of, for example, binge watching [13].

INRIA contributes to this Objective with a paper presenting the following: First, they put forward a mining methodology to extract mobile applications behaviors. To mine such behaviors, they used data mining, machine learning and different clustering algorithms such as K-means and Self-Organized Maps [14]. This is typically an enabler for automated configuration and repair since those mining results and extracted behaviors will be used to adequately choose and configure the security functions chains in this approach. Secondly, they extended the underlying mathematical model of this approach in order to model the security function compositions and characterize them with respect to several factors such as the used resources (CPU, battery) on the device-side when employing such compositions, and the induced network latency. The resulting model leverages graph theory in order to quantify the potential benefits and caveats of using the security compositions, and allows us to efficiently deploy those across the cloud architecture and device(s) to protect.

UCL developed different algorithms to control the placement of a set of distributed managers and controllers in a software-defined based management framework for fixed backbone networks [5].

OBJECTIVE 7 - To evaluate automated configuration and repair approaches as being part of the autonomic control loops: The collaboration **iMinds-UT-QoS** developed distributed autonomic control loops that monitor the current state of a dynamic network and use this information together with hierarchically distributed knowledge to efficiently divide the resources among various HAS clients [15].

iMinds has developed a fair HAS client able to achieve smooth video playback, while coordinating with other HAS clients in order to improve the fairness of the entire system. This goal is reached with the aid of a hierarchical in-network-based system of network nodes, in charge of collecting measurements on the network conditions. This information is then used by the HAS clients to refine their quality decision process and develop a fair behavior [16]. The collaboration **UCL-iMinds-Cache** has evaluated the new cache management strategies developed as part of a control loop for ISP content delivery services. The evaluations focused on both network and caching performance indicators, as well as management cost and complexity. **UCL** also developed a software-defined based management framework and demonstrated how this can be used to satisfy the requirements of two specific applications for adaptive load-balancing and energy management purposes [5].

OBJECTIVE 8 - To apply policy-based and semantic-based approaches for automated configuration and repair: The collaboration **UCL-iMinds-Cache** developed new pro-active cache management strategies to efficiently manage the utilization of network resources. In contrast to reactive policies (LRU/LFU), these aim at controlling the placement of content and the server selection based on user request characteristics in terms of content popularity and geographical distribution of the interests.

OBJECTIVE 9 - To propose and study automated configuration and repair in the context of the management of clouds (especially Inter-Clouds): Management of Clouds, especially Inter-Clouds has been addressed during Y3 in the collaboration **UT-UniBwM-IDS** by analysing the exchange of security events in a multi provider environment, with additional cloud specific observations. Furthermore, an architectural approach enhancing an IDS environment with SDN technology specific for this application domain has been developed. **UniBwM** has proposed an approach to utilize cloud based IDSs through the use of SDN mechanisms [17] in order to inspect network traffic locally and in cloud environments taking into account privacy concerns.

OBJECTIVE 10 - To apply the developed approaches to several application domains such as of (i) wireless sensor networks, (ii) cloud-based services, (iii) content-aware networking and (iv) software defined networking: **iMinds** has proposed an SDN-based framework to help HAS clients avoiding video freezes under scarce bandwidth conditions. The main element of this framework is an SDN-controller, which has the fundamental role of prioritizing the delivery of particular HAS segments in order to avoid video freezes. This decision is based on feedback collected from the HAS clients and on measurement data collected from the network nodes. The proposed framework has been implemented using OpenFlow, which currently represents the most important SDN protocol [18].

iMinds has applied their developed approaches on HAS and Scalable Video Coding that were developed in previous years to information centric networks. The pointed out new challenges that have been identified in deploying streaming applications on top of ICN and proposed how to solve these issues [19].

iMinds has developed a caching algorithm focusing on the characteristics of segmented video content and the trends in user behaviour for video on demand services. By taking into account the temporal structure in segmented video streams in the binge watching phenomenon for video on demand, the caching strategy can be significantly optimized in respect of the QoE compared to the state-of-the-art [13].

The new cache management approaches developed by **UCL-iMinds-Cache** focusing on multi-tenant content placement and server selection refer to the application domain of content-aware networking.

The software-defined based management framework, as well as the controller/manager placement algorithms, developed by **UCL** fall within the SDN application domain [5].

The collaboration between UPC and iMinds, developed algorithms for efficient and automated management of resources in SDN-based virtualized networks [10]

Mitigation techniques in the application domain of wireless sensor networks were developed and evaluated by the collaboration **INRIA-JUB-RPL**. This work [20] focuses on topological inconsistency attacks in RPL-based low-power lossy networks.

Techniques to detect zero day exploits using cloud-based services in combination with online social media are proposed by **UniBwM** [21]. Additional work facing online social networks was proposed in [22, 23]. Furthermore, **UniBwM** has proposed a solution to combine the possibilities of a cloud-based IDS with local privacy concerns in [17].

2.3 Tasks and Objectives Mapping

S.M.A.R.T objectives related to WP6 (Section 2.1) and WP6-specific objectives (Section 2.2) are summarized in Table 3. For each of the addressed objectives, Table 3 indicates if the objective has been achieved (S.M.A.R.T. objectives) or if there are WP activities that are contributing to the objective (WP6-specific objectives). For the WP6-specific objectives, Table 3 shows to which of the tasks in the DoW the objective is contributing to. Finally, the table acts as a guideline for the reader to locate the sections of this deliverable that provide additional information on a specific objective. Furthermore, Table 4 presents a summary of all objectives and their progress (Y1 to Y3).

Table 3: Objectives and tasks.

Objective	Task 6.1	Task 6.2	Task 6.3	Status	Details
S.M.A.R.T. Objective 1				Achieved	Section 3.1, D 8.3
S.M.A.R.T. Objective 2				Achieved	D 8.3
WP Objective 1				Ongoing	Section 3.1
WP Objective 2				Ongoing	D2.3
WP Objective 3	X			Ongoing	Section 4.1.1, 4.2
WP Objective 4			X	Ongoing	Section 4.1.3, 4.2
WP Objective 5	X			Ongoing	Section 4.1.1, 4.1.2, 4.1.3
WP Objective 6		X		Ongoing	Section 4.1
WP Objective 7	X	X		Ongoing	Section 4.2,4.1.3
WP Objective 8	X		X	Ongoing	Section 4.2,4.1.3
WP Objective 9	X		X	Ongoing	Section 4.1, 4.2
WP Objective 10			X	Ongoing	Section 4.2,4.1

2.4 Key contributions of WP6 in Y3

In Y3 WP6 mainly addressed the use cases security and content delivery. Collaborations working on these use cases have been very well active in Y3 and thus enabled WP6 to build the generic FLAMINGO integrative architecture with a promising added-value.

As a scientific output, we can report a total of 73 papers, where 40 are strongly related to WP6, that have been already published in Y3, and several other papers that are currently under review. WP6 also targeted top conferences and journals in the network management field and high-end conferences and journals in the field of networking and measurements as suggested by the reviewers during the last evaluation. Therefore, papers have been published at IEEE INFOCOM 2015, IEEE Communications and Network Security (CNS) 2015, IEEE International Symposium on Cyberspace Safety and Security (CSS 2015), ACM Multimedia Conference (ACM MM) 2015, ACM Transactions on Multimedia Computing, Communications and Applications (ACM TOMM) 2015 and IEEE International Conference for Internet Technology and Secured Transactions (ICITST) 2015. In addition, WP6 also organized the IJNM special issue “Measure, Detect and Mitigate Challenges and Trends in Network Security”, a collaboration between UT, UZH, UniBwM and CAIDA.

In Section 4.2, we summarized selected highlights of the research conducted in WP6 during Y3.

In addition, WP6 actively participated with presentations at the 2015 EuCNC Workshop on NFV and Programmable Networks and presented a tutorial during the FED4FIRE-GENI research experiment summit (FGRE 2015).

Table 4: Progress in Y3

Objective	Y1 activities	Y2 activities	Y3 activities
S.M.A.R.T. Obj 1	7 Ph.D.	14 Ph.D.	17 Ph.D.
S.M.A.R.T. Obj 2	37 papers	50 papers	73 papers
WP6 Objective 1	AIMS; Dagstuhl IM; CNSM; Coll. EU level	Dagstuhl; AIMS; EuCNC MCIS; Coll. EU level; TNSM ”Efficient Mgmt. SDN/NFV”	NetSoft; EuCNC; FGRE; AIMS;CNSM; Coll. EU
WP6 Objective 2	not addressed	wikipedia: SDN, NetFlow, sFlow	wikipedia: SDN, QoE, Adaptive bitrate streaming, 6LoWPAN, Schengen Routing
WP6 Objective 3	inventory of architectures, IDS exchange protocols	inventory of enablers, RPL attacks, SDN network attacks	Network Function Virtualization, Exchange Protocols for Security Events
WP6 Objective 4	guidelines for cloud-based services, content-aware routing	limitations in cloud-based services, content-aware routing with respect to enablers	hybrid cache strategies NFV for delivery of multimedia services SDN requirements for ress. management applications
WP6 Objective 5	architecure for cloud-based security services	architecure for VoD, automated intrusion detection	SDN-IDS architecture, IoT monitoring architecture
WP6 Objective 6	enablers for IDS, RPL security, network virtualization, QoE, QoS, traffic estimation	enablers for line card load balancing, HAS resource allocation, ressource allocation with machine learning	online virtual function mapping and scheduling, hybrid caching approach for proactive content placement
WP6 Objective 7	adaptive, energy-aware resource managent	ILP in cache management, network resource utilization mgmt.	distributed autonomic control loops coordinating proxies
WP6 Objective 8	content placement in CDNs according to policies	proactive content placement in multi-tenant scenarios	pro-active cache management strategy
WP6 Objective 9	inter-cloud security systems, VoIP security in cloud scenarios	architectural approaches for cloud-based security, mobile cloud security	security architecture including cloud-based IDS
WP6 Objective 10	RPL on IEEE 802.15.4 + 6LoWPAN, TelosB, WSN, content placement in CDNs	Internet of Things, cloud-based multi-layered intrusion detection, virtualized networks in SDN, SDN-based security mechanism	SDN-based framework to support HAS HAS in ICN, SDN-based management framework including controller/manager placement algorithms RPL incosistency attack mitigation

3 PhD Collaborations

The integration of PhD students is one of the S.M.A.R.T. objectives within this WP. Section 3.1 gives an overview of FLAMINGO collaborations that are ongoing, have been ended or are in the process of starting during this year. Furthermore, collaborations envisioned for the next year of FLAMINGO are shown.

In the overall FLAMINGO approach, monitoring (WP5) forms the basis for any automated configuration and repair action (WP6), while in parallel both activities are conducted within the boundaries of economic, legal and regulative constraints (WP7). Y3 has lead to several collaborations between these three WPs.

A detailed description about the currently ongoing collaborations and the recently completed ones, which are strongly related to WP6, can be found in Section 3.2. Thus, collaborations that are more related to WP5 do not appear in Section 3.2. The detailed description about these can be found in Deliverable D5.3. In Table 6 all ongoing WP5/WP6 collaborations are listed. However, only WP6 related collaborations are described in Section 3.2.

All fully integrated PhD students are listed in D8.3, including their co-supervisors and affiliation.

3.1 PhD Student Collaborations in Y3

The integration of PhDs into FLAMINGO allows valuable and fruitful joint research in the area of network and service management. The bottom-up approach was continued to integrate experienced researchers as well as new researchers not necessarily paid by FLAMINGO. Table 6 summarizes the collaborations, the affiliations involved and their respective status. Each collaboration can have one of the following status: **ONGOING, ENDED, STARTED, PLANNED**. **ENDED** applies to collaborations started in Y1 or Y2 of FLAMINGO and ended in Y3 because the research goals have been reached or they have branched into new collaborations. A collaboration is called **ONGOING** if started during Y1, Y2 or Y3 and progress is already reported (e.g. measurement results, planned papers, . . .). **STARTING** collaborations are in the process of defining their topic, research interests and goal of the collaboration, and drafting a plan how to reach their goal. The last type of collaborations with the status **PLANNED** have defined mutual interest in working jointly together, but did not define a concrete topic.

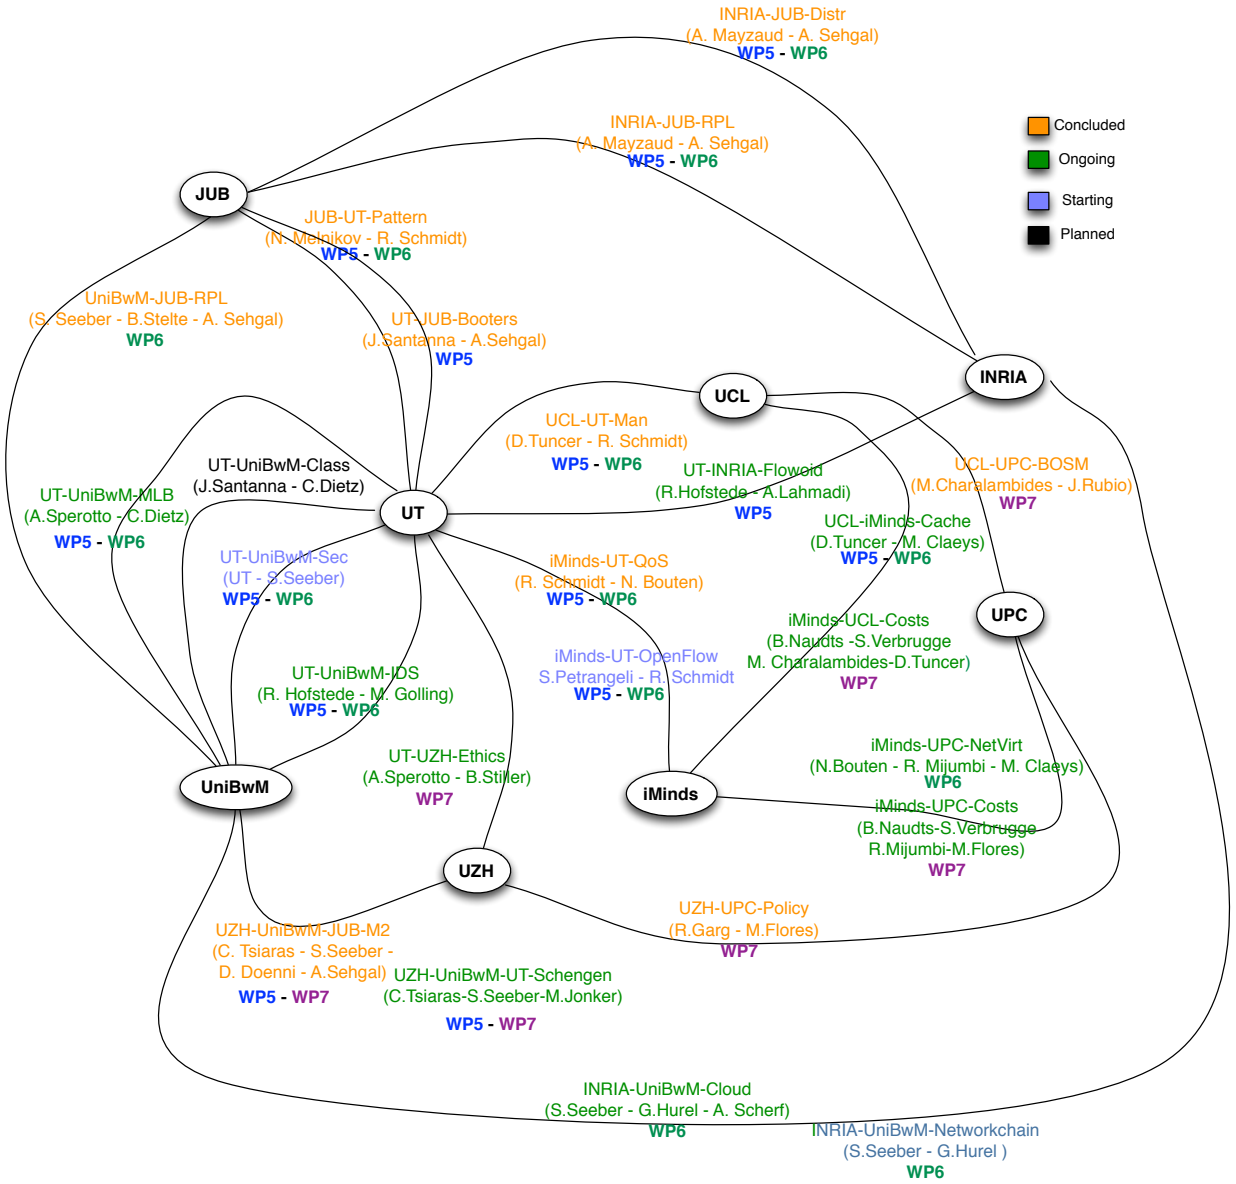

Figure 1: Overview of PhD collaborations in Y3

3.2 Description of the collaborations

This section presents the currently ongoing and recently ended collaborations between WP5 and WP6. Each collaboration description roughly follows the same structure. At first, the topic of each collaboration is explained. Subsequently, the progress and achievements in Y3 are highlighted. Depending on the status of a collaboration further steps are described. At the end each collaboration highlights the contribution to each WP.

3.2.1 Security of RPL Networks (INRIA-JUB-RPL)

The collaboration Security of RPL Networks between INRIA and JUB has led to several sub-collaborations due to the existence of various types of attacks in RPL networks. In the following a summary of each individual topic is provided.

Table 5: Overview of the FLAMINGO Collaborations, as in Figure 1

Acronym	Researchers	WPs	Status
INRIA-JUB-RPL	A. Mayzaud - A. Sehgal	WP5, WP6	Ended
INRIA-JUB-Distr	A. Mayzaud - A. Sehgal	WP5, WP6	Ended
INRIA-UniBwM-Cloud	S.Seeber - G.Hurel	WP6	Ongoing
UCL-iMinds-Cache	D.Tuncer - M. Claeys	WP5, WP6	Ongoing
UT-UniBwM-IDS	R. Hofstede - M. Golling	WP5, WP6	Ongoing
UT-INRIA- Flowoid	R. Hofstede - A. Lahmadi	WP5	Ongonig
UZH-UniBwM-JUB-M2	C. Tsiaras - S.Seeber D. Doenni - A.Sehgal	WP5, WP7	Ended
iMinds-UPC-NetVirt	N. Bouten - R. Mijumbi M. Claeys	WP6	Ongoing
INRIA-UniBwM-Chain	G. Hurel - S. Seeber	WP6	Ongoing
UZH-UniBwM-UT-Scheng	C. Tsiaras, M. Jonker-S. Seeber, L. Stiemert	WP5, WP6	Ongoing
UniBwM-UT-MLB	C. Dietz-A. Sperotto	WP5, WP6	Ongoing
UT-UniBwM-Class	J. Santanna - C. Dietz	WP5	Planned
UT-UniBwM-Sec	A. Pras - S.Seeber	WP5, WP6	Starting
iMinds-UT-OpenFlow	S. Petrangeli, R. Schmidt	WP5, WP6	Planned
JUB-UT-Pattern	N. Melnikov - R. Schmidt	WP5, WP6	Ended
UT-JUB-Booters	J. Santanna - A. Sehgal	WP5	Ended
UniBwM-JUB-RPL	S. Seeber - B. Stelte - A. Sehgal	WP6	Ended
UCL-UT-MAN	D. Tuncer - R. Schmidt	WP5, WP6	Ended
iMinds-UT-QoS	R. Schmidt - N. Bouten	WP5, WP6	Ended

Mitigating DODAG inconsistency attacks is the first sub-collaboration. Fundamentally, RPL utilizes DODAGs, a directed graph like structure, to organize the routing topology in a network. The methodology used to detect and repair possible inconsistencies in DODAG can be manipulated by malicious nodes to harm the network.

The aim of this sub-collaboration is to develop methodologies to mitigate such attacks. An approach that dynamically adapts parameters of an adaptive threshold has been developed.

The second sub-collaboration is called **RPL Version number attacks**. Version numbers are used by the RPL DODAG root in order to keep track of the latest version of the topology. If a node detects that it is part of an older version, it is required to join the new version. However, due to the lack of security mechanisms, this method could be utilized by malicious nodes to attack the topology, and possibly even hijack nodes to join its own network.

The aim of this study is to evaluate the effectiveness of attacks based on manipulating version numbers, and also study the already proposed solutions. Based on the study a new approach that overcomes existing shortcomings would be developed.

Mitigating Black-hole and Sink-hole attacks is the last sub-collaboration. The goal is to develop a mechanism that mitigates black-hole and sink-hole attacks in RPL networks, by establishing inferred trust between neighbors.

This work is currently in the development phase, with network metrics contributing towards the trust metric already identified. An implementation of the preliminary approach is currently pending.

Table 6: PhD students involved in Ongoing WP5/WP6 collaborations

Name	Affiliation	Collaborations	Acronym
Anthéa Mayzaud	INRIA	JUB	INRIA-JUB-RPL INRIA-JUB-Distr
Gaetan Hurel	INRIA	UniBwM	INRIA-UniBwM-Cloud
Rick Hofstede	UT	UniBwM, INRIA	UT-INRIA-Flowid UT-UniBwM-IDS
Ricardo Schmidt	UT	UCL	UCL-UT-Man
Mario Golling	UniBwM	UT	UT-UniBwM-IDS
Sebastian Seeber	UniBwM	UZH, JUB, INRIA	UZH-UniBwM-JUB-M2 INRIA-UniBwM-Cloud INRIA-UniBwM-Chain UT-UniBwM-Scheng
Rashid Mijumbi	UPC	iMinds	iMinds-UPC-NetVirt
Anuj Sehgal	JUB	INRIA, UT, UniBwM	INRIA-JUB-RPL INRIA-JUB-Distr UZH-UniBwM-JUB-M2
Christos Tsiaras	UZH	UniBwM, JUB	UZH-UniBwM-JUB-M2
Daniel Dönni	UZH	UniBwM, JUB	UZH-UniBwM-JUB-M2
Niels Bouten	iMinds	UPC	iMinds-UPC-NetVirt
Maxim Claeys	iMinds	UCL	iMinds-UPC-NetVirt UCL-iMinds-Cache
Mattijs Jonker	UT	UniBwM	UT-UniBwM-Scheng
Christian Dietz	UniBwM	UT	UniBwM-UT-MLB

Aspects relative to WP5 Monitoring of the RPL network and the identification of possible attacks in an RPL network contributes to WP5.

Aspects relative to WP6 The automated repair and mitigation of detected attacks in RPL networks contributes to WP6. Furthermore, the outcome of the collaboration contributes to the application area of wireless sensor networks. In general, RPL can be used as a communication protocol in the Internet of Things, which is done explicitly in this collaboration.

3.2.2 Distributed Monitoring Architecture for the Internet of Things (INRIA-JUB-Distr)

A generic distributed monitoring architecture is being designed for application in the Internet of Things (IoT) area. The goal of the architecture is to be able to monitor events and network flows passively without having any impact upon the resource constrained nodes that participate in such a network.

This monitoring architecture will be further developed to meet criteria towards anomaly detection and correction (including security aspects).

Aspects relative to WP5 Developing a distributed monitoring architecture of the IoT infrastructure can be seen as a part of WP5.

Aspects relative to WP6 An automated repair of detected anomalies in the IoT and the application of developed approaches to wireless sensor networks contributes mainly to WP6.

3.2.3 Cloud Security (INRIA-UniBwM-Cloud)

The aim of the joint research activity cloud security between INRIA and UniBwM is to investigate recently available SDN-based mechanisms for delivering security in different network scales, ranging from home networks to data centers. In the first step the scope of the study is focused on the analysis of several well-known network attack areas such as denial-of-service, information gathering and malware propagation and their distribution in cloud environments. The aim of the collaboration is to explore previous SDN attempts - such as ForCES and Active Networks - in their behavior mitigating such kind of attacks. Furthermore, the proposed approaches will be compared with the ones that are found nowadays in traditional networks (i.e. non-SDN enhanced environments). In addition, several well-known OpenFlow controllers are evaluated to identify the most suited ones for implementing security solutions in SDN networks.

3.2.4 Cache Management (UCL-iMinds-Cache)

In this collaboration, iMinds and UCL have been extending previous work in which the design and development of a proactive cache management approach for multi-tenant caching infrastructures have been investigated. The work was extended in two main directions.

The first research direction focused on the design of a hybrid cache management approach, where proactive cache reconfiguration is combined with distributed reactive cache replacement. This allows the optimization of content placement decisions, based on predicted request characteristics, while simultaneously providing reactivity to unexpected changes in the request pattern. The evaluation based on a request trace of the VoD service of a leading European telecom operator showed that the hit ratio can be increased by 40% and 19% and the bandwidth usage reduced by 5% and 7% compared to purely reactive and proactive approaches, respectively.

The second research direction addressed the limitations of the previously developed ILP-based cache management approach (in terms of scalability and complexity) by focusing on a distributed approach to control the placement of content in the available caching points. This relies on the parallelization of the decision-making process and the use of network partitioning to cluster the distributed decision-making points, which enables fast reconfiguration and limits the volume of information required to take reconfiguration decisions. The evaluation showed that a significant gain in terms of management overhead and complexity reduction can be achieved.

Aspects related to WP5 The analysis performed on the VoD traces to develop new models of prediction of the request patterns falls within the scope of WP5. In addition, preliminary work carried out to extract some pattern from the traces that could be used to model the geographical distribution of the interests for the different content items also falls within the scope of WP5.

Aspects related to WP6 The development of both the hybrid cache management approach and the distributed content distribution strategy falls within the scope of WP6.

3.2.5 Management of Virtualized Networks (iMinds-UPC-NetVirt)

This joint research activity, a collaboration between Universitat Politècnica de Catalunya (UPC) and iMinds, is referred by iMinds-UPC-NetVirt. It is aimed at managing resources in virtualized networks and functions. NFV [12, 7] is being proposed as a path towards cost efficiency, reduced time-to-markets, and enhanced innovativeness in telecommunication service provisioning. NFV leverages advances in virtualization technology to consolidate many network equipment types onto high volume servers, switches and storage, which could be located in datacentres, network nodes and in end user premises. Therefore, Service Providers (SPs) depend on virtual networks (VNs) to deploy their virtualized network functions (VNFs) in the cloud whose resources, in form of substrate networks (SNs), are owned by Infrastructure providers (InPs). However, efficiently running virtualized functions is not trivial as, among other initialization steps, it requires first mapping virtual networks onto physical networks (also known as virtual network embedding [6]), and thereafter mapping and scheduling VNFs onto the VNs. This collaboration is divided into two sub-tasks, each of which is focused to one of the above problems.

Virtual network embedding (VNE) allocates physical network resources to virtual nodes and links based on the specification in the VN requests. In the online VNE, one VN request arrives and is mapped at a time. It is therefore possible that VN requests with a low revenue per constrained resource are accepted and use up resources of the constrained node or link at the expense of VN requests that arrive later and have a higher revenue per constrained resource. The first task is to define a dynamic pricing approach that uses historic information about the resources to find the optimal price that should be charged per constrained resources based on the arrival rate, utilization rate and the number of resources requested of the constrained node or link.

In addition, since the actual loading of substrate networks varies with time [24, 8], we can combine these aspects to ensure that the revenue of infrastructure providers is maximized. The second task is based on the observation that it is possible to over-sell the SN resources with the objective that the mapped VNs load the substrate network in an efficient way, and hence improve the profitability of InPs. To this end, the proposal is to continuously forecast expected demand for SN resources, and based on this, to make both dynamic SN resource pricing decisions, as well as an evaluation of an opportunity cost that can be used to either accept or reject VN request. The main difference between the focus of this work and the state-of-the-art is that the decision to accept or reject VN requests is not only based on the availability or otherwise of resources. This means that an InP could decide to reject a VN request even if resources are available, if this will result into better profitability from the projected future VN requests. The contribution of this collaboration sub-task will be three-fold: (1) a user demand modelling approach that can be used as a basis for forecasting VN resource demand, (2) a dynamic pricing scheme that uses virtual network traffic predictions and hence expected opportunity cost (with respect to InP profit from VNE) to price substrate nodes and links, and (3) a virtual network embedding algorithm that uses future demand forecasts other than actual resource constraint to accept or reject virtual network requests.

Function Placement and Scheduling: One of the objectives of NFV is to achieve fast, scalable, on-demand and dynamic composition of network functions to a service. However, since a network service requires a number of VNFs, achieving a NFV environment raises two questions; (1) how to define and implement network services, and (2) how to efficiently map and schedule the VNFs of a given service onto a physical network. The European Telecommunications Standards Institute (ETSI) through its NFV technologies group is partnering with network operators and equipment vendors to promote the NFV approach and are currently progressing with regard to the first question above. Specifically, they have already defined the NFV problem, some use cases and a reference framework and architecture [7].

The second task of this collaboration is formulating the online virtual function mapping and schedul-

ing problem and proposing algorithms for solving it. We propose three greedy algorithms and a tabu search-based heuristic. We carry out evaluations of these algorithms considering parameters such as successful service mappings, total service processing times, revenue, cost, etc, under varying network conditions. Simulations show that the tabu search-based algorithm performs only slightly better than the best greedy algorithm. In particular, we propose some algorithms that perform the mapping and scheduling of VNFs based on a greedy criterion such as available buffer capacity for the node or the processing time of a given VNF on the possible nodes. The algorithms perform both mapping and scheduling at the same time (one-shot), i.e. at the mapping of each VNF, it is also scheduled for processing. In addition, we propose a local search algorithm based on tabu search (TS) [7]. The TS algorithm starts by creating an initial solution randomly, which is iteratively improved by searching for better solutions in its neighborhood. Finally, we also propose a an optimal mixed integer linear programming formulation of the problem, and a heuristic approach based on hard variable fixing. We also tackle the problem of placing and assigning servers that can be used to run virtualized functions [25]. This algorithms are aimed at being used as benchmarks for future algorithms in this area.

3.2.6 Network Service Chain Verification (INRIA-UniBwM-Chain)

In the context of the collaboration between UniBwM and INRIA, the focus is on the verification of network service chains. So far, a state of the art regarding this topic is established, which is a highly relevant issue taking into account the advantages of SDN and NFV. The collaboration also will design a first sketch of methodology using chained cryptographic signatures (e.g. HMAC) on network packets in order to ensure that traffic targeting this chain goes through the expected service composition. More specifically, each network function inside the service chain will have a secret key and would sign processed packets. This also allows to determine the order of the processing.

4 Automated Configuration and Repair

Given the strong involvement of several partners in the corresponding research fields, Y3 focused on the use cases of (i) security and (ii) content delivery. Work done in Y3 was based on the knowledge gained during Y1 (inventory of architectures) and Y2 (collection of enablers). Furthermore, the extension during Y2 in respect of including SDN as an additional application domain, lead to promising results in Y3, e.g. dynamic and adaptive resource management in SDN environments (see Section 4.2, 4.2).

Considering the security use case Y3 lead to promising results in respect of the construction of attack graphs, novel geolocation approaches, IDS enhancements (see Section 4.2) and mitigation strategies for RPL networks (see Section 4.2). The content delivery use case has been concentrated around HAS (see Section 4.2) and cache management in CDN (see Section 4.2). Furthermore, WP6 work addressed the development of generic approaches (see Section 4.2) towards the FLAMINGO integrative architecture (see Figure 2).

The first three subsections show an overview of research done in Y3, whereas subsequent sections present our research highlights for Y3 in more detail.

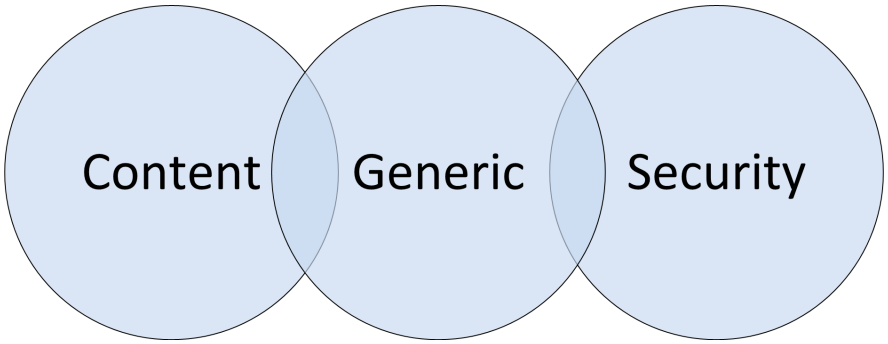

Figure 2: Relationship between use cases and generic approaches

4.1 FLAMINGO Automation Architecture, Status Y3

The FLAMINGO Automation Architecture arises from the strong connection and collaborative work of the research WPs, mainly WP5 and WP6. Where WP5 aspects focus on the monitoring, WP6 provides approaches for automated configuration and repair actions based on monitoring and analysis results. Figure 3 shows the distribution of research work done during Y3 of FLAMINGO and a separation between WP5 and WP6. This section is structured as follows: In Section 4.1.1 and Section 4.1.2 the FLAMINGO Automation Architecture is explained for two use cases which have been in the focus of Y2: (i) **security** and (ii) **content delivery**. These two use cases serve as building blocks for developing the generic FLAMINGO Automation Architecture (see Section 4.1.3). Afterwards highlights of the Y3 in the context of WP6 are explained.

The subsequent Sections 4.1.1, 4.1.2, 4.1.3 follow the same scheme of starting with a description of analytics that are build upon the results of WP5 and follows the way increasing the degree of automated configuration and repair and finally describes the inter-cloud developments within Y3.

The developed approaches withing WP6 are build upon the monitoring functionalities and datasets provided by WP5.

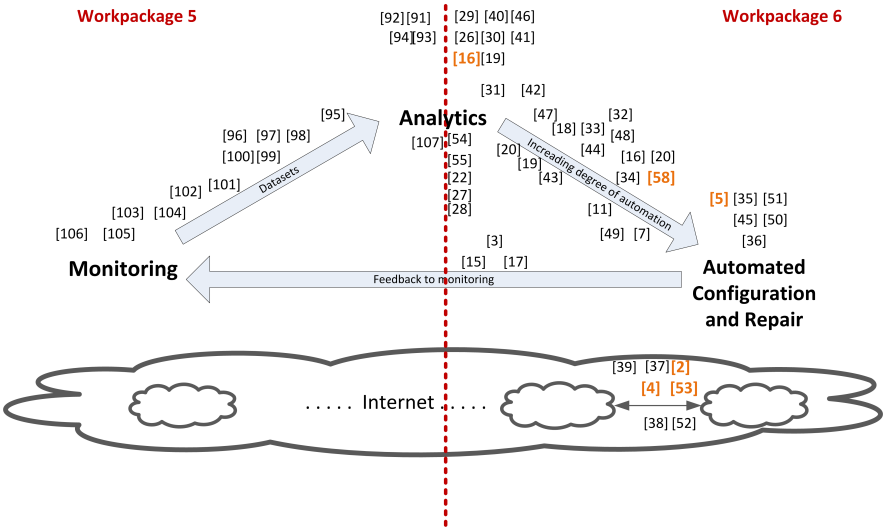

Figure 3: Research work done in Y3 separated between WP5 and WP6

4.1.1 Use Case: Security

Due to the strong involvement of several partners in the security research field, it was obvious to develop building blocks for the automation architecture in that field.

Figure 4 shows the distribution of research work during Y3 specifically for the security use case.

The work of [26] starts with a decentralized agent based framework that is able to reconstruct attack path by using the Security-by-Design paradigm. By doing so, it is possible to check plausibility for every detection result. The implementation of this approach is done by standardized protocols to prevent vendor-locks and supports a fine-grained role-based access management (e.g. different permissions for investigators and auditors). Every exchange of detection results or forensic artifacts needs a plausibility check to get acknowledged. An approach for this has been developed. Furthermore, the approach is build on RESTful web services that allow easy integration in existing solutions.

The approach described in [27] enhances the decentralized framework with geolocation capabilities. The geolocation is based on maintaining a self-build geolocation database which is based on active measurement and monitoring of e.g. BGP updates. The goal of this approach is to support the pre-incident network forensic process. The work in [28] picks up this strong relation between IT forensics and geolocation since attack traceability and attribution are two of the main tasks of IT forensics. Since primarily focused on ordinary logging, the approach proposes to take a deeper look at both degree and characteristics of logging, based on geolocation, to gather and store more evidence in advance. The additional information can be used later to reconstruct attack path to identify and to analyze distributed attacks.

Research work in [29] and [30] are solutions which support IT forensics in order to identify the source of an attack. Therefore, the approach introduces a generated ID in the IP header of an IP packet. The ID in this case identifies the outgoing interface of the source of the IP packet. Combining the information of all outgoing IDs of an IP packet allows the destination to reconstruct the path of the IP packet.

Motivated by the benefits of real-time distributed information sharing for the purpose of fast and reliable decision-making, numerous nations have been working hard over the past years to implement Network Centric Warfare (NCW). Following these considerations [31] analyses capabilities and

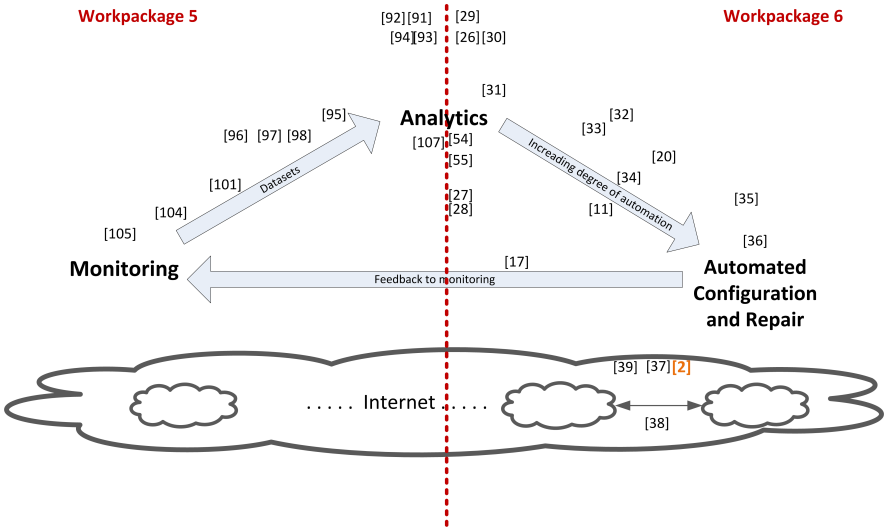

Figure 4: Research work done in Y3 for the use case security

weaknesses of NCW. Based on that, recommendations in order to strengthen the performance and accuracy for the further development of NCW are given.

Taking into account the increasing usage of online social networks [21] and [23] build a prototype to identify existing zero day attacks. The idea is to monitor online social network and in parallel maintain a database of characteristics of online services (e.g. webserver version, provided services). Correlating messages about unavailability of a specific web services the approach is able to identify possible follow-up targets for this kind of attack.

Moreover, the work presented in [32] surveys existing solutions in collaborative attack mitigation and response to identify open issues in this area. The work presents insight into processes, structures and capabilities of ISPs to mitigate and respond to network-based attacks. In the area of wireless sensor network (WSN) the detection and exchange of security events is still challenging. Therefore, [20] explores an implementation to mitigate specific types of attacks in RPL networks. The work focuses on the mitigation of topology inconsistency attacks, which allows nodes to dynamically adapt against a topological inconsistency attack based on the current network conditions. First results show that the approach outperforms the fixed threshold and mitigates these attacks without significant overhead. Details about this approach are described in Section 4.2. The issue of exchanging security events in IoT environments is addressed by the work of [33]. This work introduces a two-way authentication for IoT based on existing Internet standards (Datagram Transport Layer Security (DTLS) protocol). Relying on established standards and existing implementations as well as existing security infrastructure and engineering techniques the solution enables an easy security uptake.

Taking into account the new emerging technology of SDN the following two approaches provide detection and mitigation of security attacks grounded on cloud-based solutions. The research in [34] proposes a new approach for outsourcing mobile security functions and building transparent in-path security compositions for mobile devices. Outsourced functions are dynamically activated, configured and composed on demand. The underlying control entity uses SDN and virtualization capabilities. The security compositions are formalized in a mathematical model. In addition, an extensive set of experiments was realized during the evaluation process.

Enhancing the functionality of existing IDS using SDN is presented in [11]. The work introduces a new approach for redirecting suspicious traffic taking advantage of properties of OpenFlow in an

SDN environment. Using this, the approach is able to redirect identified suspicious traffic to various IDSs for further inspection in a dynamic and adaptive way. Furthermore, the solution is able to drop bogus traffic as well as forwarding DDoS related traffic to special DDoS capable IDS in the cloud (e.g. Cloudflare).

Fostering the further development of automated cyber defence approaches by concentrating on existing established analytics for detection purposes (see above), the following approaches provide automated configuration and repair capabilities to the underlying infrastructure. Research in [17] focuses on detecting network attacks by processing data from core network components taking advantage of properties of OpenFlow in an SDN environment. Based on this, with this approach it is able to collect metadata about forwarded traffic in an immediate and effective way. Due to SDN, it is able to steer network traffic to specific detection entities as well as actively modifying the detecting process. To preserve the protected network from attacks, that are able to drop down the connection to the Internet, the approach is able to redirect suspicious detected traffic with a high severity to cloud-based detection and mitigation solutions in a privacy preserving manner. The aim of steering the traffic in this way is to keep away suspicious traffic as much as possible from the attack target network.

Research in [35] and [36] focus on the automated trust establishment in IoT environments. Therefore, the work proposes a secure and efficient key management which is necessary to protect the authenticity, integrity, and confidentiality of multicast messages. It develops two group key establishment protocols for secure multicast communications among resource-constrained devices in IoT. Since this work focuses on multicast communication, the aim of further research is to broaden the supported type of communications. The previously explained approaches are focused mainly on single domains or cloud environments. In the following paragraph we will explain the progress made in the context of multi domain / inter-cloud security approaches.

On of the main challenges in a multi domain / inter-cloud environment is the exchange of security events that occur in the respective networks and the coordinated reaction to streamline the mitigation of an attack. Therefore, [37] and [2] focus on the exchange of security events of flow-based IDSs. The work proposes a new exchange format, called Flow-based Event Exchange Format (FLEX). It is placed in high-speed networks that use links S/MIME signature with a speed of 10 Gbps and higher, and use flow export technologies (e.g. Cisco NetFlow, IPFIX) to identify, track and mitigate malicious traffic. Further, FLEX is intended to facilitate the cooperation among network operators and focus on an automated threat information exchange. In addition, FLEX messages are disseminated using SMTP, FLEX is easy to deploy and it integrates with existing infrastructure.

Since not only flow-based events need to be exchanged, it is also necessary to compare existing security exchange mechanisms. A detailed description of this work is presented in Section 4.2. Taking into account the capabilities of SDN, in [38] the approach is primarily measurement-based, in which measurements are first focused on assessing the applicability of OpenFlow-enabled devices for DDoS mitigation. The approach in [39] maintains a logically centralized database that provides latest security related information about each system or service. Using this knowledge base, it ponders a systems' security score, security requirements given by the systems' owners and the cloud provider, and reconfigures the network accordingly to meet the security requirements for every system. In addition, the reconfiguration process can be used to redirect traffic to additional security systems, in order to obtain more detailed information about a system and therefore increase the accuracy of the specific systems' security score. The aim of abstracting to these scores is to enable an exchange of vulnerability and security related information between multiple providers and cloud environments, without disclosing the internal network structure and vulnerability details.

4.1.2 Use Case: Content Delivery

Due to a strong content delivery related background of partners in FLAMINGO, the content delivery use case has been addressed. The research scope varies from multimedia delivery to Voice over IP and content distribution. Figure 5 shows the distribution of approaches developed during Y3 of FLAMINGO for this use case.

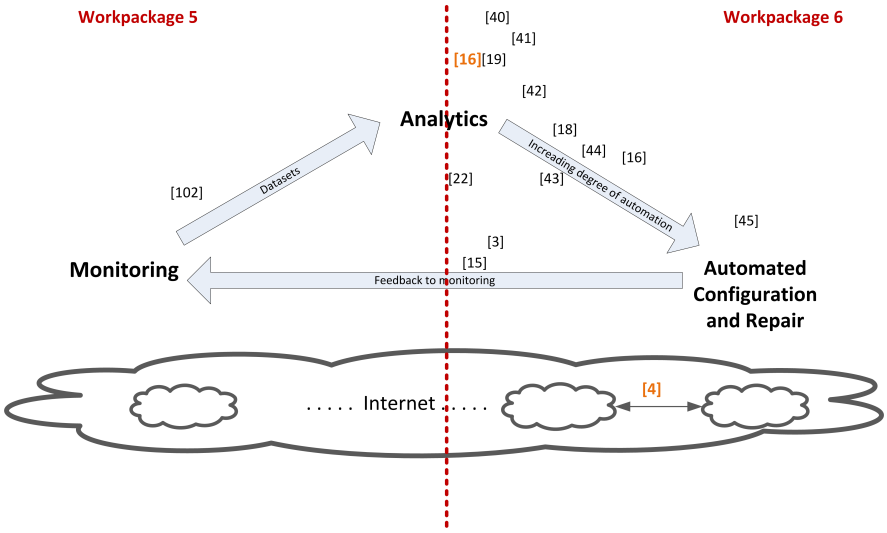

Figure 5: Research work done in Y3 for the use case content delivery

Research in [19] includes an analytics part that analysis the varying bandwidth availability to adapt the quality of the requested content accordingly. Since the work focuses on Information Centric Networking (ICN), which was the recently proposed disruptive architecture that could solve the issue of the optimal support for multimedia delivery, where the focus is given to the content rather than to end-to-end connectivity. Taking into account the bandwidth unpredictability which is typical for ICN, standard Advanced Video Coding(AVC)-based HAS performs quality selection sub-optimally, which leads to a poor Quality of Experience (QoE). The work proposes using Scalable Video Coding (SVC) instead. The research focuses on individuating the main advantages of SVC-based HAS over ICN and outlines the research challenges to be addressed to optimize the delivered QoE. Furthermore, the approach developed in [16] addresses the issues arising in a multi-client setting. Particularly, the work presents a fair HAS client able to achieve smooth video playback, while coordinating with other clients in order to improve the fairness of the entire system. This specific goal is addressed by the support of an in-network-based system of COORDINATION PROXIES, which are in charge of collecting measurements on the network conditions. In the next step this information is used by the clients to refine their quality decision process and develop a fair behavior. Further details about this work are described in Section 4.2.

In the area of content distribution, the work [40] and [41] focuses on the facility location problem which is still a well-known challenge in logistics and proven to be NP-hard. In this work a simulation of the geographical placement of facilities to provide adequate service to customers took place. Generally, the objective is to place the central nodes such that all customers have convenient access to them. Several existing approaches were compared and a new heuristic for the problem was proposed, which shows a significant improvement.

In the area of VoIP, the work [42] proposed an approach to calibrate the Deterministic QoE model (DQX), which can be used to capture end-user's QoE in VoIP services. Such a calibration of the model is essential to adapt it to the particular service and its technical and non-technical conditions

in which it is used. Furthermore, those DQX results achieved are compared with those results of the IQX Hypothesis and the E- Model, being proposed by the ITU-T. Thus, it is finally shown that DQX can capture more accurately end-user's QoE in VoIP scenarios. To provide best QoE for end-users, it is necessary to identify the user who is utilizing a specific service to start an automated configuration of the network which is delivering the content most likely requested by the user. To identify the user, the work [22] matches profiles of online social network users with respective geolocation tags. In addition, the work proposes the usage of timestamps, which are generated by the social network, and device-generated geo-tags. Furthermore, a comparison of various approaches for the implementation of profile matching algorithms is presented. The management of advanced multimedia services is still challenging, because the Internet was not designed to deliver such real-time, bandwidth-consuming applications.

A serious challenge is posed on how to efficiently provide the best service to the users. Therefore, [16] reviews the main challenges and tackles the field of end-to-end QoE optimization of video streaming services and HAS solutions, which are quickly becoming the de-facto standard for video delivery over the Internet. This work envisions the placement of in-network nodes to collect feedback regarding the network and clients' conditions and influence client's behavior. The main advantage of this approach is three-fold. First, there is no communication needed among the clients and consequently no significant overhead introduced. Second, the quality level selection is still performed locally and independently by each client. Third, the approach is robust toward network failures, as the clients can also operate (at a sub-optimal level) without the in-network system. Current approaches in this area are based on the implementation of a rate adaptation heuristic based on a multi-agent version of the Q-Learning algorithm. The principle of HAS is becoming the de-facto standard for video streaming services over the Internet. In HAS, each video is segmented and stored in different qualities. To adapt and change to a new quality level heuristics are used to allow the client to request video segments dynamically based on the current network condition. Current heuristics under-perform when sudden bandwidth drops occur, therefore leading to freezes in the video play-out, the main factor influencing users' Quality of Experience (QoE). Furthermore, [18] proposes an OpenFlow-based framework capable of increasing clients' QoE by reducing video freezes. An OpenFlow-controller is in charge of introducing prioritized delivery of HAS segments, based on feedback collected from both the network nodes and the clients. Furthermore, a novel mechanism is introduced to inform the clients about the prioritization status of the downloaded segments without introducing overhead into the network. This information is then used to correct the estimated bandwidth in case of prioritized delivery. The approach was evaluated through emulation, under varying network conditions and in several multi-client scenarios. The results show a reduction of freezes up to 75% compared to state-of-the-art heuristics.

Using a learning-based algorithm the work [43] improved bandwidth awareness of adaptive streaming clients. Current quality selection heuristics are generally hard coded. Fixed parameter values are used to provide an acceptable QoE under all circumstances, resulting in suboptimal solutions. Furthermore, many commercial HAS implementations focus on a video-on-demand scenario, where a large buffer size is used to avoid play-out freezes. When the focus is on a live TV scenario however, a low buffer size is typically preferred, as the video play-out delay should be as low as possible. Hard coded implementations using a fixed buffer size are not capable of dealing with both scenarios. The approach introduces the concept of reinforcement learning at client side, to adaptively change configuration. In addition, this approach takes into account bandwidth characteristics during the decision process, in order to improve the client's bandwidth-awareness.

Focusing on the live experience of adaptive streaming [44] follows an approach to improve the adaptive streaming using HTTP/2 methods. Due to their advantages compared to traditional techniques, HAS-based protocols are widely used for over-the-top (OTT) video streaming. However, they are yet to be adopted in managed environments, such as ISP networks. A major obstacle is

the purely client-driven design of current HAS approaches, which leads to excessive quality oscillations, suboptimal behavior, and the inability to enforce management policies. The work [45] addresses these issues and facilitates the adoption of HAS in managed networks. Therefore, several centralized and distributed algorithms and heuristics are proposed that allow nodes inside the network to steer the HAS client's quality selection process. The algorithms are able to enforce management policies by limiting the set of available qualities for specific clients. The work starts with a formal definition of the in-network rate adaption problem. Based on this an optimal centralized algorithm is proposed that solves the problem as an Integer Linear Program (ILP). Afterwards, a scalable variant of the algorithm is introduced that can be distributed across multiple logically hierarchical intermediary proxies. Finally, a heuristic with significantly lower computational complexity is proposed. Using packet sampling measurements the work [15] focuses on QoE-driven in-network optimization for Adaptive Video Streaming. Based on the concept of HAS the client can autonomously decide, based on the current buffer filling and network conditions, which quality representation it will download. Each of these players strives to optimize their individual quality, which leads to bandwidth competition, causing quality oscillations and buffer starvations.

A solution is presented to alleviate these problems by deploying in-network quality optimization agents, which monitor the available throughput using sampling-based measurement techniques and optimize the quality of each client, based on a HAS Quality of Experience (QoE) metric. This in-network optimization is achieved by solving a linear optimization problem. Supported by the promising idea of NFV [3] investigates how existing service chains from datacenter network can be mapped onto NFV-based Service Function Chains (SFC). Furthermore, the different alternative SFCs are explored and their impact on network and datacenter resources (e.g., bandwidth, storage) are quantified. The approach proposes to use these findings to cost-optimally distribute datacenters across an Internet Service Provider (ISP) network. In the area of inter-cloud content delivery networks the placement of caches in the network remains still challenging, since the management of resources is still an open issue. Therefore, [4] proposed a proactive cache management system for Internet Service Provider (ISP)-operated multi-tenant Telco CDNs.

In this approach, a central manager periodically computes a caching configuration based on predicted values of the future request pattern. This covers the allocation of caching capacity across the network for the multiple tenants, the proactive placement of content and the server selection strategy (where to serve each request from). This problem was modeled as an Integer Linear Programming (ILP) problem with the objective of minimizing the bandwidth usage inside the ISP network. Details about this approach are described in Section 4.2.

4.1.3 Generic Approaches

This section concentrates on the description on generic approaches developed in Y3 of FLAMINGO that are usable in various use cases. Figure 6 shows the research developed during Y3 of FLAMINGO that contributes to this topic.

Presenting a novel approach to traceback IP packets for data-flow analysis [46] investigates how to identify the exact path that packets are routed in the network. The concept, named Tracemax, allows a detailed analysis of traffic and the transmission paths through the network. It consists mainly of a marking scheme and a reconstruction method. The routers are marking packets on the path during the transmission. The reconstruction method determines the path of a packet afterwards. The increasing demands on traffic and the current trend of network and services virtualization calls for effective approaches for optimal use of network resources.

In the Future Internet multiple virtual networks will coexist on top of the same physical infrastructure, and these will compete for bandwidth resources. Link dimensioning can support fair share and

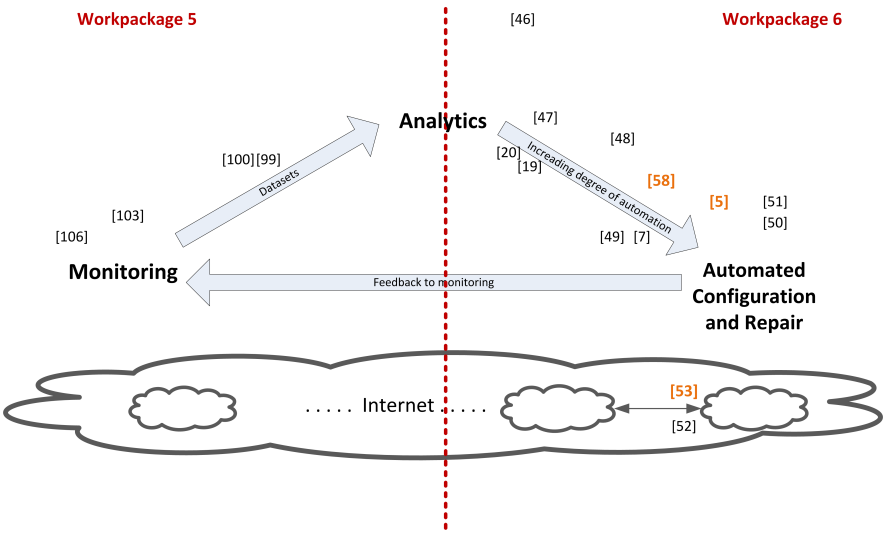

Figure 6: Research work done in Y3 for generic approaches

allocation of bandwidth. Therefore, [47] developed and validated link dimensioning approaches that estimate the needed traffic statistics from measurement data obtained via technologies that are largely found in today's networks (namely, sFlow and NetFlow/IPFIX). Furthermore, a link dimensioning approach is envisaged that uses measured data from the recent and already widely available OpenFlow. In addition, the quality of flow-level measurements in current implementations of OpenFlow is evaluated, which showed that these are not yet accurate enough for link dimensioning purposes. Automatic network configuration, especially in the case of P2P networks, where millions of nodes worldwide in environments that range from static to very dynamic and therefore exhibit different churn levels, can make a system more adaptable to changing environments and reduce manual configuration tasks. Therefore, the approach [48] proposes an automatic replication configuration based on churn prediction that automatically adapts its replication configuration to its environment. The mechanism termed dynamic replication mechanism (dynamic RM) developed and evaluated in this work is based on exponential moving averages to predict churn that is used itself to determine a replication factor meeting a certain reliability threshold.

Focusing on resource management in virtualized networks [58] proposes to use the SDN control plane to efficiently manage resources in virtualized networks by dynamically adjusting the virtual network (VN) to substrate network (SN) mappings based on network status. The work extended an SDN controller to monitor the resource utilisation of VNs, as well as the average loading of SN links and switches, and uses this information to proactively add or remove flow rules from the switches. Details about this work are presented in Section 4.2.

The challenging aspects of mapping and scheduling of virtual network functions are part of the work in [7]. Efficiently running virtualized services is not trivial as, among other initialization steps, it requires first mapping virtual networks onto physical networks, and thereafter mapping and scheduling virtual functions onto the virtual networks. The approach formulates the online virtual function mapping and scheduling problem and proposes a set of algorithms for solving it. The main objective is to propose simple algorithms that may be used as a basis for future work in this area. Furthermore, three greedy algorithms and a tabu search-based heuristic are presented.

In the area of self-management of virtual network resources [49] presents a reinforcement learning-based neuro-fuzzy algorithm that perform dynamic, decentralised and coordinated self-management of substrate network resources. The objective is to achieve better efficiency in the utilisation of substrate network resources while ensuring that the quality of service requirements of the virtual

networks are not violated. The proposed algorithms are evaluated through comparisons with a Q-learning-based approach as well as two static resource allocation schemes. Following the self-management principle [50] contributes to efficient resource sharing in network virtualisation by dividing the resource management problem into three sub-problems: virtual network embedding (VNE), dynamic resource allocation (DRA), and virtual network survivability (VNS). In addition, the work proposes a solution for each one of them. Specifically, a path generation-based approach for VNE, machine learning-based self-management approaches for DRA, and a multi-entity negotiation algorithm for VNS is presented. While using a centralized controller within SDN with a network-wide view has the benefit of facilitating a fairly straight forward implementation of the control logic, it also presents limitations, especially in terms of scalability as the size and dynamics of the network increase. In addition, resource management in fixed networks is usually performed by external offline centralized systems, which is not adequate to support applications that adapt to traffic and network dynamics. In [5] a SDN-based network resource management and control framework that can support both static and dynamic management applications is developed. A key characteristic of the proposed framework is its modular structure, which resides in the separation between the management logic and the control logic. Details about this management and control framework are outlined in Section 4.2.

In the area of WSN [51] developed a framework that supports mobility requirements and incorporates online database storage, access control management, and visualization with responsive design for different screen sizes of mobile devices (e.g., smartphones, tablets). The goal of this framework is to allow users to configure and manage their WSNs, but also to monitor them independent of the user's location. Implementing automated configuration and repair independent on a specific use case requires a strong abstraction of the environment where the action should take place. Besides this generalization can also be achieved in applying this automated configuration and repair at a low level below the application layer. This is especially necessary in inter-cloud environments, where network virtualisation continues to receive attention and recent proposals have advocated for survivability in network virtualisation environments (NVEs).

However, research work within the same area has mainly focused on the single provider environment, leaving network survivability in multi-domain environments largely unexplored. In particular, survivability in heterogeneous physical networks raises questions with regard to the negotiation between competing parties so as to form coalitions for resource provisioning. Therefore, [52] proposes a distributed negotiation algorithm which uses a system of entities to support survivability in a multi-domain NVE. The objective is to make each of the virtual network providers adaptive and dynamic by modeling them with capacity to perform QoS aware resource back-ups and/or restorations for physical link failures. Another form of generic automated configuration and repair can be abstracted to geometric routing schemes which are proposed as an alternative for lookup-based routing algorithms. Although they were initially designed for Unit Disk Graphs (UDGs), their application to scale-free complex networks has been demonstrated. [53] explores the possibilities for combining low stretch with load balancing behavior. The main contribution is a family of routing schemes called Forest Routing (FR). These algorithms are capable of adapting their routing behavior to varying traffic intensities by using a generalized distance function incorporating link load information. A more detailed description of this approach can be found in 4.2.

To sum up, the above three Sections 4.1.1, 4.1.2, 4.1.3 present the steps towards the FLAMINGO integrative architecture for automated configuration and repair. Due to the two main pillars in the use cases security (Section 4.1.1) and content (Section 4.1.2) the FLAMINGO integrative architecture for automated configuration and repair is based on strong and well proven approaches. Not only the number of publications in these two areas, also the quality and spread among conferences and journals out of the typical network and service management community point out the performance of WP6 in FLAMINGO in Y3. Based on these two use cases and the developed generic

approaches an essential step towards the FLAMINGO integrative architecture for automated configuration and repair was done.

4.2 Research Highlights of Y3

In this section and in the following sections we summarized selected highlights of the research conducted in WP6 during Y3.

Dynamic Resource Management in SDN-based Virtualized Networks [54]

Network virtualization has emerged as a promising technology for the Future Internet in which network deployment and management are separated from service provision [55]. Specifically, an infrastructure provider (InP) owns, controls and manages physical resources in form of substrate networks (SNs), which may be used by one or more service providers (SPs) to create virtual networks (VNs) to provide services to end-users. However, hosting multiple VNs and supporting their complete isolation raises resource management (RM) challenges for the InP, e.g. the need to efficiently allocate SN resources to multiple VNs.

SDN [56] is an appealing platform for network virtualization environments (NVE), since each VN's control logic can run on a controller rather than the physical switches [57]. SDN allows for a flexible and easier way of defining VNs, say, by representing each virtual link as a flow and hence defining a VN as a set of flow rules in different switches. This way, SDN's control plane can be used to achieve important resource allocation policies such as SN load balancing, VN resource cost minimization, e.t.c. For instance, Flowvisor [58] and XNetMon[59] allow multiple tenants to share an SDN substrate through virtualization by allowing for isolation and sharing of network slices.

However, current proposals for virtualized SDNs are silent about the RM requirements that have to be taken into account in such an environment. For example, an important step in initializing VNs is the mapping of virtual nodes and links to substrate nodes and links. While this mapping is a well-studied problem [60], as shown in this paper, some of the resource mapping approaches such as path splitting [61] that have been shown to lead to better resource utilisation in VNs create another problem in an SDN environment. When a virtual flow is split into multiple sub-flows, each sub-flow would need flow rules in each of the switches along the substrate path that supports it, hence requiring more ternary content-addressable memory (TCAM), which is expensive to build, consumes a lot of power and dissipates a high level of heat [62]. In addition, if performed in a static way, virtual to substrate resource mapping leads to high resource fragmentation at the SN layer [63]. Therefore, dynamic RM leads to better resource utilisation efficiency [64], since VN requests arrive and depart in a dynamic manner. Current approaches to dynamic RM in virtualized networks are mainly based on link migration [60], which is aimed at balancing the load on substrate links without considering the effect on the substrate node resources. As already mentioned, given the cost and power dissipation [65], [66] of node resources in SDN environments, it is necessary for a RM approach to also be node resource aware and manage them.

In this paper, the authors propose a flow migration approach to dynamically manage link and switch resources in an SDN-based virtualized environment which does not only consider link resources, but also node resources and VN resource costs. To this end, they extend a floodlight controller [67] by adding an application module which monitors the resource costs of mapped virtual links, as well as average load of the substrate links and switches. This information, coupled with that about arrivals and departures of VN requests is used to determine which virtual links can profitably be migrated. The module then proactively modifies (adds and/or deletes) flow rules (which represent

virtual links) from the affected switches. The idea of the proposal is that due to the dynamic arrival of VN requests, some virtual flows may utilize more resources at the time of mapping, but when some VNs leave, more efficient flows can be established.

Adaptive Resource Management and Control in Software Defined Networks [5]

The heterogeneous nature of applications, technologies and multi-vendor equipment, which form todays networking landscape, have made the management of network infrastructures a very complex task. Over the past few years, the SDN paradigm has gained a significant interest from both the industry and the research community, who envision SDN-based solutions as a key enabler towards simplifying the management processes. In the SDN architecture, control functions are moved away from the network devices, which are treated as basic forwarding elements, towards external dedicated software-based components, referred to as the controllers, forming a unified control platform. This can be seen as a logically centralized control plane which operates on a global network view and implements a range of functions. While using a centralized controller with a network-wide view has the benefit of facilitating a fairly straight forward implementation of the control logic, it also presents limitations, especially in terms of scalability as the size and dynamics of the network increase. In addition, resource management in fixed networks is usually performed by external offline centralized systems, which is not adequate to support applications that adapt to traffic and network dynamics.

To overcome these limitations, UCL has developed in [5] a SDN-based network resource management and control framework that can support both static and dynamic management applications. A key characteristic of the proposed framework is its modular structure, which resides in the separation between the management logic and the control logic. More specifically, the framework follows a hierarchical architecture and relies on three layers. The bottom layer concerns the physical network infrastructure, represented by a set of network devices and network links. The middle layer represents the distributed management and control layer in which local managers (LMs) and local controllers (LCs) (software components) form separate management and control planes. These are responsible for managing and controlling the configuration of the network resources. Finally, the top layer concerns the central management system, which is responsible for longer term operations, for example those that pertain to the life cycle of LMs and LCs. The interaction between the components of the architecture is realized through a set of intra and inter layer interfaces as represented in the Figure 7.

In the proposed framework, short to medium term management operations are realized through the LMs, which implement the logic of management applications (e.g. online monitoring/sampling, adaptive traffic engineering, etc.). Each LM is in charge of a subset of network resources. Based on information monitored locally or obtained from other LMs in the management plane, each LM is responsible for determining the changes to apply in order to (re)configure the settings of the network devices under its super-vision. The output of the reconfiguration is then passed to the LC(s), which define(s) and plan(s) the sequence of actions to be enforced for updating the relevant network parameters. The actions are then mapped to instructions sent-to and executed-by the network devices. This modular structure of the framework enabling the separation between management and control functionality offers significant benefits. First, it provides more flexibility in terms of deployment as changes can be applied to LMs in an operational environment independently of the LCs and vice versa. In addition, it facilitates the integration of new applications, as management and control functions are not tight to specific implementations and can therefore evolve independently.

A key challenge when implementing the proposed framework concerns the degree of distribution of the entities in each of the management and control planes. In practice, the number of LMs and LCs

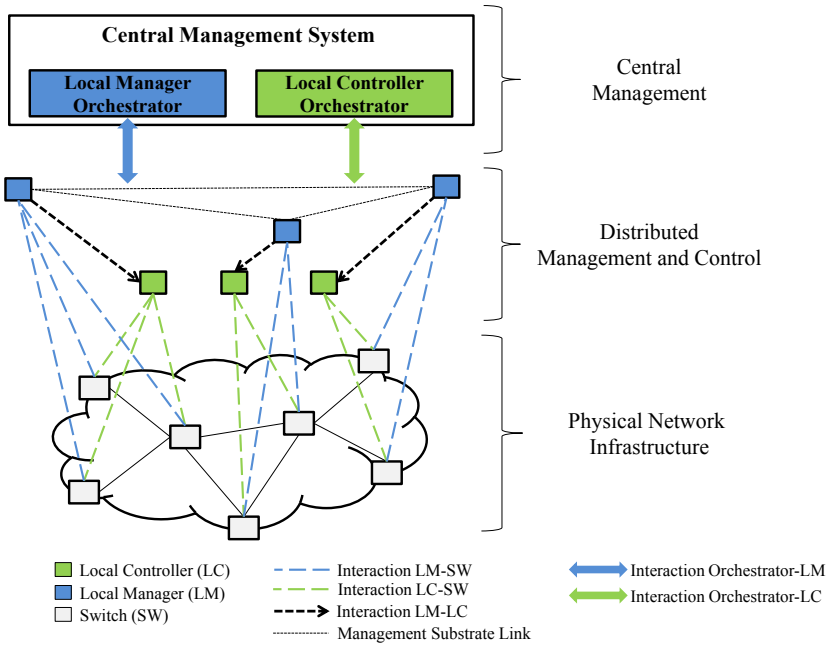

Figure 7: Component interaction

to deploy, as well as the association between the two, should be driven both by the characteristics of the physical infrastructure as well as the type of management applications to consider. In [5], UCL has developed an approach to determine the allocation (number and location) of LMs and LCs in the proposed distributed management and control layer. Given a network topology, the objective of the proposed approach is to compute the number of LM/LCs to deploy, their location, as well as the devices these are connected to, with the objective of minimizing the distance (in terms of hop count) between the network devices and the LM/LCs.

The proposed placement algorithm follows a greedy approach, so that LM/LCs are iteratively added one-by-one in the network. The new location selected at each iteration is the one that leads to the highest reduction in terms of average distance between the LM/LCs and the devices. The output of the algorithm provides the number of LM/LCs to deploy, their location, as well as the configuration of their mapping to network devices. To control the output, the algorithm relies on two tunable parameters: a) the initial placement metric, which is used to determine the location of the initial LM/LC and is defined based on properties taken from graph theory, and, b) the ending threshold, which is used to decide when the algorithm terminates and is based on a measure of the average distance reduction at each iteration. The value of the algorithm parameters is configured based on both topological factors and the requirements of the considered management applications.

The benefits of the proposed framework is demonstrated based on two distributed adaptive resource management applications for adaptive load-balancing and energy management whose performance is evaluated in terms of resource utilization reduction. The results showed that a significant gain in terms of link utilization and energy consumption can be achieved in a scalable manner.

QoE-Driven Rate Adaptation Heuristic for Fair Adaptive Video Streaming [16]

Nowadays, video streaming applications are responsible for the largest portion of the traffic exchanged over the Internet. Particularly, HAS protocols have become very popular due to their

flexibility, and can therefore be considered as the de-facto standard for video streaming services. Microsoft's Smooth Streaming (MSS), Apple's HTTP Live Streaming, Adobe's HTTP Dynamic Streaming and MPEG Dynamic Adaptive Streaming over HTTP (DASH) are examples of available HAS technologies. In a HAS architecture, video content is stored on a server as segments of fixed duration at different quality levels. Each client can request the segment at the most appropriate quality level on the basis of the local perceived bandwidth. In this way, video playback dynamically changes according to the available resources, resulting in a smoother video streaming experience. The main disadvantage of current HAS solutions is that the heuristics used by clients to select the appropriate quality level under-perform in a multi-client scenario [68],[69],[70]. In a real scenario, multiple clients simultaneously request content from the HAS server. Often, clients have to share a single medium and issues concerning fairness among them appear, meaning that the presence of a client has a negative impact on the performance of others. As reported by Akhshabi et al. [68], fairness issues are not due to TCP dynamics, but mainly arise from the rate adaptation algorithms, as they decide on the actual rate to download. When multiple clients stream a video at the same time, wrong bandwidth estimation can occur, due to the temporal overlap of the activity-inactivity periods of different clients. This wrong estimation subsequently affects the bit rate selection and thus the clients' Quality of Experience (QoE). This problem is aggravated by the uncoordinated nature of current HAS heuristics. This entails they are not aware of the presence of other clients nor can they adapt their behavior to deal with it.

In this paper [16], they investigate the aforementioned problems arising in a multi-client setting. Particularly, they present a fair HAS client able to achieve smooth video playback, while coordinating with other clients in order to improve the fairness of the entire system. This goal is reached with the aid of an in-network-based system of COORDINATION PROXIES, in charge of collecting measurements on the network conditions. This information is then used by the clients to refine their quality decision process and develop a fair behavior.

The main contributions of this paper are three-fold. First, a new HAS heuristic called FINEAS (Fair In-Network Enhanced Adaptive Streaming) able to select the best quality depending on network conditions, in order to provide a smooth video streaming and improve fairness is presented. Particularly, our heuristic is able to increase the average requested quality level compared to current HAS heuristics and avoid video freezes, while guaranteeing similar QoE to the all the clients streaming video, i.e. fairness. Second, they design an in-network-based system to help clients coordinate their behavior, which does not require explicit client-to-client communication or a centralized decision process. Consequently, the quality level selection can still be performed locally and independently by each client, without any modification to the general HAS principle. Third, detailed simulation results are presented to characterize the gain of the proposed framework compared to state-of-the-art HAS heuristics.

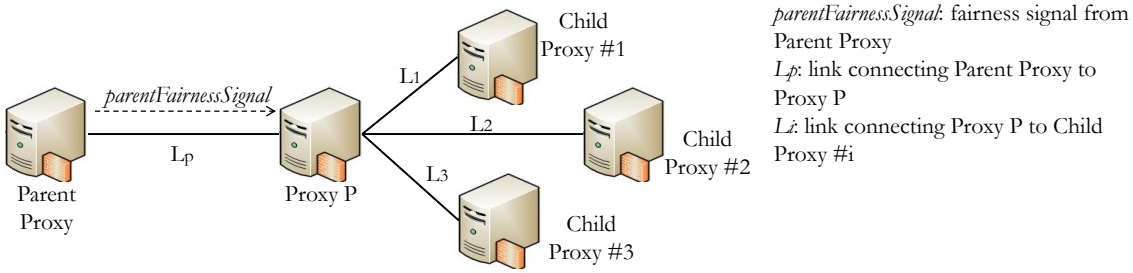

Figure 8: Schematic representation of the coordination proxies' architecture

The problem proposed to investigate in this paper is two-fold. First, clients have to obtain the highest possible video quality. Second, they have to show similar performance if they share bottle-

neck links, i.e. fairness. Based on this consideration, all the clients sharing the same bottlenecks should act fairly, even if they belong to different networks. In order to maximize the QoE delivered to the clients and achieve fairness, they present the FINEAS (Fair In-Network Enhanced Adaptive Streaming) heuristic. The FINEAS heuristic runs at the clients and performs the quality level selection based on three inputs: the local perceived bandwidth, the video player buffer status and the so-called FAIRNESS SIGNAL. The fairness signal is an additional measure introduced to achieve fairness, obtained when the client downloads a segment. The fairness signal is computed by a system of network nodes, called COORDINATION PROXIES, and represents an estimate of the fair bandwidth share of all the clients streaming video.

As introduced previously, the system of coordination proxies is in charge of helping clients achieving fairness, by computing an estimate of the fair bandwidth share for all the clients streaming video, even if they belong to different networks. In order to maintain scalability, the computation of the fairness signal is performed periodically and in a hierarchical way by the coordination proxies (see Figure 8). A generic coordination proxy P receives an estimate of the fairness signal from its parent node and computes a new estimate of the fairness signal for each of its child proxies. This estimate is computed by monitoring the available bandwidth for HAS traffic on the links connecting proxy P to its child nodes. In order to limit overhead, the calculated fairness signal can be added as an HTTP header field and returned to the clients when delivering the next segment to play. Particularly, the clients translate the fairness signal into a REFERENCE quality level, representing the theoretical quality level to request in order to obtain perfect fairness among the clients. This reference gives an indication on the best quality level to achieve fairness, rather than determining the actual quality to be requested. The reason for this behavior is two-fold. First, directly requesting the reference quality level would be optimal from the fairness point of view but not from the QoE point of view, because of the frequent switches that would occur. Second, directly requesting the reference quality level would alter the classical HAS principle, as the decision on the quality level to download would no longer be carried out by the clients.

The main advantage of this hybrid approach is three-fold. First, no communication is needed among clients and consequently no significant overhead is introduced. Moreover, no client-to-proxy communication is required. The proxies are TRANSPARENT to the clients, as the clients only need to know how to access the fairness signal but not how it is created. Second, the computation and delivery of the fairness signal do not negatively affect the behavior of existing clients. Third, the approach is robust toward proxy failure, as the clients can also operate without the fairness signal.

As far as the coordination proxies positioning is concerned, the proxies should be located at the main aggregation points of the network, in order to monitor the links where a bottleneck can occur. Potential bottlenecks can be identified by analyzing the underlying network architecture or at runtime by monitoring link conditions (e.g., if the traffic exceeds a certain percentage of the link capacity, a coordination proxy can automatically become active). Since network operators have full control of their delivery infrastructure, they can easily identify which are the most sensible paths in their networks where a bottleneck could occur. This way, they can perform an initial placement of the coordination proxies on network nodes. Note that this assumption does not impact the flexibility of the solution, since in a real scenario the network architecture is given and does not significantly change over time. Furthermore, as coordination proxy functionalities can be implemented via software, proxies can be flexibly relocated in case network conditions consistently change over time. Moreover, coordination proxies can be placed liberally on network nodes, without negatively impacting the fairness signal computation even if a bottleneck does not occur. In this case, a proxy only receives the fairness signal from its parent node and forwards it to its child proxies, without performing any operation on it. In other words, if a bottleneck does not occur, the considered proxy becomes transparent with respect to the computation of the fairness signal.

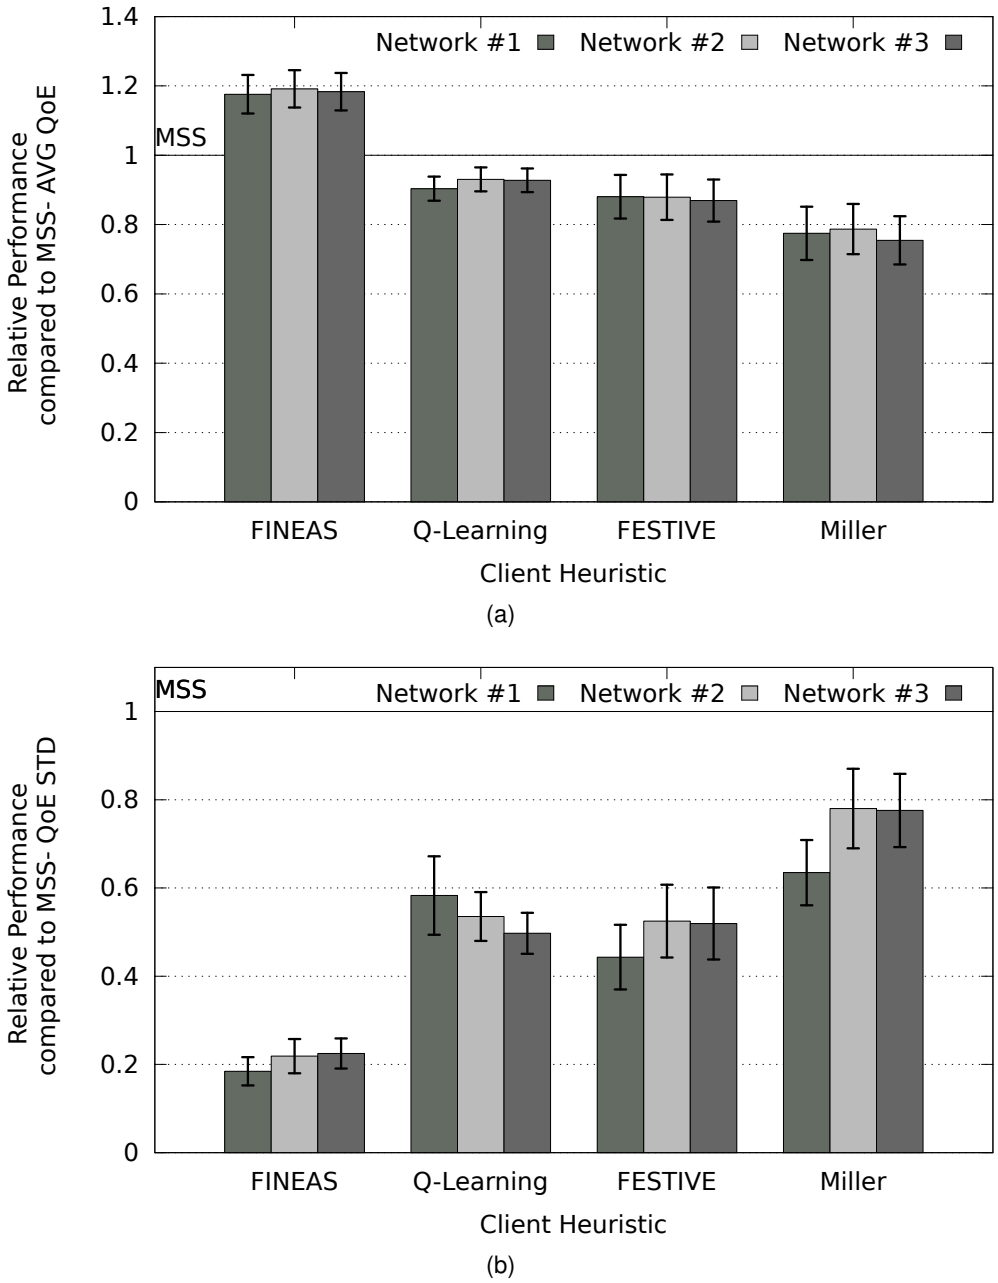

Figure 9: Comparison between the different clients, from a QoE perspective, for a variable bandwidth scenario. Each network contains 30 clients streaming video. The graphs report the relative performance of the considered clients in terms of (a) average QoE and (b) its standard deviation compared to MSS. The standard deviation of clients' QoE is used as fairness metric.

Figure 9 shows some results from the presented proposal. In order to provide an extensive benchmark of the FINEAS algorithm, they compare their results to those obtained using four other HAS clients. Particularly, they choose a proprietary HAS client, the MSS client, the Q-Learning-based client described by Claeys et al. [71] and the client developed by Miller et al. [72]. They also studied the performance of the FESTIVE algorithm, one of the first algorithms developed to explicitly deal with a multi-client scenario [73]. As far as the performance evaluation is concerned, fairness is computed as the standard deviation of clients' QoE. The QoE is a metric in the same range of the Mean Opinion Score and can be computed as described by Claeys et al. [71]. An NS-3-based simulation framework has been used to evaluate the developed multi-client framework. The simulated network topology is composed by three networks, each containing 30 HAS clients streaming video at the same time.

They considered MSS as reference client and computed the ratio between the average QoE of the analyzed client and that of MSS. Figure 9(a) reports the average value of this ratio, together with the confidence intervals at 95%. Figure 9(b) reports the average value of the ratio between the QoE standard deviation of the MSS algorithm and that of the analyzed client, together with the confidence intervals at 95%. The QoE standard deviation is used as fairness metric. The proposed solution is able to increase the average QoE by almost 20% for each of the three networks and to improve fairness with almost 80% when compared to MSS. Also the FESTIVE and Miller clients improve fairness, but consistently reduce the average QoE. This entails that the final QoE at the clients with these two heuristics is lower than that obtained using the MSS heuristic. These results show the sub-optimality of these two heuristics in case of a variable bandwidth, caused by frequent quality switches and video freezes. As far as the Q-Learning client is concerned, it improves fairness by about 50% with respect to MSS, but with a loss of 8–10% in terms of average QoE. This negative behavior is mainly due to the mutual influence among the learning processes of the clients and the uncoordinated nature of Q-Learning. When a client selects a certain quality level, it uses a portion of the shared bandwidth. This decision has an impact on the performance of the other clients and thus also on their learning process. Since the clients do not share any information, this leads to a sub-optimal quality adaptation policy.

How to Exchange Security Events? Overview and Evaluation of Formats and Protocols [2]

Network-based attacks pose a strong threat to the Internet landscape. Recent approaches to mitigate and resolve these threats focus on cooperation of Internet service providers and their exchange of security event information. A major benefit of a cooperation is that it might counteract a network-based attack at its root and provides the possibility to inform other cooperative partners about the occurrence of anomalous events as a proactive service.

This research provides a structured overview of existing exchange formats and protocols. The authors review the exchange formats and protocols used in context of intrusion detection and incident management. They analyze both the data representation and the use case scenario of the exchange formats. Further, they review existing exchange protocols and explain their intended use. As the authors identified the key position of ISPs in detection and mitigation of cyber-criminal activities, they develop various criteria to assess the exchange formats and protocols specifically in context of high-speed networks. Moreover, they assess the exchange formats for the use in conjunction with flow-based data, because a previous study from Steinberger et al. [74] stated that ISPs focus on detection of anomalous events based on aggregated network data (e.g. NetFlow, IPFIX).

The goal of this research is to provide network operators a hands-on selecting an exchange format and protocol suitable to use in their network. Therefore the main contributions of this paper are:

Table 7: Evaluation summary of the exchange formats

Criterion	CIDF	IODEF	CAIF	IDMEF	ARF	CEE	X-ARF		Syslog	
							v0.1	v0.2	RFC 3164	RFC 5425
Interoperability	−	−	−	−	+	+	+	+	+	+
Extensibility	+	+	+	+	+	+	+	+	+	+
Scalability	−	−	−	−	−	−	−	−	−	−
Aggregability	−	−	+	0	−	−	−	+	−	−
Protocol independency	−	0	+	0	+	0	+	+	+	+
Human readability	−	−	−	−	+	+	+	+	+	+
Machine readability	+	+	+	+	+	+	+	+	−	+
Integrity & Authenticity	−	−	−	−	−	−	−	+	−	−
Confidentiality	−	−	−	−	−	−	−	+	−	−
Practical application	−	0	0	0	0	−	0	0	+	+

Legend: high (+), medium (0) and low (−)

Table 8: Evaluation summary of the exchange protocols

Criterion	CIDF	RID	XEP-0268	IDXP	SMTP	CLT	Syslog	
							RFC 3164	RFC 5425
Confidentiality	+	+	+	+	−	+	−	+
Authenticity	+	+	+	+	−	+	−	+
Integrity	+	+	+	+	−	+	−	+
Reliable message transport	+	+	+	+	+	+	−	+
Interoperability	−	+	−	+	+	−	+	+
Scalability	+	−	+	+	+	+	+	+
Practical application	−	0	−	0	+	−	+	+

Legend: high (+), medium (0) and low (−)

(i) a comprehensive literature survey of 10 exchange formats and 7 exchange protocols that can be used to share security event related information in context of intrusion detection and incident handling, (ii) a structured overview that can be used by network operators when they have to decide what format and protocol should be used, (iii) an assessment of the exchange formats for the interoperability with flow- based data, (iv) a qualitative evaluation and comparison of the formats and protocols in context of high-speed networks and finally an investigation of how to exchange potentially sensitive information.

In the following, we report some of the key findings of this research. For the complete overview, we refer the reader to [32].

Table 7 lists the considered exchange formats and the identifed evaluations criteria, and it summarizes the overall evaluation. Similarly, Table 8 covers the exchange protocols.

The research highlights that the use of flow-based data within the XML-based exchange formats IODEF (Incident Object Description Exchange Format), CAIF (Common Announcement Interchange Format) and IDMEF (Intrusion Detection Message Exchange Format) requires a new XML scheme or the AdditionalData element needs to be used. The MIME-Message based formats transmit the same information multiple times without providing new knowledge.

Most of the exchange formats are machine readable. This ensures that security events can be handled automatically. In case the network operator focuses on machine readability, all exchange formats except syslog RFC 3164 are suitable. If a network operator focuses on an exchange format

that is human readable and does not require an additional parser, the network operator should focus on the MIME based exchange formats ARF and x-arf. But also CEE and syslog might be of interest, as they provide an free form message part. However, CEE has never been finalized, so actions should first be taken into this direction before the exchange format can be used in practice.

With respect to the interoperability with flow-based data, the exchange formats ARF (Abuse Reporting Format), CEE (Common Event Expression), x-arf (extended ARF) and syslog are suitable to use in conjunction with flow-based data. To exchange sensitive data, however, the network operator might focus on mechanisms to sign or encrypt a security event. Except the exchange format x-arf v0.2, none of the exchange formats provide mechanisms to sign or encrypt a security event. Finally, they note that to establish a collaboration between exchange peers, a well-known and established format should be used. Even though, a lot of exchange formats were published in the last years, only the exchange format syslog provides a widespread use.

To transmit a security event, the network operator might focus on a high-security level. The exchange protocols CIDF (Common Intrusion Detection Framework), RID (Real-time Inter-network Defense) , XEP-0268[13], IDXP (Intrusion Detection Exchange Protocol), CLT (CEE Log Transport) and syslog RFC 5425 provide a high-security level. SMTP and syslog RFC 3164 were not designed to ensure the four key aspects of information security (confidentiality, integrity, authenticity and non-repudiation). However, syslog and SMTP have the advantage that they are widely spread. Therefore they identify here a tradeoff between security level and easiness to deploy. Even though SMTP has never been updated, the use of the S/MIME standard provides the ability to digitally sign messages and to encrypt message contents to overcome these missing security aspects. In case a network operator focuses on an exchange protocol that should be used in high-speed networks, all exchange protocols are suitable except RID. RID does not scale in high-speed networks because it was designed as point-to-point protocol.

Cache Management for Telco-operated Content Delivery Network [4]

In previous work [75], [76], they proposed a proactive cache management system for Internet Service Provider (ISP)-operated multi-tenant Telco CDNs. In this approach [4], a central manager periodically computes a caching configuration based on predicted values of the future request pattern. This configuration covers the allocation of caching capacity across the network for the multiple tenants, the proactive placement of content and the server selection strategy (where to serve each request from). This problem was modeled as an Integer Linear Programming (ILP) problem with the objective of minimizing the bandwidth usage inside the ISP network.

The evaluation results highlighted that the performance of this proactive cache management approach strongly depends on the quality of the future request prediction. As such, more advanced prediction strategies have been investigated. However, analysis of the considered Video-on-Demand (VoD) request trace has shown that the predictability of future requests is strongly limited due to the high volatility of video content popularity and the large number of new content requested on a daily basis. Therefore, in the collaboration (**UCL-iMinds-Cache**), a hybrid cache management strategy is proposed as an extension of the proactive caching approach. In the hybrid approach, periodical proactive cache allocation and content placement is combined with distributed reactive cache replacement. In this way, content placement and server selection can be optimized across the network and tenants, based on predicted content popularity and the geographical distribution of requests, while simultaneously providing reactivity to unexpected changes in the request pattern. In addition, they have also adjusted the ILP model for the proactive placement in order to optimize the cost incurred by the placement update.

[13] specification of the Extensible Messaging and Presence Protocol (XMPP)

The proposed approach has been thoroughly evaluated in a simulated environment on a VoD use case, for which a request trace of the VoD service of a leading European telecom operator was used. Its performance has been compared to both a reactive caching approach, using the Least Recently Used (LRU) replacement, and the purely proactive approach, proposed in previous work. Evaluations have shown that in a realistic scenario, the hit ratio can be increased with 44% compared to a purely reactive approach. Furthermore, routing paths are shortened with 8% on average, resulting in an average reduction of 9% in terms of bandwidth usage. Compared to the proactive approach, the hit ratio is increased with 11% with 40% less migration overhead.

While the proposed ILP-based cache management approach enables the joint optimization of content placement and server selection, this may have scalability limitations as both the content catalogue and the caching infrastructure increase and as such, may not be suitable for short reconfiguration cycles.

To address these limitations, they also developed a novel scalable and efficient distributed approach to control the placement of content in the available caching points [4]. The proposed approach relies on parallelizing the decision-making process and the use of network partitioning to cluster the distributed decision-making points, which enables fast reconfiguration and limits the volume of information required to take reconfiguration decisions. More specifically, the network nodes are partitioned into independent clusters that could each execute an instance of the placement problem. Different clustering heuristics were investigated and a placement procedure was developed to determine the configuration of the different in-network caches.

The performance of the proposed approach was evaluated based on a wide range of parameters and the results show that network and cache performance similar to the ones obtained with sequential decision-making process can be achieved while significantly reducing management overhead and complexity. In particular, the number of iterations and the number of exchanged messages can be divided by up to a factor 10^4 and 10^6 when using network partitioning.

Mitigation of Topological Inconsistency Attacks in RPL based Low Power Lossy Networks [20]

The Routing Protocol for Low-power Lossy Networks (RPL) [77], designed for constrained devices and networks, is expected to find application in multiple areas of the Internet of Things (IoT). Being suitable for various fields like, Industrial Networks [78], Home and Building Automation [79] and Advanced Metering Infrastructure (AMI) Networks [80], it is evident that RPL will be exposed to multiple different operating scenarios, some of which will expose it to malicious attacks.

The RPL protocol The Routing Protocol for Low-power Lossy Networks (RPL) has been designed by the IETF [77] to address resource constraints of embedded devices. This protocol enables a distance-vector routing based on IPv6. RPL forms a loop-free tree like topology termed a Destination Oriented Directed Acyclic Graph (DODAG). A network can operate one or more RPL instances which consist of multiple DODAG graphs. When a loop occurs, RPL provides the *data path validation* mechanism to detect and repair rank related DODAG inconsistencies. This mechanism works by carrying the following flags in the RPL IPv6 header options [81] of multi-hop data packets:

- The *'O' flag* — indicates the expected direction of a packet. When set, the packet is intended for a descendant. Otherwise it is intended for a parent, towards the DODAG root.
- The *'R' flag* — indicates that a rank error was detected by a node forwarding the packet. A mismatch between the direction indicated by the 'O' flag and the rank of sending/forwarding node causes the flag to be set.

A DODAG inconsistency exists if the direction indicated by the ‘O’ flag does not match the rank relationship of the node from which the packet was received [77]. The ‘R’ flag is used to repair this problem by setting it, in case it was not set previously, and forwarding the packet. Upon receiving a packet with the ‘R’ flag already set an inconsistency is detected, the packet is discarded and the trickle timer used by RPL is reset [82]. This detection mechanism can be exploited by a malicious node to attack the network.

Attack description The data path validation can be misused either to harm a targeted node directly, or to manipulate packet headers and cause the next-hop node to drop the modified packet.

A malicious intruder can directly attack its neighborhood by sending packets that have the ‘R’ flag and the wrong direction set. For instance, if a parent is targeted, the attacker can send packets with the ‘O’ and ‘R’ flags set, since packets with ‘O’ flag are intended for descendant nodes. The parent will detect an inconsistency and thus, drop the packet and restart the trickle timer. This causes control messages to be sent more frequently which leads to local instability in the network. This increased control message overhead reduces channel availability and increases energy consumption which can lead to a shortened network lifetime in case nodes are battery operated. Since nodes in RPL networks are likely to be resource constrained, they are unlikely to support multitasking or large packet buffers. As such, time spent on processing malicious packets could lead to loss of genuine ones.

A malicious intruder can also modify the IPv6 header of packets it forwards such that the ‘R’ flag and the ‘O’ flag representing the wrong direction are set. The receiving node assumes that a DODAG inconsistency has taken place and discards the packet. As a result, the malicious node succeeds in forming a black-hole at the next-hop node. This attack could either be carried out on all packets forwarded by the malicious node, or selectively based on source, destination, or even type of message. In general this approach is a good strategy for the attacker to force another node to drop the packets. Furthermore, if the control packets originating from the malicious node are normal, then the malicious activity is completely hidden. In this scenario, not only does the delivery ratio decrease, but the control overhead of RPL nodes also increases along with deteriorating channel availability and increasing energy consumption.

Results of attack mitigation The data path validation can be misused either to harm a targeted node directly, or to manipulate packet headers and cause the next-hop node to drop the modified packet.

A malicious intruder can directly attack its neighborhood by sending packets that have the ‘R’ flag and the wrong direction set. For instance, if a parent is targeted, the attacker can send packets with the ‘O’ and ‘R’ flags set, since packets with ‘O’ flag are intended for descendant nodes. The parent will detect an inconsistency and thus, drop the packet and restart the trickle timer. This causes control messages to be sent more frequently which leads to local instability in the network. This increased control message overhead reduces channel availability and increases energy consumption which can lead to a shortened network lifetime in case nodes are battery operated. Since nodes in RPL networks are likely to be resource constrained, they are unlikely to support multitasking or large packet buffers. As such, time spent on processing malicious packets could lead to loss of genuine ones.

A malicious intruder can also modify the IPv6 header of packets it forwards such that the ‘R’ flag and the ‘O’ flag representing the wrong direction are set. The receiving node assumes that a DODAG inconsistency has taken place and discards the packet. As a result, the malicious node succeeds in forming a black-hole at the next-hop node. This attack could either be carried out on

all packets forwarded by the malicious node, or selectively based on source, destination, or even type of message. In general this approach is a good strategy for the attacker to force another node to drop the packets. Furthermore, if the control packets originating from the malicious node are normal, then the malicious activity is completely hidden. In this scenario, not only does the delivery ratio decrease, but the control overhead of RPL nodes also increases along with deteriorating channel availability and increasing energy consumption.

Attack mitigation The default DODAG inconsistency attack mitigation strategy of RPL consists in a fixed threshold. Upon receiving a packet with an inconsistency, the node drops it and resets its own trickle timer. To limit the effects of an attack, the number of trickle timer resets is limited to the recommended constant 20 [81]. Upon reaching this threshold, malformed packets are dropped but the trickle timer is not reset. The counter used is reset every hour, allowing attackers to once again have a higher impact. This approach limits the impact of a DODAG inconsistency attack, but the value of the threshold is arbitrarily set. No reasoning is provided to justify this choice or how performance could be improved in case of varying attack scenarios. Also since the packets are still dropped this approach does not mitigate the indirect attack scenario.

In order to take into account the current network state and react to varying attack patterns they developed an adaptive threshold (AT) [83], which determines when to stop resetting the trickle timer. Instead of a constant, a decreasing exponential function is used with fixed parameters. The adaptive threshold causes the threshold to change based on network conditions. If an attacker is aggressive, the threshold drops quickly and increases slowly once the attacks stop. Unlike with the fixed threshold, the counter of 'R' flag packets is not reset every hour, but rather allowed to increase in the absence of attacks. As such, not only is this approach likely to be better than a fixed threshold within the first hour of an attack, but it should perform significantly better against long running attacks. To counter the packet manipulation DODAG inconsistency attack, an extension was made to the adaptive threshold. Nodes behave normally until the number of messages indicating an inconsistency becomes greater than the threshold obtained from the function. This situation indicates either an attack against the node, or malfunction of the node forwarding such packets. To rectify the situation, the node clears the 'O' and 'R' flags before forwarding the packets normally.

The adaptive threshold approach relies on set parameters, which a particular RPL implementation needs to choose. This can lead to sub-optimal optimizations and so they have improved the presented mitigation approach via the design of a fully dynamic threshold, which is based on network characteristics. The new threshold used to determine whether the trickle timer should be reset is similar to the previous one (decreasing exponential function). However, the parameters for this function are based on node specific characteristics (number of neighbors). It is possible for multiple packets with an 'R' flag to arrive as a result of the same inconsistency. Resetting the trickle timer each time a malfunctioning node sends packets with 'R' flags leads to unnecessary overhead, especially since a single trickle timer causes aggressive transmissions of DIOs anyway. To avoid this situation, a convergence timer is introduced. This timer is used to ensure that no further trickle timer resets take place within the amount of time it takes for an RPL neighborhood to typically converge. A new counter that keeps track of the number of trickle timer resets is introduced and compared with the calculated threshold to determine when the trickle timer should be reset. Like the adaptive threshold approach, this mitigation strategy should perform better against long running attacks. This dynamic threshold approach not only does away with arbitrary constant thresholds, as in the case of the default strategy, but by being based purely upon network characteristics it does away with the need for constant parameters to be chosen before deployment [83] and thereby is more useful in case of unforeseen network conditions as well. Also the solution was

adapted to counter the packet manipulation scenario by allowing a node to forward packets with the inappropriate flags under certain conditions.

They showed through several experiments that the presented strategies were efficient to mitigate both scenarios of DODAG inconsistency attack without having a significant cost on the deployed nodes.

Robust Geometric Forest Routing with Tunable Load Balancing [53]

Geometric routing schemes are proposed as an alternative for lookup-based routing algorithms. Although they were initially designed for Unit Disk Graphs (UDGs) [84], their application to scale-free complex networks has been demonstrated [85]. This form of routing makes use of a *graph embedding*, the assignment of coordinates in a mathematical space to every network vertex. This embedding, together with an appropriate distance function, forms the core of geometric routing, allowing packets to be transmitted along a distance-decreasing path towards their destination.

The main advantage of geometric routing is its low state complexity. A node only requires information about its neighbors, rather than being dependent on the state of the whole network. In contrast to more traditional routing schemes based on lookup tables, geometric routing thus restricts the required router memory overhead. A large disadvantage, however, is their lack of load balancing characteristics, which is essential in avoiding traffic congestion in large-scale networks. Lookup-based schemes can easily add this by incorporating multiple alternative routes in their lookup tables. How load balancing characteristics and stretch can be combined and traded off in geometric routing is still an open research question.

In this work [53] they explore the possibilities in combining low stretch with load balancing behavior. The main contribution is a family of routing schemes called Forest Routing (FR). These algorithms are capable of adapting their routing behavior to varying traffic intensities by using a generalized distance function incorporating link load information. Additionally, in the absence of network failures they attain a 100% success ratio, while having a high resiliency to node and link failures. Although designed for, Forest Routing is not restricted to complex scale-free networks.

Methodology The FR algorithms employ a spanning tree $T = (V, E')$ of the underlying graph $G = (V, E)$ to construct an embedding by making use of the vertex labeling procedure described by Korman et al. [86] and a metric representing the shortest path length in T, in number of hops, as described by Chávez et al. [87]. The embedding target space S is denoted as the *tree space* $\mathbb{T}$, defined as

$$\mathbb{T} = \bigcup_{n \in \mathbb{N}} \Big((0)^\frown \mathbb{N}^n \Big), \tag{1}$$

in which the function $(^\frown) : \mathbb{N}^m \times \mathbb{N}^n \to \mathbb{N}^{m+n}$ represents the concatenation of two tuples. Now, one can interpret the assigned labels as coordinates in the space $\mathbb{T}$. They say that these coordinates form a greedy graph embedding [87], denoted as $\mathcal{T}$. As such, each vertex $v \in V$ corresponds to a point in $\mathbb{T}$ identified by the coordinates $\mathcal{T}(v)$.

The distance function δ for $\mathbb{T}$ is defined as

$$\delta(u, v) = |u| + |v| - 2|\phi(u, v)|, \tag{2}$$

with $\phi : \mathbb{N}^n \times \mathbb{N}^m \to \mathbb{N}^*$ a function that returns the largest common prefix of two tuples; $|u|$ is the length of the coordinate tuple of vertex u. This leads to a distance function $\delta : \mathbb{T} \times \mathbb{T} \to \mathbb{R}^+$ that, combined with the space $\mathbb{T}$, forms the metric space $(\mathbb{T}, \delta)$. Now, this metric space can be

used according to the principles of geometric routing. This means that every vertex knows the coordinates of its neighbors and the coordinates of the target vertex are encoded in each packet header. As such, using this header, each node forwards packets along a distance-decreasing path to their destination.

A straightforward way of routing with multiple embeddings is to allow each vertex to alternate freely between them, making use of their individual greediness. However, this naive forwarding mechanism is unreliable because it can introduce routing cycles. Routing along a distance-decreasing path in $\mathcal{T}_i$ may increase the distance in a different embedding $\mathcal{T}_j$. At a certain vertex along a packet's routing path, it may be sent back to its origin, resulting in a routing cycle. A cycle avoidance solution requires that each vertex along the routing path decreases the packet's minimum distance (over the k embeddings) to the destination. This way of working is similar to the TCGR mechanism [88]. For this reason a new distance function $\epsilon : \mathbb{T}^k \times \mathbb{T}^k \to \mathbb{R}^+$ is defined as

$$\epsilon(u, v) = \min_{0 \leq i < k} \{\delta_i(u, v)\} \ \forall u, v \in V, \tag{3}$$

which replaces the original distance function δ [89]. The k embeddings into $\mathbb{T}$ can now be treated as a single k-dimensional embedding into $\mathbb{T}^k$. This allows us to adhere to the principles of geometric routing by using a semimetric space $(\mathbb{T}^k, \epsilon)$. As such, each node u will forward packets to the neighbor with the lowest distance $\epsilon(u, d)$, with d the destination node. Therefore, it is a form of *greedy routing*, which is geometric routing in which a node always forwards to the neighbor leading to the largest decrease in distance. In case multiple neighbors have an equal ϵ-distance, a random choice is made among them.

To supplement the passive load balancing behavior emerging from GFR, an active load balancing approach was developed called *Load Balanced Forest Routing* (LBFR). This system can be seen as a special case of the final HFR routing scheme. In LBFR, vertices $u \in V$ make use of traffic load information about their incident edges $e \in I(u)$. This information is used to select the neighbor v for which the edge (u, v) has the lowest load. Solely using local link information is advantageous as it is scalable by nature and therefore fitting for a large-scale distributed setting. LBFR relaxes the greedy requirements of GFR by allowing routing alternately via different embeddings $\mathcal{T}_i$ independently. Because naive switching between embeddings may introduce cycles, routing is guided by an auxiliary function κ. This function acts as a routing restriction by requiring that its value decreases at each hop, similarly to how the δ-distance must decrease in standard geometric routing (or ϵ in GFR). All neighbors fulfilling this requirement are added to a set $S(u)$, the set of nodes that can be considered as next hops.

This function κ makes use of an additional function δ^* that outputs a k-tuple storing the minimal distance to the destination d attained by a packet so far along its routing path P_u, before arriving at the current node u, for each of the k embeddings. They denote the i-th element of δ^* as δ_i^* and the union of all possible paths in the network as $\mathcal{P}$. Assuming a packet has been routed along a path $P = \langle p_0, p_1, \ldots, p_n \rangle$ towards a destination vertex d, then κ is of the type $\mathcal{P} \times \mathbb{T}^k \to \mathbb{N}$ and is defined as

$$\kappa(P_{p_n}, d) = \sum_{i=0}^{k-1} \delta_i^* \left(P_{p_n}, d\right), \tag{4}$$

with $\delta^*(P_{p_n}, d)$ a function of the type $\mathcal{P} \times \mathbb{T}^k \to \mathbb{N}^k$ that is defined recursively $\forall i \in \{0, \ldots, k-1\}$ as

$$\delta_i^*(P_{p_0}, d) = \delta_i(p_0, d) \tag{5}$$

$$\forall n > 0 : \delta_i^*(P_{p_n}, d) = \min\{\delta_i^*(P_{p_{n-1}}, d), \delta_i(p_n, d)\}. \tag{6}$$

Herein P_u represents the path P until u has been reached, consisting of the vertices that a packet arriving at u has reached. Furthermore, p_0 is the source vertex of the path P. The minimum distances of each of the k embeddings is thus represented by an element $\delta_i^*(P_u, d)$. The LBFR system will enforce the restriction that κ has to decrease strictly monotonically along the routing path: $\kappa(P_{p_n}, d) < \kappa(P_{p_{n-1}}, d) < \ldots < \kappa(P_{p_0}, d)$. When forwarding, a node u will select those neighboring nodes which have a strictly decreasing κ-value and add them to the previously mentioned set $S(u)$. Next, u will select a vertex $v \in S(u)$ as the next hop for which the current traffic load of the link (u, v) is minimal compared to its other incident links $I(u)$.

The following three theorems prove its robustness:

Theorem 1. *Let $G = (V, E)$ be a graph with k embeddings $\mathcal{T}_i$ for $0 \leq i < k$ into the metric space $(\mathbb{T}, \delta)$. Let d be the destination node. Then, for every path $P \in \mathcal{P}$ (in G) with a last element $v \in V$, for which d has not yet been reached, thus $d \notin P$, the set of neighbors $S(v)$ for which the value of the κ-function strictly decreases is not empty.*

Theorem 2. *The path followed by a packet routed on a graph $G = (V, E)$ by LBFR is never a cycle.*

Theorem 3. *A packet routed according to the principles of LBFR on a graph $G = (V, E)$ will arrive at its destination.*

In terms of stretch and load balancing, GFR and LBFR are two opposites: GFR attains low stretch, but has no load balancing technique, while LBFR achieves load balancing, but pays no attention to stretch. They combine the best of both worlds into one algorithm called *Hybrid Forest Routing* (HFR). HFR makes a trade-off between stretch and load balancing by replacing the GFR distance function ϵ by a cost function that combines link load information with the ϵ-distance to the destination. This cost function $C : V^3 \to \mathbb{R}^+$ is defined as

$$C(u, n, d) = \gamma \cdot \hat{L}(u, n) + (1 - \gamma) \cdot \epsilon(n, d) \tag{7}$$

for $n \in N(u)$, with the ϵ-function defined by Eq. (3). The function $\hat{L}(u, v)$ represents the normalized traffic load of the edge between u and v. This is the traffic load of link (u, v) divided by the average load of all the node's incident links $I(u)$. This normalized load is defined as

$$\hat{L}(u, n) = d_G(u) \cdot L(u, n) \Big/ \sum_{v \in N(u)} L(u, v), \tag{8}$$

with $d_G(u)$ the degree of vertex u. The factor $\gamma \in [0, 1]$ is a weight factor which allows scaling between greedy and load balanced routing. As can be seen, HFR also uses the semimetric space $(\mathbb{T}^k, \epsilon)$, but because the cost function is an extension of the ϵ-function, HFR is not greedy routing. Even more, it is not necessarily distance-decreasing in ϵ.

To guarantee packet delivery, the κ-function from LBFR is used to steer packets towards their destination, relying on the LBFR theorems from the previous section. HFR attains strong load balancing while keeping the stretch down, which is shown in the following section. GFR and LBFR can now be seen as two special instances of HFR on opposite sides of the spectrum. On the one hand, when $\gamma = 1$, the LBFR mechanism is recreated. On the other hand, when $\gamma = 0$, the HFR reverts to GFR.

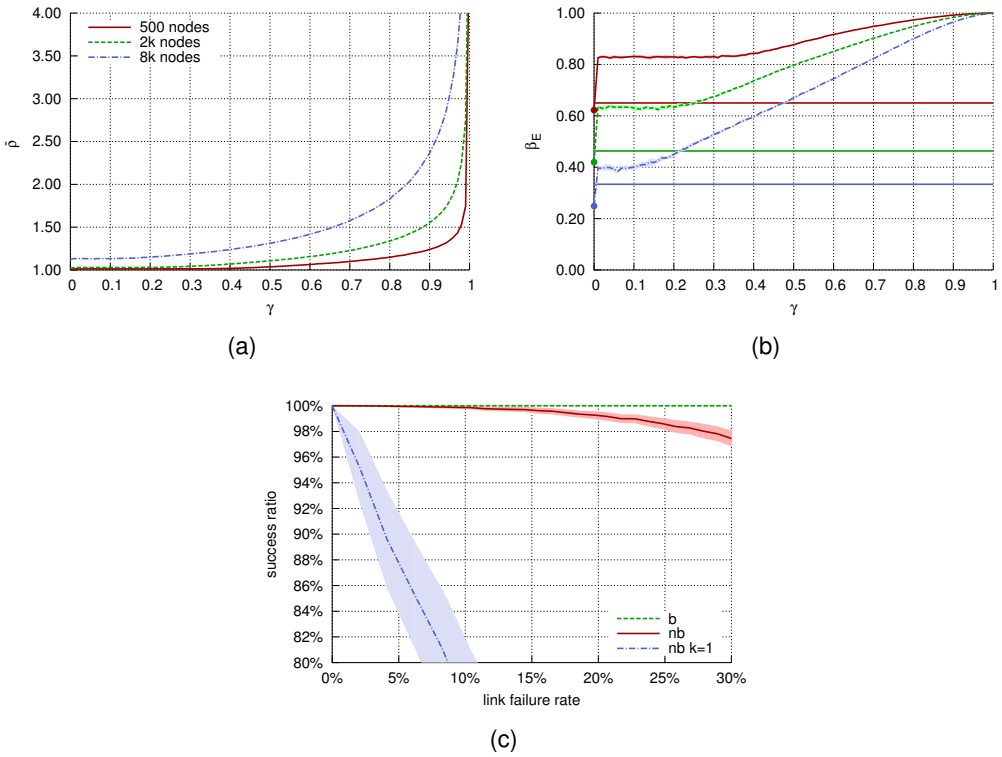

Figure 10: **Figures 10(a) and 10(b)**: HFR: average stretch $\bar{\rho}$ (a) and link load balancing metric β_E (b) in function of a varying γ-value. The solid horizontal lines in the bottom figure represent the load balancing β_E-value of shortest path routing with lookup tables. The shaded background represents the average, plus and minus the standard deviation. **Figure 10(c)**: Fault-tolerance: the horizontal axis depicts the fraction of links removed from the total number of links that can be removed ($|E| - |V| + 1$) without disconnecting the graph $G = (V, E)$. HFR ($\gamma = 0.1$, $k = 15$) with (b) and without (nb) backup mechanism is tested, along with a single tree-based geometric routing algorithm (RTP [87]). The shaded background represents the average success ratio, plus and minus the standard deviation.

Results and discussion HFR, which unites both GFR and LBFR, is evaluated. When the parameter γ in Eq. (7) is shifted towards 0, the GFR system is recreated, while shifting γ to 1 results in LBFR. Therefore, in Figures 10(a) and 10(b) a sensitivity analysis of the γ-parameter is conducted for $k = 15$. For each value of γ, 10 runs consisting of 10^5 generated routing paths are executed. Figure 10(a) shows that at low γ-values $\bar{\rho}$ (1: shortest path routing; goes to $+\infty$ as routing paths increase in length) becomes equally low as with GFR. Afterwards, $\bar{\rho}$ starts to incline as $\gamma \to 1$, consistent with HFR approximating LBFR.

The β_E-values (1: perfect load balancing; 0: no load balancing) in Figure 10(b) indicate that shifting γ between its two extremes gives the expected results regarding load balancing. However, something interesting can be noticed at the right side of $\gamma = 0$. A step occurs such that β_E suddenly rises. Although, when inspecting Figure 10(a) one can see no such step in the $\bar{\rho}$-curve. This can be explained as follows. Upon a node's forwarding decision, many potential candidates will have an equal distance to the destination. So within this set it does not matter which node to forward to in terms of stretch. When taking into account load balancing, a huge improvement can be made by prioritizing those links with a low current load. Figure 10(b) also shows the load balancing behavior

of routing with lookup-tables for the different scale-free graphs, which is depicted by the horizontal solid lines. It can be noticed that HFR offers much better load balancing with only a minor increase in stretch.

In Figure 10(c) the routing success rate of the HFR system with $\gamma = 0.1$ and $k = 15$ is shown, alongside HFR combined with the backup routing system Gravity-Pressure (GP) routing [90], and the single tree-based geometric routing algorithm RTP [87], exercised on a scale-free graph with 500 nodes. Links are removed probabilistically such that link failures are spread out evenly over the network to avoid random failure concentration in a certain area. This shows that HFR can be easily equipped with a backup routing mechanism, allowing it to achieve a 100% success rate even in severe failure scenarios. However, even without the GP mechanism, HFR attains a success rate of over 97% when 30% of the removable links are deleted. This is a huge improvement over more basic tree-based geometric routing algorithms based on a single tree where the success ratio quickly declines as the number of failed links increases (a success ratio of less than 50% at a 30% link failure rate).

Conclusion In this work a theoretical framework is built which serves as a foundation for the developed family of geometric routing systems, called Forest Routing (FR). Combining a strictly greedy approach, Greedy Forest Routing (GFR), with a load balanced routing scheme, Load Balanced Forest Routing (LBFR), results in Hybrid Forest Routing (HFR). In HFR, path stretch can be traded against load balancing behavior, two features until now not perceived to be compatible. Due to its local routing decision making procedure it is highly scalable regarding router memory requirements, making it robust towards network growth. Furthermore, the HFR system has favorable characteristics such as inherent fault-tolerance and guaranteed packet delivery. It can deal even with a highly deteriorated network topology, and is as such able to guarantee success ratios as high as 97% at link failure rates of 30%.

5 Conclusions and Outlook

Deliverable 6.3 describes the achievements of WP6 in the third year with respect to S.M.A.R.T. as well as work package specific objectives. This deliverable documents also the full achievement of these objectives. For more details about the achievement of the S.M.A.R.T. objectives, the reader is referred to Deliverable D8.3; the achievements of work package specific objectives in the third year have been reported in this deliverable in Section 2.2.

The work package specific objectives in the third year have been centered around the two main pillars of WP6: security use case and content delivery use case. Based on these two use cases and the developed generic approaches an essential step towards the FLAMINGO integrative architecture for automated configuration and repair was done.

Nonetheless, the achievements of the last two years of FLAMINGO encouraged WP6 to build the integrative architecture inline with the applicability guidelines and inventory of enablers and architectures in the area of automated configuration and repair.

A strong integration of PhD students (both, fully payed and not fully payed by FLAMINGO) and PhD collaborations is the basis for the immense and excellent scientific output achieved in Y3.

In this year the research work packages published 73 papers at major conferences as well as in journals, and exceeds the expected number of papers. 40 of these publications are strongly related to WP6.

The close and still ongoing PhD collaborations, especially between WP5 and WP6, are an enabler for excellent research results in the next year as well.

To recall our highlights in Y3 cover dynamic and adaptive resource management in SDN, overview and evaluation of formats and protocols for the exchange of security events and robust geometric forest routing with tunable load balancing, to mention only a few examples.

Y4 of FLAMINGO will focus on further work towards the FLAMINGO integrative architecture, including the strong use cases in security and content delivery. Furthermore, WP6 will work on generic approaches usable in various use cases. The extreme, continuous success of PhD collaborations also in the fourth year, is a guarantee to obtain excellent research results in the next year as well, and to foster the joint PhD collaborations.

Acknowledgments

This deliverable is based on input from the WP6 Partners of the FLAMINGO consortium. A particular acknowledgment goes to all the PhD students that have not only provided textual input, but that are working on a daily basis on the challenging research topics that we report.

6 Abbreviations

$3GPP$	3rd Generation Partnership Project
$6LoWPAN$	IPv6 over Low Power, Wireless Networks
ACE	Autonomic Communication Element
ACR	Automated Configuration and Repair
AE	Autonomic Element
AME	Autonomic Management Entities
ANA	Autonomic Network Architecture
$ANEMA$	Autonomic Network Management Architecture
ANM	Autonomic Network Management
ANN	Artificial Neural Networks
API	Application Programming Interface
AQM	Active Queue Management
AS	Autonomous System
AWS	Amazon Web Services
BP	Back-Propagation
$CASCADAS$	Component-ware for Autonomic, Situation-aware Communications and Dynamically Adaptable Services
CDN	Content Delivery Network
CLI	Command Line Interface
CPS	Cyber Physical Systems
CVE	Common Vulnerabilities and Exposures language
$CVSS$	Common Vulnerability Scoring System
$DACoRM$	Decentralised and Adaptive Network Resource Management Framework
$DASH$	Dynamic Adaptive Streaming over HTTP
DCE	Direct Code Execution
$DDoS$	Distributed Denial of Service attack
$DNS-SD$	DNS Based Service Discovery
$DODAG$	Destination Oriented Directed Acyclic Graph
DoW	Description of Work
DRA	Dynamic Resource Allocation
$ECMP$	Equal-Cost Multi-Path
EF	Expedited Forwarding
$EMANICS$	European Network of Excellence for the Management of Internet Technologies and Complex Services
ESB	Enterprise Service Bus
EU	European Union
FI	Future Internet
FN	False Negative
$FOCALE$	Foundation, Observation, Comparison, Action, Learning, rEason
FP	False Positive

GPS	Global Positioning System
HAS	HTTP Adaptive Streaming
$HTTP$	Hyper-text Transfer Protocol
$ICMP$	Internet Control Message Protocol
IDS	Intrusion Detection System
$IDMEF$	Intrusion Detection Message Exchange Format
$IETF$	Internet Engineering Task Force
ILP	Integer Linear Program
InP	Infrastructure Provider
$INRIA$	Institut National de Recherche en Informatique et Automatique
IoT	Internet of Things
IP	Internet Protocol
$IPFIX$	Internet Protocol Flow Information Export
ISP	Internet Service Provider
$ITU-T$	International Telecommunications Union - Telecommunications Standardization Sector
JNI	Java Native Interfaces
JUB	Jacobs University Bremen
LLN	Low-power and Lossy Networks
LRU	Least Recently Used
MAS	Multi-Agent System
$mDNS$	multicast DNS
MDP	Markovian Decision Processes
$MITM$	Man-In-The-Middle
MNO	Mobile Network Operator
MOS	Mean Opinion Score
$MP2P$	Multipoint-to-Point
$MPEG$	Moving Picture Experts Group
MTR	Multi-Topology Routing
MSS	Microsoft ISS Smooth Streaming
MTU	Maximum transmission unit
$NETCONF$	Network Configuration Protocol
$NFQL$	Network Flow Query Language
NFV	Network Functions Virtualization
$NIST$	National Institute of Standards and Technology
NSC	Network Simulation Cradle
$OSPF$	Open Shortest Path First
OTT	Over-the-Top
$OVAL$	Open Vulnerability and Assessment Language
OWL	Web Ontology Language
PHB	Per Hop Behavior
QoE	Quality-of-Experience
QoS	Quality-of-Service
RED	Random Early Detection
$RDBMS$	Relational Database Management System
$RMSE$	Root Mean Square Error
$ROLL$	Routing Over Low Power Lossy networks
RPL	Routing Protocol for Low power and Lossy Networks
$SACK$	Selective Acknowledgments
$SCAP$	Security Content Automation Protocol

SDN	Software-defined networking
SLA	Service Level Agreement
$S.M.A.R.T.$	Specific Measurable Achievable Relevant Timely
SN	Substrate Network
$SNMP$	Simple Network Management Protocol
$SPARQL$	SPARQL Protocol and RDF Query Language
SOA	Service-oriented architecture
SP	Service Provider
SSH	Secure Shell
$SWRL$	Semantic Web Rule Language
$P2P$	Peer-to-Peer
$P2MP$	Point-to-Multipoint
RL	Reinforcement Learning
TCP	Transmission Control Protocol
TD	Time Difference
TN	True Negative
$TNSM$	Transactions on Network and Service Management
TP	True Positive
TPM	Trusted Platform Module
TSP	Travelling Salesman Problem
$UniBwM$	Universität der Bundeswehr München
UCL	University College London
UDP	User Datagram Protocol
UPC	Universitat Politecnica de Catalunya
UT	University of Twente
UZH	University of Zürich
$VDBE$	Value-Difference Based Exploration
VoD	Video-on-Demand
VoS	Value of Service
VM	Virtual Machine
VNP	Virtual Network Provider
VPN	Virtual Private Networks
$WLAN$	Wireless Local Area Network
WP	Work Package
WSN	Wireless Sensor Network
XML	Extensible Markup Language

References

[1] R. Mijumbi, J. Serrat, J. Gorricho, N. Bouten, F. De Turck, and R. Boutaba. Network function virtualization: State-of-the-art and research challenges. *IEEE Communications Surveys & Tutorials*, Nov 2015.

[2] J. Steinberger, A. Sperotto, M. Golling, and H. Baier. How to exchange security events? Overview and evaluation of formats and protocols. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 261–269. IEEE, 2015.

[3] Niels Bouten, Jeroen Famaey, Rashid Mijumbi, Bram Naudts, Joan Serrat, Steven Latré, and Filip De Turck. Towards nfv-based multimedia delivery. *IEEE International Symposium on Integrated Network Management(IM)*, May 2015.

[4] Daphne Tuncer, Vasilis Sourlas, Marinos Charalambides, Maxim Claeys, Jerome Famaey, George Pavlou, and Filip De Turck. Scalable cache management for isp-operated content delivery services. *Under submission to IEEE JSAC Special Issue on Video Distrubution over Future Internet*, 2015.

[5] Daphne Tuncer, Marinos Charalambides, Stuart Clayman, and George Pavlou. Adaptive resource management and control in software defined networks. *Network and Service Management, IEEE Transactions on*, 12(1):18–33, 2015.

[6] Rashid Mijumbi, Joan Serrat, Juan-Luis Gorricho, and Raouf Boutaba. A path generation approach to embedding of virtual networks. *Network and Service Management, IEEE Transactions on*, 12(3):334–348, 2015.

[7] Rashid Mijumbi, Joan Serrat, Juan-Luis Gorricho, Niels Bouten, Filip De Turck, and Steven Davy. Design and evaluation of algorithms for mapping and scheduling of virtual network functions. *IEEE Conference on Network Softwarization (NetSoft). University College London*, April 2015.

[8] Rashid Mijumbi, Joan Serrat, and Juan-Luis Gorricho. Self-managed resources in network virtualization. *IFIP/IEEE International Symposium on Integrated Network Management (IM)*, 2015.

[9] Rashid Mijumbi, Juan-Luis Gorricho, Joan Serrat, Javier Rubio-Loyola, and Ramon Aguero. Survivability-oriented negotiation algorithms for multi-domain virtual networks. In *Network and Service Management (CNSM), 2014 10th International Conference on*, pages 276–279. IEEE, 2014.

[10] Rashid Mijumbi, Joan Serrat, Javier Rubio-Loyola, Niels Bouten, Filip De Turck, and Steven Latré. Dynamic resource management in sdn-based virtualized networks. In *Network and Service Management (CNSM), 2014 10th International Conference on*, pages 412–417. IEEE, 2014.

[11] S. Seeber, L. Stiemert, and G. D. Rodosek. Towards an SDN-Enabled IDS Environment. In *Communications and Network Security (CNS), 2015 3th International Conference on*. IEEE, 2015.

[12] Rashid Mijumbi, Joan Serrat, Juan-Luis Gorricho, Niels Bouten, Filip De Turck, and Raouf Boutaba. Network function virtualization: State-of-the-art and research challenges. *IEEE Communications Surveys and Tutorials*, 2015.

[13] Maxim Claeys, Niels Bouten, Danny De Vleeschauwer, Werner Van Leekwijck, Seven Latre, and Filip De Turck. An announcement-based caching approach for video-on-demand streaming. In *Network and Service Management (CNSM), 2015 11th International Conference on*. IEEE, 2015.

[14] Gaëtan Hurel, Remi Badonnel, Abdelkader Lahmadi, and Olivier Festor. Behavioral and Dynamic Security Functions Chaining for Android Devices. In *Proceedings of the 10th International Conference on Network and Service Management, CNSM 2015, Barcelona, Spain, November 9-13, 2015*, 2015.

[15] Niels Bouten, Ricardo de O Schmidt, Jeroen Famaey, Steven Latré, Aiko Pras, and Filip De Turck. Qoe-driven in-network optimization for adaptive video streaming based on packet sampling measurements. *Computer networks*, 81:96–115, 2015.

[16] S. Petrangeli, J. Famaey, M. Claeys, S. Latre, and Filip De Turck. Qoe-driven rate adaptation heuristic for fair adaptive video streaming. In *ACM Transactions on Multimedia Computing, Communications and Applications (ACM TOMM)*. IEEE, 2015.

[17] S. Seeber and G.D. Rodosek. Towards an Adaptive and Effective IDS Using OpenFlow. In *Proc of the 9th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2015)*, pages 134–139. Springer, 2015.

[18] Stefano Petrangeli, Tim Wauters, Rafael Huysegems, Tom Bostoen, and Filip De Turck. Network-based dynamic prioritization of http adaptive streams to avoid video freezes. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 1242–1248. IEEE, 2015.

[19] Stefano Petrangeli, Niels Bouten, Maxim Claeys, and Filip De Turck. Towards svc-based adaptive streaming in information centric networks. In *In proceedings of the Workshop on Multimedia Streaming in Information-Centric Networks (MuSIC)*, 2015.

[20] Anthéa Mayzaud, Anuj Sehgal, Rémi Badonnel, Isabelle Chrisment, and Jürgen Schönwälder. Mitigation of topological inconsistency attacks in rpl-based low-power lossy networks. *International Journal of Network Management*, 2015.

[21] Dennis Kergl, Robert Roedler, and Gabi Dreo Rodosek. Detection of zero day exploits using real-time social media streams. In *Computational Aspects of Social Networks, 7th International Conference on*. Advances in Intelligent and Soft Computing, 2015.

[22] Robert Roedler, Dennis Kergl, and Gabi Dreo Rodosek. Profile matching across online social networks based on geo-tags. In *Computational Aspects of Social Networks, 7th International Conference on*, 2015.

[23] Dennis Kergl. Enhancing network security by software vulnerability detection using social media analysis extended abstract. In *Data Mining (ICDM), 2015 IEEE 15th international conference on*. IEEE, 2015.

[24] Rashid Mijumbi, J Serrat, JL Gorricho, Maxim Claeys, Filip De Turck, and Steven Latré. Design and evaluation of learning algorithms for dynamic resource management in virtual networks. In *14th IEEE/IFIP Network Operations and Management Symposium (NOMS 2014)*, 2014.

[25] Rashid Mijumbi, Joan Serrat, Juan-Luis Gorricho, Javier Rubio-Loyola, Steven Davy. Server placement and assignment in virtualized radio access networks. In *Conference on Network and Service Management (CNSM)*, November 2015.

[26] R. Koch, M. Golling, F. Tietze, S.D. Hein, M. Kretzschmar, and G.D. Rodosek. An Agent-based Framework for a Decentralized Reconstruction of Attack Paths. In *DFN-Forum Kommunikationstechnologien*. GI, 2015.

[27] M. Golling, R. Koch, L. Stiemert, F. Tietze, V. Eiseler, and G.D. Rodosek. A Decentralized Framework for Geolocation-based Pre-Incident Network Forensics. In *7th International Symposium on Cyberspace Safety and Security (CSS 2015)*. IEEE, 2015.

[28] R. Koch, M. Golling, L. Stiemert, and G.D. Rodosek. Using Geolocation for the Strategic Preincident Preparation of an IT Forensics Analysis. *Systems Journal, IEEE*, To be published, 2015.

[29] Frank Tietze, Peter Hillmann, and Gabi Dreo Rodosek. Strategie zur Verfolgung einzelner IP-Pakete zur Datenflussanalyse. In *8. DFN-Forum Kommunikationstechnologien, Beiträge der Fachtagung, 08.-09. Juni 2015, Lübeck*, volume 243 of *LNI*, pages 47–56. Gesellschaft für Informatik (GI), GI, 2015. Nominated for Best Paper Award.

[30] P. Hillmann, R. Tietze, and G. Dreo Rodosek. Strategies for Tracking Individual IP Packets Towards DDoS. In *PIK Magazine*, 2015.

[31] R. Koch and M. Golling. Blackout and Now? - Network Centric Warfare in an Anti-Access Area-Denial Theatre. In *7th International Conference on Cyber Conflict (CyCon)*. IEEE, 2015.

[32] J. Steinberger, A. Sperotto, H. Baier, and A. Pras. Collaborative attack mitigation and response: A survey. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 910–913, May 2015.

[33] C. Schmitt, M. Noack, W. Hu, T Kothmayr, and B. Stiller. Two-way authentication for the internet-of-things. In *Securing the Internet of Things through Progressive Threat Detection and Management, H. Alzaid, B. Alomair, S. Almotiri, N. Nasser (Edts.), Book Series on Advances in Information Security, Privacy, and Ethics (AISPE), IGI Global*, 2015.

[34] Gaëtan Hurel, Rémi Badonnel, Abdelkader Lahmadi, and Olivier Festor. Towards cloud-based compositions of security functions for mobile devices. In *IFIP/IEEE International Symposium on Integrated Network Management (IM'15)*, page 6, 2015.

[35] P. Porambage, A. Braeken, C. Schmitt, A. Gurtov, B. Ylianttila, and B. Stiller. Group key establishment for source multicasting in iot-enabled wireless sensor networks. In *40th IEEE Conference on Local Computer Networks (LCN 2015)*, 2015.

[36] C. Schmitt and B. Stiller. Secure and Efficient Wireless Sensor Networks. *ERCIM News - Special Issue: The Internet of Things and The Web of Things*, 2015(101):18–19, apr 2015.

[37] J. Steinberger, A. Sperotto, H. Baier, and A. Pras. Exchanging Security Events of flow-based Intrusion Detection Systems at Internet Scale. In *Coordinating Attack Response at Internet Scale (CARIS) Workshop*, June 2015.

[38] M. Jonker and A. Sperotto. Mitigating DDoS Attacks Using OpenFlow-Based Software Defined Networking. In *Proc. of the 9th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2015)*, pages 129–133. Springer, June 2015.

[39] S. Seeber and G. D. Rodosek. Improving network security through SDN in cloud scenarios. In *Network and Service Management (CNSM), 2014 10th International Conference on*, pages 376–381. IEEE, 2014.

[40] P. Hillmann, T. Uhlig, G. Dreo Rodosek, and O. Rose. Geographical Placement of Warehouses based on the K-Center Problem. In *Proceedings of the Winter Simulation Conference (WSC)*. IEEE, 2015.

[41] Peter Hillmann, Tobias Uhlig, Gabi Dreo Rodosek, and Oliver Rose. A Novel Approach to Solve K-Center Problems with Geographical Placement. In *Proceedings of the 10th International Conference on Service Operations and Logistics, and Informatics (SOLI)*. IEEE, 2015.

[42] C. Tsiaras, M. Rösch, and B. Stiller. VoIP-based Calibration of the DQX Model. In *14th IFIP International Conferences on Networking (Networking 2015)*, 2015.

[43] J. van der Hooft, S. Petrangeli, M. Claeys, J. Famaey, and F. De Turck. A Learning-Based Algorithm for Improved Bandwidth-Awareness of Adaptive Streaming Clients. In *International Symposium on Integrated Network Management (IM 2015)*, pages 131–138. IEEE, 2015.

[44] R. Huysegems, J. van der Hooft, T. Bostoen, P. R. Alface, S. Petrangeli, T. Wauters, and F. De Turck. HTTP/2-Based Methods to Improve the Live Experience of Adaptive Streaming. In *In proceedings of the ACM Multimedia Conference (ACM MM), Brisbane, Australia*, 2015.

[45] N. Bouten, S. Latré, J. Famaey, W. Van Leekwijck, and F. De Turck. In-network quality optimization for adaptive video streaming services. *Multimedia, IEEE Transactions on*, 16(8):2281–2293, 2014.

[46] Peter Hillmann, Frank Tietze, and Gabi Dreo Rodosek. Tracemax: Single Packet IP Traceback Strategy for Data-Flow Analysis. In *Proceedings of the 40th IEEE Conference on Local Computer Networks (LCN)*. IEEE, 2015.

[47] R. De O. Schmidt, H. van den Berg, and A. Pras. Measurement-based network link dimensioning. In *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM 2015),*, pages 1071–1077, May 2015.

[48] A. Lareida, T. Bocek, M. Pernebayev, and B. Stiller. Automatic network configuration with dynamic churn prediction. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, 2015.

[49] Rashid Mijumbi, JL Gorricho, J Serrat, Meng Shen, Ke Xu, and A Kun Yang. Neuro-fuzzy approach to self-management of virtual network resources. In *Journal of Expert Systems With Applications 2014*, 2014.

[50] R. Mijumbi, J. Serrat, and J.-L. Gorricho. Self-managed resources in network virtualisation environments. *Journal of Expert Systems With Applications. Volume*, 2014.

[51] C. Schmitt, M. Keller, and B. Stiller. WebMaDa: Web-based Mobile Access And Data Handling Framework for Wireless Sensor Networks. In *Proc. of the 2015 Conference on Networked Systems (NetSys 2015)*, March 2015.

[52] R. Mujumbi, J. Serrat, J. L. Gorricho, and J. Rubio-Loyola. Survivability-oriented negotiation algorithms for multi-domain virtual networks. In *CNSM 2014*, 2014.

[53] R. Houthooft, S. Sahel Sahhaf, W. Tavernier, F. De Turck, D. Colle, and M. Pickavet. Robust geometric forest routing with tunable load balancing. In *IEEE Conference on Computer Communications (INFOCOM)*, pages 1382–1390, 2015.

[54] Rashid Mijumbi, J Serrat, J Rubio-Loyola, Niels Bouten, Filip De Turck, and Steven Latré. Dynamic resource management in sdn-based virtualized networks. In *CNSM 2015, 1st International Workshop on Management of SDN and NFV Systems*, 2014.

[55] NM Mosharaf Kabir Chowdhury and Raouf Boutaba. A survey of network virtualization. *Computer Networks*, 54(5):862–876, 2010.

[56] Bruno Nunes, Manoel Mendonca, Xuan-Nam Nguyen, Katia Obraczka, Thierry Turletti, et al. A survey of software-defined networking: Past, present, and future of programmable networks. *Communications Surveys & Tutorials, IEEE*, 16(3):1617–1634, 2014.

[57] Dmitry Drutskoy, Eric Keller, and Jennifer Rexford. Scalable network virtualization in software-defined networks. *Internet Computing, IEEE*, 17(2):20–27, 2013.

[58] Rob Sherwood, Michael Chan, Adam Covington, Glen Gibb, Mario Flajslik, Nikhil Handigol, Te-Yuan Huang, Peyman Kazemian, Masayoshi Kobayashi, Jad Naous, et al. Carving research slices out of your production networks with openflow. *ACM SIGCOMM Computer Communication Review*, 40(1):129–130, 2010.

[59] Natalia Castro Fernandes and Otto Carlos Muniz Bandeira Duarte. Xnetmon: A network monitor for securing virtual networks. In *Communications (ICC), 2011 IEEE International Conference on*, pages 1–5. IEEE, 2011.

[60] Anath Fischer, Juan Felipe Botero, Michael Till Beck, Hermann De Meer, and Xavier Hesselbach. Virtual network embedding: A survey. *Communications Surveys & Tutorials, IEEE*, 15(4):1888–1906, 2013.

[61] Yong Zhu and Mostafa H Ammar. Algorithms for assigning substrate network resources to virtual network components. In *INFOCOM*, volume 1200, pages 1–12, 2006.

[62] Francis Zane, Girija Narlikar, and Anindya Basu. Coolcams: Power-efficient tcams for forwarding engines. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, volume 1, pages 42–52. IEEE, 2003.

[63] Feng Gu, Min Peng, Samee Khan, Ammar Rayes, and Nasir Ghani. Virtual network reconfiguration in optical substrate networks. In *National Fiber Optic Engineers Conference*, pages NTh4J–6. Optical Society of America, 2013.

[64] Rashid Mijumbi, Juan-Luis Gorricho, Joan Serrat, Maxim Claeys, Filip De Turck, and Steven Latré. Design and evaluation of learning algorithms for dynamic resource management in virtual networks. In *Network Operations and Management Symposium (NOMS), 2014 IEEE*, pages 1–9. IEEE, 2014.

[65] Kalapriya Kannan and Subhasis Banerjee. Compact tcam: Flow entry compaction in tcam for power aware sdn. In *Distributed Computing and Networking*, pages 439–444. Springer, 2013.

[66] Tao Feng, Jun Bi, and Ke Wang. Joint allocation and scheduling of network resource for multiple control applications in sdn. In *Network Operations and Management Symposium (NOMS), 2014 IEEE*, pages 1–7. IEEE, 2014.

[67] Syed Ahmar Shah, Jawad Faiz, Maham Farooq, Aamir Shafi, and Syed Atif Mehdi. An architectural evaluation of sdn controllers. In *Communications (ICC), 2013 IEEE International Conference on*, pages 3504–3508. IEEE, 2013.

[68] S. Akhshabi, L. Anantakrishnan, A. C. Begen, and C. Dovrolis. What happens when http adaptive streaming players compete for bandwidth? In *22nd International Workshop on Network and Operating System Support for Digital Audio and Video*, NOSSDAV '12, pages 9–14. ACM, 2012.

[69] Z. Li, X. Zhu, J. Gahm, R. Pan, H. Hu, A. C. Begen, and D. Oran. Probe and adapt: Rate adaptation for http video streaming at scale. *IEEE Journal on Selected Areas in Communications*, pages 719–733, 2014.

[70] S. Petrangeli, M. Claeys, S. Latré, J. Famaey, and F. De Turck. A multi-agent q-learning-based framework for achieving fairness in http adaptive streaming. In *2014 IEEE Network Operations and Management Symposium (NOMS)*, pages 1–9, May 2014.

[71] M. Claeys, S. Latré, J. Famaey, T. Wu, W. Van Leekwijck, and F. De Turck. Design and optimization of a (fa)q-learning-based http adaptive streaming client. *Connection Science*, 26(01):27–45, 2014.

[72] K. Miller, E. Quacchio, G. Gennari, and A. Wolisz. Adaptation algorithm for adaptive streaming over http. In *2012 International Packet Video Workshop (PV)*, pages 173–178, May 2012.

[73] J. Jiang, V. Sekar, and H. Zhang. Improving fairness, efficiency, and stability in http-based adaptive video streaming with festive. *IEEE/ACM Transactions on Networking*, 22(1):326–340, Feb 2014.

[74] Jessica Steinberger, Lisa Schehlmann, Sebastian Abt, and Harald Baier. Anomaly Detection and Mitigation at Internet Scale: A Survey. In *Proceedings of the 7th IFIP WG 6.6 International Conference on Autonomous Infrastructure, Management, and Security (AIMS 2013)*. Springer, 2013.

[75] M. Claeys, D. Tuncer, J. Famaey, M. Charalambides, S. Latre, F. De Turck, and G. Pavlou. Towards multi-tenant cache management for isp networks. In *Networks and Communications (EuCNC), 2014 European Conference on*, pages 1–5, June 2014.

[76] M. Claeys, D. Tuncer, J. Famaey, M. Charalambides, S. Latre, G. Pavlou, and F. De Turck. Proactive Multi-tenant Cache Management for Virtualized ISP Networks. In *Network and Service Management (CNSM), 2014 10th International Conference on*, 2014.

[77] T. Winter, P. Thubert, A. Brandt, J. Hui, R. Kelsey, P. Levis, K. Pister, R. Struik, J. Vasseur, and R. Alexander. RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. *IETF RFC 6550*, March 2012.

[78] T. Phinney, P. Thubert, and R. A. Assimiti. RPL Applicability in Industrial Networks. *IETF I-D <draft-ietf-roll-rpl-industrial-applicability-02>*, October 2013.

[79] A. Brandt, E. Baccelli, R. Cragie, and P. van der Stok. Applicability Statement: The use of the RPL protocol suite in Home Automation and Building Control. *IETF I-D <draft-ietf-roll-applicability-home-building-06>*, December 2014.

[80] D. Popa, M. Gillmore, L. Toutain, J. Hui, R. Ruben, and K. Monden. Applicability Statement for the Routing Protocol for Low Power and Lossy Networks (RPL) in AMI Networks. *IETF I-D <draft-ietf-roll-applicability-ami-09>*, July 2014.

[81] J. Hui and J. Vasseur. The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams. IETF RFC 6553, March 2012.

[82] Philip Alexander Levis, Neil Patel, David Culler, and Scott Shenker. Trickle: A Self Regulating Algorithm for Code Propagation and Maintenance in Wireless Sensor Networks. In *1st Symposium on Networked Systems Design and Implementation (NSDI)*, San Francisco, CA, USA, March 2004.

[83] Anuj Sehgal, Anthéa Mayzaud, Rémi Badonnel, Isabelle Chrisment, and Jürgen Schönwälder. Addressing DODAG Inconsistency Attacks in RPL Networks. In *Proc. of GIIS conference*, 2014.

[84] Brad Karp and H. T. Kung. GPSR: greedy perimeter stateless routing for wireless networks. In *Proceedings of the 6th annual international conference on Mobile computing and networking*, MobiCom '00, pages 243–254, 2000.

[85] M. Boguñá, F. Papadopoulos, and D. Krioukov. Sustaining the Internet with hyperbolic mapping. *Nature Communications*, 1(62), 2010.

[86] Amos Korman, David Peleg, and Yoav Rodeh. Labeling schemes for dynamic tree networks. In *STACS 2002*, volume 2285 of *Lecture Notes in Computer Science*, pages 76–87. Springer, 2002.

[87] Edgar Chávez, Nathalie Mitton, and Héctor Tejeda. Routing in wireless networks with position trees. In *Ad-Hoc, Mobile, and Wireless Networks*, volume 4686 of *Lecture Notes in Computer Science*, pages 32–45. Springer, 2007.

[88] Mingdong Tang, Hongyang Chen, Guoqing Zhang, and Jing Yang. Tree cover based geographic routing with guaranteed delivery. In *Communications (ICC), 2010 IEEE International Conference on*, pages 1–5, 2010.

[89] Rein Houthooft, Sahel Sahhaf, Wouter Tavernier, Filip De Turck, Didier Colle, and Mario Pickavet. Fault-tolerant greedy forest routing for complex networks. In *RNDM'14 - 6th International Workshop on Reliable Networks Design and Modeling (RNDM 2014)*, Barcelona, Spain, November 2014.

[90] A. Cvetkovski and M. Crovella. Hyperbolic embedding and routing for dynamic graphs. In *INFOCOM 2009, IEEE*, pages 1647–1655, 2009.

[91] R. Hofstede and L. Hendriks. Unveiling SSHCure 3.0: Flow-based SSH Compromise Detection. In *Proc. of the 2015 Conference on Networked Systems (NetSys 2015)*, March 2015.

[92] Mario Golling, Robert Koch, Peter Hillmann, and Volker Eiseler. On the Evaluation of Military Simulations: Towards A Taxonomy of Assessment Criteria. In *Proceedings of the annual Military Communications and Information Systems (MilCIS) Conference*. IEEE, 2015.

[93] J.J. Santanna, R. Durban, A. Sperotto, and A. Pras. Inside booters: An analysis on operational databases. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 432–440, May 2015.

[94] J.J. Santanna, R. van Rijswijk-Deij, R. Hofstede, A. Sperotto, M. Wierbosch, L. Zambenedetti Granville, and A. Pras. Booters – An analysis of DDoS-as-a-service attacks. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 243–251, May 2015.

[95] J.J. Chromik, J.J. Santanna, A. Sperotto, and A. Pras. Booter websites characterization: Towards a list of threats. In *Proc. of the XXXIII Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (SBRC 2015)*, May 2015.

[96] L. Hendriks, A. Sperotto, and A. Pras. Characterizing the IPv6 Security Landscape by Large-Scale Measurements. In *Proc. of the 9th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2015)*, pages 145–149. Springer, June 2015.

[97] D. van der Steeg, R. Hofstede, A. Sperotto, and A. Pras. Real-time DDoS attack detection for Cisco IOS using NetFlow. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 972–977, May 2015.

[98] O. van der Toorn, R. Hofstede, M. Jonker, and A. Sperotto. A first look at HTTP(S) intrusion detection using NetFlow/IPFIX. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 862–865, May 2015.

[99] Peter Hillmann, Lars Stiemert, Gabi Dreo Rodosek, and Oliver Rose. Modelling of IP Geolocation by use of Latency Measurements. In *Proceedings of the 11th International Conference on Network and Service Management (CNSM)*. IEEE, 2015.

[100] Peter Hillmann, Lars Stiemert, Gabi Dreo Rodosek, and Oliver Rose. Dragoon: Advanced Modelling of IP Geolocation by use of Latency Measurements. In *Proceedings of the 10th International Conference for Internet Technology and Secured Transactions (ICITST-2015)*. IEEE, 2015.

[101] M. Jonker, R. Hofstede, A. Sperotto, and A. Pras. Unveiling flat traffic on the Internet: An SSH attack case study. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 270–278, May 2015.

[102] G. Machado, T. Bocek, A. Filitz, and B. Stiller. Measuring interactivity and geographical closeness of online social network users in support of social recommendation systems. In *10th International Conference on Network and Service Management CNSM 2014*, pages 187–192. IEEE, 2014.

[103] W. de Vries, J.J. Santanna, A. Sperotto, and A. Pras. How Asymmetric Is the Internet? A Study to Support the use of Traceroute. In *Proc. of the 9th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2015)*, pages 113–125. Springer, June 2015.

[104] R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras. The Internet of Names: A DNS Big Dataset. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (SIGCOMM 2015)*, pages 91–92, 2015.

[105] R. van Rijswijk-Deij, A. Sperotto, and A. Pras. DNSSEC and Its Potential for DDoS Attacks: A Comprehensive Measurement Study. In *Proceedings of the 2014 Internet Measurement Conference (IMC 2014)*, pages 449–460, Nov 2014.

[106] V. Bajpai, S.J. Eravuchira, and J. Schönwälder. Lessons Learned from using the RIPE Atlas Platform for Measurement Research. *SIGCOMM Computer Communications Review*, 45(3):35–42, July 2015.

[107] D. Dönni, G. S. Machado, C. Tsiaras, and B. Stiller. Schengen Routing: A Compliance Analysis. In *9th International Conference on Autonomous Infrastructure, Management, and Security (AIMS 2015), Lecture Notes in Computer Science, Springer*, 2015.