

Siren: A Platform for deploying Virtual Network Services in the Cloud to Fog Continuum

Lyndon Fawcett, Matthew Broadbent, and Nicholas Race

School of Computing and Communications, Lancaster University, UK
{l.fawcett1, m.broadbent, n.race}@lancaster.ac.uk

Abstract—The burden put on network infrastructures is changing. The increasing number of connected devices, along with growing demand, are creating an unsustainable future for the Internet. The recently introduced concept of Fog computing predicts a future Internet where general compute power is ubiquitous, extending the Cloud right the way to the network edge. In turn, this acts as a catalyst for Network Functions Virtualisation (NFV), increasing the potential infrastructure locations for deploying new services, specifically ones that can cater to the demands of the changing Internet. However, current realisations of NFV typically host network functions in homogeneous, centralised servers in Cloud infrastructures. This is in contrast to the Fog where environments are both distributed and heterogeneous, thus current management and orchestration platforms suffer from suboptimal service deployment. With the use of a multiple use cases, and a novel auctioning orchestration method, this paper presents *Siren*, which is an orchestrator for network functions in the Cloud to Fog continuum.

I. INTRODUCTION

The Internet is a critical infrastructure that underpins many of the services that we have grown accustomed to in modern life. However, it is a constantly shifting and increasingly used asset. Considering the expenditure and lead time of upgrading of the constituent network components, it is important any existing resource are used efficiently for the greatest possible length of time. Network operators are increasingly faced with challenges in management scalability, resource allocation, and configuration complexity [22], [23]. This takes place against a background of ever changing demand and usage profiles, making prediction and planning hard to achieve. If not correctly anticipated, a failure in this respect leads to a lower quality of service received by end-users, as congestion and contention has an impact.

Despite recent advancements in security and protection, the network and services on it can still be vulnerable to attacks, with the threat of total outages as the result of a successful attack. Importantly attacks from inside of the network are often more threatening than remote ones, resulting in orders of magnitude greater cost [24] when successful. However, this segment of the network is often neglected due to the relative difficulty and cost of deploying here. Nonetheless, edge networks still need to be protected effectively to protect businesses and homes [11]. This motivates a combined approach to network security, which moves away from the current trend of protecting networks on the at the gateway to the network. We propose that the network

should be secured at multiple and varied points, including within the network itself. When we consider that networks are increasingly being used to underpin industrial processes, and with the move towards Industry 4.0, these networks are becoming more ubiquitous and connected, the attack surface is not only increasing, but the impact becoming more severe.

This problem is likely to worsen in years to come as the Internet grows with the adoption of the Internet of Things (IoT) [27]; the Cisco Visualisation Networking Index [21] predicts that traffic demands will increase by an order of magnitude. Collectively, this is placing pressure on Internet Service Providers (ISPs) to reflect on their approach to building both network and system infrastructures and ultimately look towards more cost-effective solutions. Edge networks, contain a wide variety of heterogeneous devices such as, Sensors, actuators, Laptops, Printers, Machines etc. This diversity, matched with the scale of such devices creates even more vectors to be exploited by malicious actors. As such, an adaptable system that can react to new attack types, and rapidly address exploits is necessary.

To successfully achieve this, approaches such as Software Defined Networking (SDN) and Network Functions Virtualisation (NFV) can be applied. These offer the potential for reduced costs, better resiliency, and quicker time to market [16]. So far, implementations of these technologies have found moderate uptake in cloud-based environments and datacenters [26]. For a security application where transmission is both sensitive and limited, security functions should be placed as appropriately as possible. In other aspects, such as latency and quality of service, this also applies to other network functions such as name servers and content caches. With emerging computing architectures such as those introduced by Fog computing [8], [34], the benefits from NFV can be expanded beyond the Cloud [20], making way for new services that were previously not feasible as well as enhancing existing over-the-top network services.

Due to the latency and cost of communication, the traditional cloud architecture behind conventional NFV would be severely limited to network functions that required very little to no reactive behaviour. Many other works [17], [18], [19], [20] argue that the edge must be used in order to successfully operate certain NFV functions, and is one of the key motivators behind Fog Computing. These functions are a clear way of managing the aforementioned challenges

with the future Internet, improving network security as well as reducing the traffic burden on the network.

Fog computing has become a viable platform due to the increase in compute performance at low costs. Figures 1 and 2 show an analysis of 4,140 Customer Premise Equipment (CPEs) sold between 1998–2017 demonstrates the upward trend in device capabilities. Although there are now many devices that can now reasonably be considered as compute hosts for a Fog-NFVI, they clearly differ vastly in their capabilities relative to those found in cloud datacentres. Nonetheless, they would be perfectly capable of running the services described later in this paper and explicitly in Table I. The source of these results are available at [3].

Given the desire to deploy in many and varied locations for total network security and service optimisation, it is necessary to leverage the full continuum of resources available. This includes spanning both Cloud and Fog computing, and using resources at all points within that spectrum. Taking this approach enables the edge of the network to be adequately protected, whilst maintaining the current trend to protect networks at the gateway. Furthermore, the resource rich environment of the datacenter can also be utilised to host resource-intensive tasks such as analysis.

Yet with this continuum comes challenges around the fair allocation of resources. This arises because of the mixture of platforms that need to be used to across this continuum. It is likely that these will be platforms which allow many different parties to use the resources. Given this openness, it is envisaged that there will be a competitive process to use the underlying physical substrate, which will enable a service to be deployed in a particular and specific location. The contention present for the use of resource is only exacerbated as you move closer to the edge, where resource availability is even more limited. This creates a competitive environment, whereby many potential parties will be vying to run a variety of different services on the actual hardware. Although the hardware is likely to support multi-tenancy through virtualisation, it is nonetheless possible that the demand may outweigh the availability.

To best handle this price competitive situation, and to ensure that the infrastructure provider is fairly recompensed for the use of their resources, we propose the use of an auction-based resource allocation function. Auctions are often adopted when a fair market value cannot be easily determined; buyers indicate the price they are willing to pay, and the winner is the buyer that bids the highest amount. Auction mechanisms are in widespread use today, including cloud resource provisioning.

The work in this paper presents the prototype platform named *Siren*, which provides a solution for placement of virtual network services in the Cloud to Fog Continuum.

II. BACKGROUND AND RELATED WORK

Siren is designed as a platform and market for deploying NFV network services to the Cloud to Fog continuum. In this section we highlight similar pieces of work as well as existing tools within this space.




Fig. 1: Growth of CPU Clock Rates




Fig. 2: Advancements in Network Throughput Capability

Management And Orchestration (MANOs) frameworks are one of the primary components to deploying NFV. They are used for the automated deployment of networks as well as the ability to react to a change in the network without operator intervention [26]. The majority of the currently available MANOs are designed with a focus on deploying services to the Cloud. As previously discussed these systems do not optimally support heterogeneous environments. This said, various research efforts from ONLabs [28], [29] and multiple Universities [10], [31], [32] have already been conducted into NFV-MANO solutions hosted outside of the cloud datacentre in locations such as the operator’s edge with the purpose to virtualise the customer network functions. Collectively, these papers motivate the need to move away from Cloud computing by highlighting a major challenge of NFV-MANO, which is creating an infrastructure that can scale to run millions of VNFs [25]. However, the above solutions solve challenges of NFV orchestration by creating a brand-new infrastructure and not considering what a future Fog infrastructure may contain. NFV is not intended to require a clean slate installation, it aims to be a gradual change between existing infrastructures. This deployment will include installation into many heterogeneous environments [26].

The closest system so far to a non-Cloud based deployment solution is Central Office Rearchitected as a Datacenter (CORD) [4]. This proposes a decentralised architecture, deploying network functions to the central office. This again assumes a full shift in network and virtualisation architecture for operation. Also, whilst this provides network services closer to the networks using them, there is still value in placing certain functions closer, as well as providing an opportunity for a variety of service providers to benefit from the infrastructure. This is the primary aspect that we are researching with *Siren*.

Siren is not in direct competition with these other MANOs;

TABLE I: Survey of service requirements

Service type	Bandwidth	Latency	CPU	Storage
DPI	High	Fast (<10ms)	High	Medium
DNS	Low	Fast (<10ms)	Low	Medium
CDN	High	Slow (<100ms)	Medium	High

It’s purpose is to provide a new method of managing the network and the infrastructure that is inclusive of the Cloud to Fog continuum, of which findings could be benefit existing MANOs. *Siren* deploys initial connectivity in the network, but conceptually services for basic network connectivity are the network providers responsibility and are being solved by current MANO systems. Services deployed through *Siren*’s auction are additional services that are either benefiting the network or improving the customers experience.

Further impacting NFV deployment to the Fog is that many current attempts at making a NFV systems [13], [12] run a single network function in its own virtual machine, this is done to simplify orchestration and to make sure that each VNF is isolated. However, this is wasteful; depending on the VNF in use, a significant amount of the compute resource is consumed by the host operating system rather than the VNF. Recent maturity of containerisation can reduce the performance impact and overall cost of running VNFs verses additional VMs [5], [15], [18]. Multi-tenancy can be used to further conserve compute resource by running multiple VNFs on one machine, or even in a single process. This is point has been realised within *Siren* with the use of Docker and private docker repositories for hosting VNFs.

In the predecessors for this work [18], the use of edge/Fog networks were motivated for the placement of media content caches through a prototype system. Furthermore, in this demo [17] a preliminary auction system was designed. The work in this paper brings these two works together, evolving the system to handle more complex scenarios and other use cases, including network protection, DNS, and content caches.

III. DESIGN

This section describes the design choices and architecture use in the system. The code for all of the components discussed, including the overall prototype system are available at [2].

Figure 3 shows the various components of the architecture. The highest layer contains the *Service Providers*. These are the parties interested in using the infrastructure. They communicate directly with an *Auctioneer*, who provides details of the resources available and details of forthcoming auctions. Figure 4 is also an illustrative example of how different resources are auctioned off. It shows different resources (network and compute) and context (administrative domain) on each device that have been reserved by different service providers/bidders. In the case of this paper, there are three bidders used.

Figure 6 shows the operators interface through a web GUI, hosted by *Siren*. This shows that the devices are split up by locations (detail is more fine grained to bidders) and that

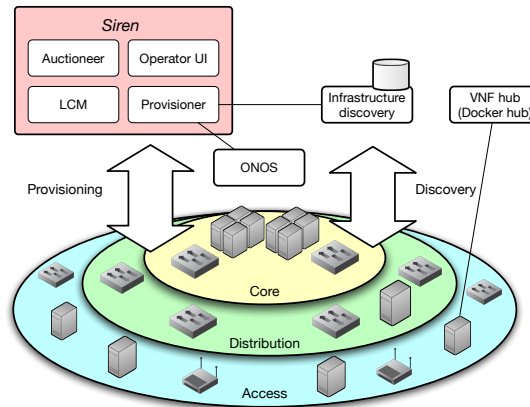


Fig. 3: System Overview

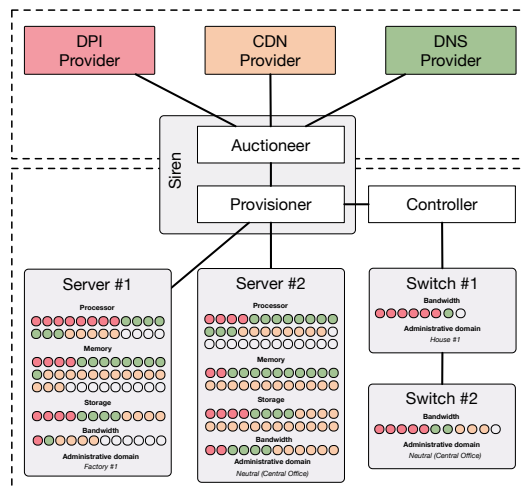


Fig. 4: Auction-Based Resource Allocation Platform




Fig. 5: Tunneling/Virtual Circuits Pipeline

there are different amounts of free resource in each pool. This also shows the ongoing live auctions and their associated packages.

A. Network Controller

The SDN functions in *Siren* are provided by the industry standard SDN controller, ONOS [7].

1) *Tunneling and virtual circuits*: Tunneling is a core requirement of SDN and NFV; tunnels are created between network functions to form service chains. However, ONOS and other SDN controllers lack a generic and transparent way of tunneling. Users must rely on vendor specific solutions that require additional configuration and are a halfway house solution to a pure SDN solution. To solve this, we created a simple tunneling solution that required no modification of existing apps within the controller.

In order to Connect the the deployed network services to the OpenFlow switches, a form of tunneling with VLANs was used. Figure 5 shows the OpenFlow pipeline we created to tunnel traffic. If traffic is marked for tunneling, it is tagged with a VLAN. Subsequent switches read this VLAN and on the packet then forward it to the out point of the tunnel. Normal untagged traffic misses the first two tables illustrated and is forwarded as normal.

It is our aim that the tunneling application presented here can be pushed back to ONOS as a support application. At the time of writing, it can be found at [2].

B. Test Network Functions

A set of example over-the-top or application layer network functions were created for demonstration purposes. These represent the functions that each service provider wants to deploy. In reality, one service provider may offer a multitude of different services.

1) *DPI*: Snort [30] is used to perform DPI on packets so that they can then be forwarded to a central control which can inform a user about. Alternatively this could have been set up in different ways, including using it as an IPS. These two alternatives have a trade-off between network overhead and remediation time. For the purposes of the results below, the DPI was in a store, check, and forward mode. *Siren* also supports using this in a mirroring configuration, whereby the user does not see any increased latency.

2) *DNS*: PiHole [1] is a popular local DNS service. This offers the ability to block certain domain names whilst providing lower latency than public DNS servers. This would be a service that could be used by either business or home users. In this example, clients are configured to use PiHole for all of their DNS requests.

3) *CDN*: A CDN based service is generally useful to home and business customers. For this we used a simple NGINX server with a big bunny file on. Through the operator interface, one can get the IP address of a new service to connect to it. This is configured through the clients video player.

C. Auctioning System

For the NFV placement problem [6], we take an approach appropriate to the heterogeneous nature of a *Fog* environment: we propose the use of a combinatorial auction [14]. This allows bidders to define bids containing combinations of discrete sets of resources. In our example scenario, each of the providers is able to bid on exactly what they require for their service. Furthermore, as resources in the Fog could

Siren visualiser

A combinatorial auctioning system for the Fog

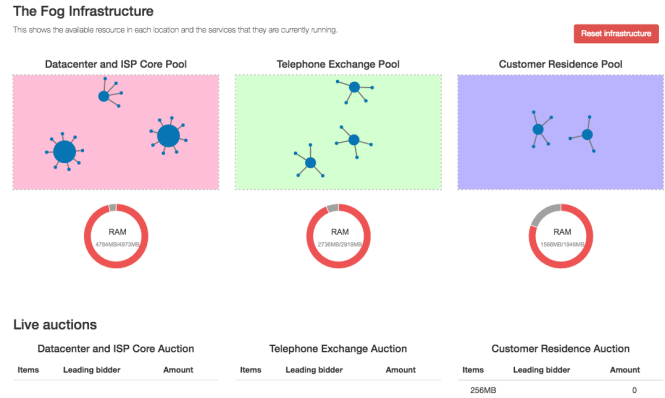


Fig. 6: Observer Interface

be located in different administrative domains, depending on the sensitivity of the service, the auction may be filtered by a resources administrative domain. For example, a bid could contain a *package* consisting of 128MB of RAM, 2 CPUs and 10GB of storage across 400 devices located specifically in local home networks. Importantly, the purpose of these bids is to reserve and allocate those resources for a fixed period of time; resources are auctioned again for other adjacent time slots. This system enables multiple providers to use the same set of resources at different times during the same day. The bidding operates in multiple phases, each bidder creates their packages in each phase. If two bidder's packages overlap, one bidder has to buy out the overlapping package. Once bidding is over, the *Auctioneer* will alert each *Service Provider* as to which of these reservations has been successful. The *Service Providers* will then supply details of the service they wish to run on their reserved infrastructure, using a supported templating language (such as Docker's *Dockerfiles*).

The next step begins the realisation of the aforementioned reservation. The *Auctioneer* informs the *Provisioner* of each successful reservation, which is then actioned on the constituent devices. This includes creating the relevant containers and services, and ensuring that they are kept within their particular resource constraints. The service is now deployed, and will remain until the reservation window expires. When this occurs, the *Provisioner* will clear the existing reservations, and make way for new services to be deployed.

IV. EXPERIMENTATION AND RESULTS

In this section the evaluation environment is laid out which is then followed by the analysis of system in different scenarios.

A. Network Provider Cost

As suggested by [9], the cost of placement of a network function is orthogonal to the cost that creates on the network. Therefore, network provider cost must be taken into account.

Equation 1 shows an example of a path cost calculation for rent for a single hour. C represents the cost set by the provider for an amount of network transfer m . A network provider may also wish to cost their network on a per hop basis, where some hops might be more expensive than others. Thus the equation supports a standard fee plus a fee for the path that the service will traverse. These values are entirely configurable within the system. For the purposes of illustration, a cost has been added to the reservation of bandwidth, however these are not necessarily representative of a true cost, as this would likely be different from one network provider to another. C_p and C_n represents 10 cents per mbps.

$$C_l = m * C_p + m * \sum_{n=1}^{path} C_n \quad (1)$$

B. Experimentation environment and scenarios

To demonstrate the purpose of selective placement and to show *Siren* in operation, we use a simplified topology which is representative of a typical ISP topology with access, distribution and core networking layers. The network topology, shown in 7 consists of 7 switches in the distribution and access layer and another 4 (not shown) in the core which are connected in a straight line, this is were the core NFVIs are located. Each link in this topology is given a 5ms delay to simulate an increased latency the further away on the network the end point is. This topology was realised using the virtual container orchestrator Mininet [33]. Mininet is an emulation tool that enables evaluation, validation and measurement of SDN applications. In order to create the network topology, Mininet instantiates LXC's (Linux containers) to act as hosts and software switches (such as Open vSwitch (OvS)). One of the significant benefits of Mininet (beyond using a simulator to create such a topology) is that each of these containers has a fully fledged and isolated networking stack. This is particularly useful when simulating separate Fog devices where we want to run various network functions, which themselves are real applications.

The testbed runs on a general-purpose server with 256GB RAM, and two Intel Xeon E5-2697v4s totalling 32 cores. The server runs Ubuntu 16.04.3 and resources were shared between Mininet (Hosts, OvS), DPI nodes, the SDN Controller (ONOS v.1.10) and *Siren*.

In order to emulate a traffic on the network, potentially similar to that seen in an ISP network, a simple web poller was used from 24 hosts. These hosts (connected to the access layer) were continuously pulling a 243MB "big bunny" video file from the core from speeds between 10-50mbps, randomly rotating between these speeds every 30 seconds. In stressing the network, we get to see queues in action, which can affect the latency of a connection, especially when large amounts of data are being transferred as part of that connection. The impact of this increases as traffic reaches the core where traffic is aggregated.

For the following results two aspects were taken into consideration: Latency, which is used to determine the level of service a service will provide at each location; Cost to the




Fig. 7: Experimentation Environment




Fig. 8: Latency Results From Mininet Experiment

network, which is used to determine if (at least from network provider fees) a service is economically viable to run at a certain location. The results on latency were ran 10 times each and then averaged.

C. Analysis

The results in Figure 8 show that the latency increases the further away from the customer the service is. This shows that the most bandwidth intensive service (DPI) increased the latency more so than the less demanding services. This is a motivating factor that the DPI service provider would use to ensure that their services were deployed as close to the customer as possible.

The results in Figure 9 shows the example calculated cost of running each network service at different locations within the network. Here we can see that if the service is demanding in terms of bandwidth, it is clearly more expensive to deploy network function further away from the customer. Whereas with a service such as DNS, where bandwidth requirements are small, assuming the latency was at an acceptable rate, it may be economical for a DNS provider to deploy to the distribution or core layers.




Fig. 9: Example Network Provider Fees Per Month

V. CONCLUSIONS AND FUTURE WORK

This paper has presented *Siren*, which is an NFV orchestration system that is used to deploy network services in a multi-tenant Cloud to Fog environment. A series of use cases were demonstrated in the paper, showing the value of a market driven solution to network service orchestration. The paper specifically focuses on the method of provisioning, auction method used, and additional commissions that are attached to each auction. For future work we are looking into the other factors that need to be considered for the winning bids. These include, total bytes the service is planning on using, minimum cost of renting an NFVI, identifying service conflicts, and verifying bids.

Furthermore, the prototype system in this paper uses a relatively simple bidder logic, in a competitive environment the bidder logic would be crucial to ascertaining a worthwhile price. This work opens up a new area of research of creating bidder logic that provides an optimal solution based on a set of requirements in a timely fashion.

ACKNOWLEDGEMENT

The authors are grateful to the UK Engineering and Physical Sciences Research Council (EPSRC) for funding the TOUCAN (EP/L020009/1) and NG-CDI (EP/R004935/1) projects, which supported the work presented in this paper, and are grateful to Lancaster University for funding Lyndon Fawcett's PhD studentship in association with TOUCAN.

REFERENCES

- [1] Dns black hole repositories. <https://pi-hole.net/>, 2017.
- [2] *Siren repositories*. <https://github.com/Siren-Project>, 2017.
- [3] *Web scraped CPE motivation results and generation*. <https://github.com/lyndon160/CPE-Scraper.git>, 2017.
- [4] A. Al-Shabibi and L. Peterson. Cord: Central office re-architected as a datacenter. *OpenStack Summit*, 2015.
- [5] P. Ann and M. Joy. Performance Comparison Between Linux Containers and Virtual Machines. (Lxc), 2015.
- [6] A. Basta, W. Kellerer, M. Hoffmann, H. J. Morper, and K. Hoffmann. Applying nfV and sdn to lte mobile core gateways, the functions placement problem. In *Proceedings of the 4th workshop on All things cellular: operations, applications, & challenges*, pages 33–38. ACM, 2014.
- [7] P. Berde, M. Gerola, J. Hart, Y. Higuchi, et al. Onos: towards an open, distributed sdn os. In *Proceedings of the third workshop on Hot topics in software defined networking*, pages 1–6. ACM, 2014.

- [8] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli. Fog computing and its role in the internet of things. In *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*, pages 13–16. ACM, 2012.
- [9] M. Bouet, J. Leguay, and V. Conan. Cost-based placement of virtualized deep packet inspection functions in sdn. In *Military Communications Conference, MILCOM 2013-2013 IEEE*, pages 992–997. IEEE, 2013.
- [10] Z. Bronstein. NFV Virtualisation of the Home Environment. *Consumer Communications and Networking Conf. (CCNC)*, (Cenc), 2014.
- [11] N. by Symantec. Norton cybercrime report, 2013.
- [12] R. Cantó Palancar, R. A. López da Silva, J. L. Folgueira Chavarría, D. R. López, A. J. Elizondo Armengol, and R. Gamero Tinoco. Virtualization of residential customer premise equipment. Lessons learned in Brazil vCPE trial. *it - Information Technology*, 57(5), 2015.
- [13] I. Cerrato, A. Palesandro, F. Risso, M. Suñé, V. Vercellone, and H. Woessner. Toward dynamic virtualized network services in telecom operator networks. *Computer Networks*, 000, 2015.
- [14] P. Cramton et al. Combinatorial auctions. 2006.
- [15] R. Cziva and D. P. Pezaros. Container network functions: bringing nfV to the network edge. *IEEE Communications Magazine*, 55(6):24–31, 2017.
- [16] N. F. V. ETSI. Introductory white paper. Technical Report, SDN and OpenFlow World Congress, 2012.
- [17] L. Fawcett, M. H. Broadbent, and N. J. P. Race. Combinatorial auction-based resource allocation in the fog. 2016.
- [18] L. Fawcett and N. Race. Siren: a platform for deployment of vnfs in distributed infrastructures. In *Proceedings of the Symposium on SDN Research*, pages 201–202. ACM, 2017.
- [19] O. C. A. W. Group et al. Openfog architecture overview. *White Paper, February*, 2016.
- [20] O. C. A. W. Group et al. Openfog reference architecture for fog computing. *OPFRA001*, 20817:162, 2017.
- [21] C. V. N. Index. Forecast and methodology, 2014-2019 white paper. Retrieved 23rd September, 2015.
- [22] H. Kim, T. Benson, A. Akella, and N. Feamster. The evolution of network configuration: a tale of two campuses. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, pages 499–514. ACM, 2011.
- [23] H. Kim and N. Feamster. Improving network management with software defined networking. *IEEE Communications Magazine*, 51(2):114–119, 2013.
- [24] T. Markham and C. Payne. Security at the network edge: A distributed firewall architecture. In *DARPA Information Survivability Conference & Exposition II, 2001. DISCEX'01. Proceedings*, volume 1, pages 279–286. IEEE, 2001.
- [25] R. Mijumbi, J. Serrat, J.-L. Gorricho, N. Bouten, F. De Turck, and R. Boutaba. Network function virtualization: State-of-the-art and research challenges. *IEEE Communications Surveys & Tutorials*, 18(1):236–262, 2015.
- [26] R. Mijumbi, J. Serrat, J.-L. Gorricho, S. Latré, M. Charalambides, and D. Lopez. Management and orchestration challenges in network functions virtualization. *IEEE Communications Magazine*, 54(1):98–105, 2016.
- [27] D. Miorandi, S. Sicari, F. De Pellegrini, and I. Chlamtac. Internet of things: Vision, applications and research challenges. *Ad Hoc Networks*, 10(7):1497–1516, 2012.
- [28] L. Peterson, A. Al-Shabibi, T. Anshutz, S. Baker, A. Bavier, S. Das, J. Hart, G. Palukar, and W. Snow. Central office re-architected as a data center. *IEEE Communications Magazine*, 54(10), 2016.
- [29] L. Peterson, S. Baker, A. Bavier, S. Bhatia, J. Nelson, M. Wawrzoniak, and J. Hartman. XOS : An Extensible Cloud Operating System. pages 23–30.
- [30] M. Roesch et al. Snort: Lightweight intrusion detection for networks. In *Lisa*, volume 99, pages 229–238, 1999.
- [31] F. Sánchez, D. Brazewell, and S. Plc. Tethered Linux CPE for IP Service Delivery. 2015.
- [32] J. Soares, M. Dias, J. Carapinha, B. Parreira, and S. Sargento. Cloud4NFV: A platform for Virtual Network Functions. *2014 IEEE 3rd International Conference on Cloud Networking (CloudNet)*, 2014.
- [33] M. Team. Mininet: An instant virtual network on your laptop (or other pc), 2012.
- [34] L. M. Vaquero and L. Roderó-Merino. Finding your way in the fog: Towards a comprehensive definition of fog computing. *ACM SIGCOMM Computer Communication Review*, 44(5):27–32, 2014.