Software-Defined-Networking-Enabled Capacity Sharing in User-Centric Networks

Bruno A. A. Nunes, Mateus A. S. Santos, Bruno T. de Oliveira, Cintia B. Margi, Katia Obraczka, and Thierry Turletti

ABSTRACT

In this article, we discuss user-centric networks as a way of, if not completely solving, considerably mitigating the problem of sharing limited network capacity and resources efficiently and fairly. UCNs are self-organizing networks where the end user plays an active role in delivering networking functions such as providing Internet access to other users. We propose to leverage the recently proposed SDN paradigm to enable cooperation between wireless nodes and provide capacity sharing services in UCNs. Our SDN-based approach allows coverage of existing network infrastructure (e.g., WiFi or 3GPP) to be extended to other end users or ad hoc networks that would otherwise not be able to have access to network connectivity and services. Moreover, the proposed SDN-based architecture also takes into account current network load and conditions, and QoS requirements. Another important feature of our framework is that security is an integral part of the architecture and protocols. We discuss the requirements for enabling capacity sharing services in the context of UCNs (e.g., resource discovery, node admission control, cooperation incentives, QoS, security) and how SDN can aid in enabling such services. The article also describes the proposed SDN-enabled capacity sharing framework for UCNs.

INTRODUCTION

The 1990s futuristic vision of "ubiquitous computing" and "anywhere, anytime connectivity" is now, only 20 years later, a reality, enabled mostly by widespread access to both portable computing devices as well as wireless communication infrastructure. Over the past few years, anywhere, anytime connectivity has resulted in exponential increase in mobile traffic, which is expected to outgrow the capabilities of current fourth generation (4G) and Long Term Evolution (LTE) infrastructure in the near future.

One possible solution to this problem would be, of course, to provision and upgrade the network infrastructure, for example, by deploying a higher number of more capable access points and base stations (e.g., conventional macro base stations or pico/femtocells). However, "throwing bandwidth at the problem" (i.e., augmenting network infrastructure at the same rate as traffic demand increases) comes with considerably high, and most of the time prohibitive, costs.

Consequently, a major challenge facing future networks is to provide ubiquitous connectivity in a scalable and resource-efficient fashion. This problem has been referred to as "network capacity sharing" [1], and has drawn considerable attention from industry and academia. *User-centric networks*, or UCNs, have emerged as a way of, if not completely solving, considerably mitigating the problem of sharing limited network capacity and resources efficiently and fairly. UCNs are self-organized networks where the end user plays an active role in networking functions such as providing Internet access to other users. As such, in UCNs, end users can act as "micro network operators," sharing their subscribed Internet access with other users, often based on some incentive mechanism. Besides extending the coverage of the Internet's backbone infrastructure at marginal cost, mitigating the capacity sharing dilemma, UCNs also improve communication services, fault tolerance, and detection, as well as load balancing. On the other hand, UCNs raise a wealth of interesting challenges themselves, ranging from providing adequate security and trust management, incentivizing users to act as micro network operators, understanding and harnessing user mobility, and coping with intermittent connectivity, to name a few.

In this article, we explore software-defined networking (SDN) as a promising approach to address some of the challenges raised by UCNs, in particular, providing efficient network capacity sharing services. The SDN paradigm has been proposed as a way to facilitate and foster Internet evolution by enabling innovation through network programmability. The main idea behind SDN is to decouple the control plane from the data plane by:

- Removing control decisions from the forwarding hardware
- Allowing the forwarding hardware to become "programmable" via an open interface

Bruno A. A. Nunes and Thierry Turletti are with INRIA.

Mateus A. S. Santos, Bruno T. de Oliveira, and Cintia B. Margi are with the University of São Paulo.

Katia Obraczka is with the University of California at Santa Cruz.

0163-6804/14/$25.00 © 2014 IEEE

- Having a separate entity called a *controller* to define, by software, the behavior of the network formed by the forwarding infrastructure, thereby creating a software-defined network

We contend that, based on its knowledge and control of the network infrastructure, the SDN controller will be able to efficiently orchestrate the capacity sharing efforts involving end-user devices as well as network access elements such as access points and base stations. In exploring SDN-enabled capacity sharing in UCNs, we describe our proposed architecture as well as functions such as mobility management, node admission control, fault tolerance, and load balancing. We also briefly discuss extending the original "logically centralized" SDN paradigm so that it can operate in distributed, decentralized UCN environments.

Background and Related Work

In this section, we provide a brief overview of UCNs and discuss some user-centric networking initiatives. We then describe the SDN paradigm and discuss its potential to enable and foster efficient capacity sharing in UCNs.

Capacity Sharing in User-Centric Networks

UCNs typically refer to wireless network deployments where end users share network resources and cooperate to provide network services. According to [2], UCN network sharing models include direct sharing, multihop networks, and user-enabled *micro-providers*.

In direct sharing, cooperation is enabled by the user if and when the user is available/willing to cooperate, for instance, by sharing network connectivity (e.g., opening access to his/her WiFi connection to other users). Resource sharing can also be enabled by the network operator; consider, for example, the case of network provider *A* allowing subscribers of network provider *B* to access *A*'s hotspots when they are in their vicinity, and network provider *B* reciprocates and allows *A*'s subscribers to access its hotspots. In fact, there are currently a number of capacity sharing services that are commercially available in the context of WiFi access networks. Notable examples include FON[1] and Whisher,[2] where users receive incentives from their Internet service providers (ISPs) to share their WiFi access.

Self-organizing, autonomous, multihop wireless networks, or mobile ad hoc networks (MANETs), have been the focus of a vast body of research since the mid-1990s [3]. The disruption tolerant networking (DTN) paradigm [4] that emerged in the early 2000s is another notable research thrust that addressed the problem of providing communication services in extreme and connectivity-challenged environments. Unlike the Internet, in these extreme environments, continuous end-to-end connectivity cannot be assumed. Originally motivated by interplanetary and deep space communication scenarios, DTNs also find applications in environmental and habitat monitoring, bridging the digital divide, emergency response, disaster relief, law enforcement, and special operations. DTNs also attracted considerable attention from the research community; in particular, the HAGGLE project [5], which was proposed in the context of pocket switched networks (PSNs) [6], was one of the early efforts recognizing the importance of user- and content-centricity in the context of network environments prone to episodic connectivity. HAGGLE explored the use of opportunistic information dissemination mechanisms, where human factors and mobility patterns play an important part.

More recently, motivated by the wide availability of portable computing devices and wireless communication infrastructure, and inspired by new user- and content-centric networking paradigms, projects such as SOCIALNETS [7] and BIONETS [8] have been proposed. SOCIALNETS considers the social interactions between users and how those can be exploited for content delivery, focusing on issues of security and trust. The main goal of the BIONETS project is to provide an integrated network and service environment that scales to large numbers of heterogeneous devices. BIONETS' scalability and adaptability are inspired by biological and social systems, in which large populations are able to reach efficient equilibrium states and develop effective collaboration strategies. Another initiative worth mentioning was the Internet Engineering Task Force (IETF)-sponsored Mobile Ad hoc Networking Interoperability and Cooperation (MANIAC) 2013 Challenge in which participants proposed, implemented, and demonstrated strategies for mobile data offloading in MANETs given cooperation incentives.

The concept of micro-providers [2] refers to end users who can act not only as consumers/producers of content, but also as providers of network access. In this context, we highlight the ULOOP [9] and PERIMETER [10] projects. ULOOP exploits how user-provided network access can help expand the coverage of a multi-access backbone infrastructure. Furthermore, ULOOP also focuses on other important aspects such as legislation implications, community-driven services, trust management, cooperation incentives, and how these aspects enable new business models for both users and access providers. The main goal of the PERIMETER project is to set a baseline for future user-centric mobility experimentation focusing on security, quality of experience (QoE), and also cooperation and trust in mobile networks.

Software-Defined Networking

The basic premise of the SDN paradigm is to decouple the network control and data planes to facilitate network protocol and service evolution, especially in production networked environments. In SDN, the network intelligence is logically centralized in software-based controllers (the control plane), and network devices become simple packet forwarding devices (the data plane) that can be programmed via an open interface (ForCES [11], OpenFlow [12], etc), which would enable programmatic control of the network's data plane.

As illustrated in Fig. 1, the separation between the forwarding hardware and the con-

In direct sharing, cooperation is enabled by the user if and when the user is available/willing to cooperate, for instance, by sharing network connectivity (e.g., opening access to his/her WiFi connection to other users). Resource sharing can also be enabled by the network operator.

[1] https://corp.fon.com

[2] http://www.whisher.com

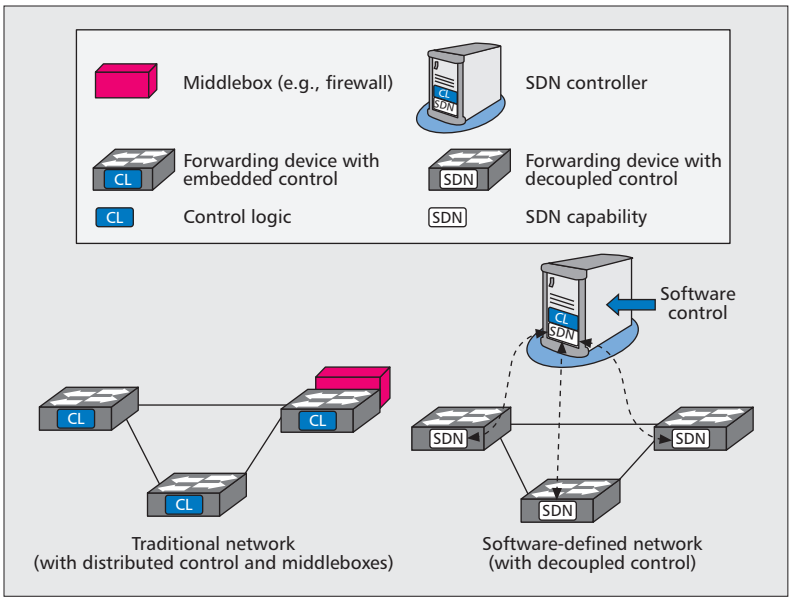

Figure 1. The SDN architecture decouples control logic from the forwarding hardware, and enables the consolidation of middleboxes, simpler policy management, and new functionalities. The solid lines define the data-plane links and the dashed lines the control plane links.

trol logic allows easier deployment of new protocols and applications, straightforward network visualization and management, and consolidation of various middleboxes into software control. Instead of enforcing policies and running protocols on a convolution of scattered devices, the network is reduced to "simple" forwarding hardware and decision-making network controller(s). A brief history of programmable networks, SDN's current state of the art, as well as current research on SDN can be found in [13].

In [14], we describe our preliminary ideas on providing capacity sharing services enabled by SDN. In this article, we go a step further and discuss a general but simple SDN-based framework and architecture for network resource sharing in user-centric networking environments. The proposed framework addresses some of the main challenges involved in enabling capacity sharing in UCNs, such as resource discovery, node admission control, support for mobility, cooperation incentives, QoS, and security.

UCN Capacity Sharing Challenges and Requirements

As mobile devices and wireless communication become prevalent, users will demand uninterrupted, high-quality communication services regardless of location or type of network access. Our premise is that self-organizing UCNs will provide cost-efficient opportunities to meet user demand for communication services in future internetworks.

Another major challenge for next-generation internets is the need to support a variety of network access technologies, applications, and end-user devices. The latter will likely be highly heterogeneous in terms of battery life, operating systems, communication range, processing and storage capabilities, and so on. Moreover, the deployed wireless infrastructure is likely to be composed of a number of radio access technologies (RATs) (WiFi, WiMAX, Third Generation Partnership Project [3GPP], LTE, etc.) that may be geographically co-located. The motivation behind this is the fact that single RAT systems are not able to offer ubiquitous coverage and QoS guarantees. Each RAT is typically composed of a set of cells, where each cell is covered by a radio access point (RAP), such as WiFi access points (APs), WiMAX routers, or 3G base stations.

User requirements and preferences also play a very important role, especially in the context of UCNs. Combined with network policy enforcement and resource availability, they will determine when and where new users can access network services. In the remainder of this section, we discuss the many challenges posed when providing capacity sharing services in UCN scenarios and the resulting functional requirements. We also explore how an SDN-enabled capacity sharing architecture can address these challenges and summarize the proposed SDN-based mechanisms in Table 1.

Resource Sharing and Allocation

Achieving efficient resource allocation is a fundamental challenge when providing network capacity sharing services in UCNs. It must account for network resource availability and usage, as well as consider user QoS requirements. The latter is discussed below.

Similar to admission control mechanisms such as joint call admission control (JCAC) for cellular networks, we propose the node admission control (NAC) mechanism. NAC will determine whether a new end user (node) can be admitted into the network and, if so, with which RAP (and thus RAT) it should be associated. Several previous JCAC algorithms consider user preferences in making RAP (and RAT) selection using, for example, multiple-objective decision making (MODM); converting imprecise variables into quantitative values; or adopting a fuzzy multiple-attributes decision making (MADM) approach. In the context of UCNs, the decision on or selection of the most suitable RAP should consider, among other things, network usage and resource availability. In an SDN-based capacity sharing architecture, the SDN controller can use its global knowledge of the network topology and conditions to decide whether new users can be admitted and, if so, how much resources can be allocated and which RAP will be used. SDN-enabled NAC can be implemented as an application running on the SDN controller. As such, the SDN controller's NAC module can decide the best available RAP for a particular user based on the controller's knowledge of the current network topology and conditions. We describe the proposed SDN-enabled NAC mechanism in more detail below.

In the case of WiFi, for example, it is worth noting that even though IEEE 802.11 specifies that a user should be associated with a single AP, and the user is the one responsible for selecting its point of attachment, standardization initiatives confirm the demand for moving away from the user-driven association model. Examples include the WiFi Alliance Hotspot 2.0,[3] which enables devices to automatically discover

and securely connect to WiFi hotspots with no user intervention, and the IETF work on network-based mobility management solutions such as Proxy Mobile IPv6 (PMIPv6) [15] and Mobile IPv6 (MIPv6) [16].

However, interactions between existing standards present compatibility issues [17], and SDN has emerged as a promising solution due to its flexibility, ease of deployment, and management. For example, the work of Dely *et al.* [18] proposes virtual SDN switches integrated into user stations in order to allow multiple AP associations. Such SDN-based schemes can be deployed to manage handovers and allow nodes to move between APs. One possibility is to have a mobility management service running on a wireless station make the decision on switching from one AP to another in a decentralized fashion, a la user-driven association. Another promising solution is the use of a centralized NAC, as we propose here, which can:

- Install flows in the virtual SDN switches of participating user stations to define the most suitable AP to be used
- On the network side, adapt the wired and wireless backhaul accordingly (e.g., establish new routes for the new selected AP)

This method can also be used to provide SDN-based services such as load balancing, fault tolerance, and mobility management.

Furthermore, when deploying any new system, backward compatibility is an important consideration. Our SDN capacity sharing framework for UCNs is no exception: it will accommodate legacy WiFi user devices by allowing them to perform traditional device-driven association. We discuss this in more detail below.

Cooperation

Social interactions and human interests are the basis for building trust; however, trust is an integral component in many kinds of human interactions, allowing individuals to act under uncertainty. Examples may include exchanging money for goods and services (e.g., reputation models in auction websites such as eBay), giving access to your property, and choosing between conflicting sources of information (e.g., wiki-pages and blogs on the web). All may utilize some sort of trust model. Trust is also the basis for end-user nodes to rely on other nodes for connectivity [2]. Recent work done in the context of ULOOP [9] and PERIMETER [10] has been focusing on trust, cooperation incentives, and reputation-based schemes for cooperation in the context of UCNs.

It is clear that incentive and trust models are necessary to ensure collaboration between nodes by incentivizing end users, acting as RAPs, to agree to forward traffic to/from other nodes. Incentive schemes may include monetary compensation, reciprocity in the form of network access credits, and so on.

QoS Support

Clearly, fulfilling end-user QoS requirements is another fundamental challenge network capacity sharing mechanisms must address. In particular, it is important to address both the requirements of the end user requesting networking services (e.g., maximum delay or minimum bandwidth requirements) as well as not deteriorating the service provided to the end-user nodes who are willing to serve as RAPs. The latter is key to ensuring cooperation on the part of currently connected end nodes, cooperation incentives aside.

Challenges	Proposed solutions
Resource sharing	Node admission control, resource discovery, and network measurements
Cooperation	Incentives
QoS support	NAC, load balancing, mobility management and handover
Security	Identity-based cryptography
Resilience	Control and data plane fault tolerance

Table 1. Challenges and solutions in enabling capacity sharing in UCNs.

For instance, it is important to limit the fraction of a node's total bandwidth as provided by the ISP that will be shared with other nodes. This allocation can be adjusted dynamically based on resource availability and allocation policies, and can be enforced by the SDN controller through ingress policies. As an example, current OpenFlow versions already allow QoS policies to be enforced by means of creating virtual ports on the switches and applying priority scheduling mechanisms such as weighted fair queuing (WFQ). QoS policies could also be used by service providers in order to offer differentiated services among users. Examples include scenarios in which customers who share their resources might get incentives to do so by means of higher ingress policies, while customers provided with temporary shared services are subject to lower bandwidth. It also allows the controller to restrict access to certain applications (e.g., deny or limit BitTorrent connections or resource-demanding applications such as video streaming) in order to preserve QoS for nodes serving as RAPs. Here, too, the use of an SDN-based architecture allows the SDN controller to implement and enforce QoS policies by employing techniques such as load balancing among RAP nodes and enforcing flow priorities.

Security

In order to control access to network resources, it is required not only to authenticate a new end-user node, but also to ascertain membership eligibility and bootstrap security services such as data confidentiality and authenticity. Clearly, security is a major concern as existing standards (e.g., 802.1x[4]) do not provide adequate security for these types of scenarios and applications. For instance, in the particular capacity sharing scenario of Fig. 2, a RAP node may need to authenticate an end-user node requesting communication services in order to make sure it is a legitimate user. Similarly, nodes should not be able to impersonate RAP nodes in order to benefit from incentives. Furthermore, RAP nodes should not be liable for misbehaving users connecting through them. At the same time, data confidentiality and data integrity should be provided to users connecting through other nodes.

[3] http://www.wi-fi.org/discover-wi-fi/wi-fi-certified-passpoint

[4] http://www.ieee802.org/1/pages/802.1x-2004.html.

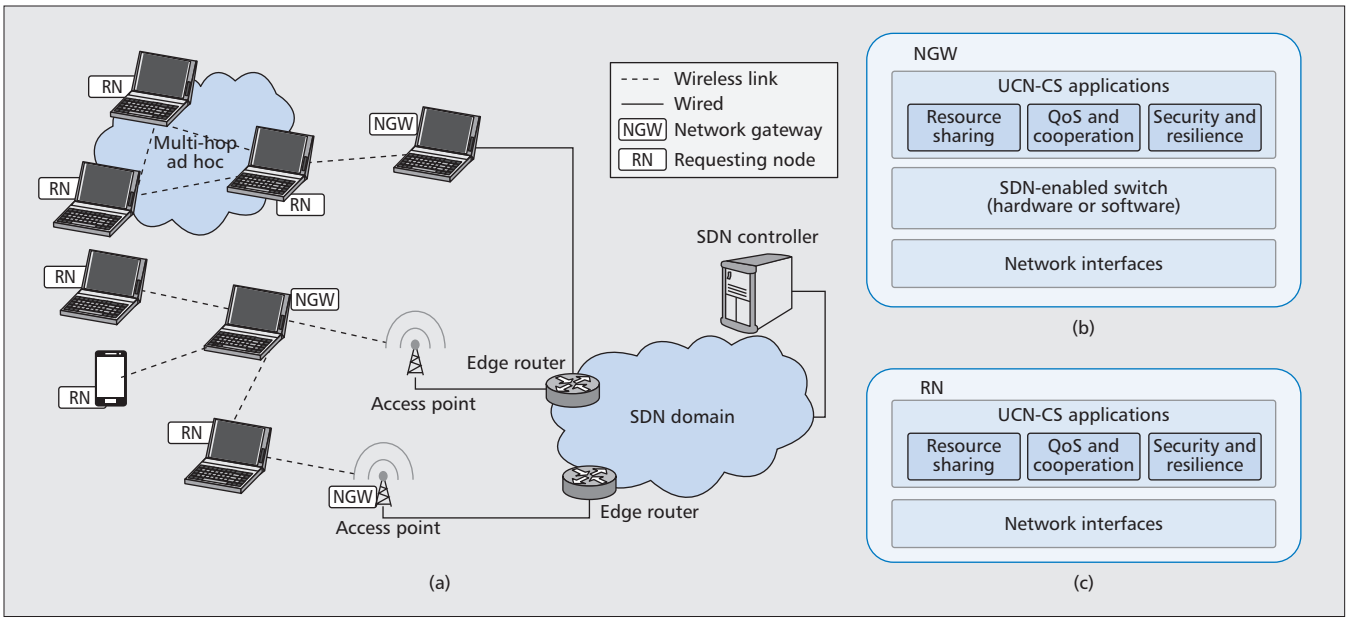

Figure 2. a) Capacity sharing architecture overview; b) network gateway (NGW); c) requesting node (RN) architectures.

Resilience

Robustness to failures in order to avoid service disruptions is also key when providing capacity-sharing-based communication services. If the current RAP node fails or is disconnected, end-user nodes connecting through it should be migrated in a seamless fashion to other RAPs. When the failed RAP comes back online, load balancing mechanisms will determine whether to migrate users back to the RAP. This seamless migration of users in response to faults can be supported quite naturally through the SDN controller in an SDN-enabled capacity sharing architecture. In fact, fault tolerance and load balancing can use common basic functions such as topology discovery, network measurement collection, and mobility management.

Control plane robustness is another fundamental challenge and can be addressed by physically distributing control. In the context of SDN, controller functions can be replicated in a number of devices that could assume control in case the current controller fails. Of course, this requires that controller replicas communicate periodically to keep their state consistency with one another, detect failures, and select a controller to take over in case of failure of the current controller.

As part of our ongoing work, we have also been exploring logically distributing the SDN controller, which addresses not only fault tolerance but also administrative decentralization, especially when considering internetworks consisting of infrastructureless self-organizing networks that are prone to episodic connectivity.

SDN-Enabled Capacity Sharing Architecture for UCNs

We contend that SDN facilitates and fosters user-centric capacity and resource sharing services by consolidating in the SDN controller network control functions as well as network structure and topology knowledge. This section provides an overview of our proposed SDN-enabled capacity sharing framework, which we call UCN — capacity sharing (UCN-CS), and how it addresses the different challenges discussed previously.

Architecture Overview

As illustrated in Fig. 2, the architectural components of the proposed UCN-CS framework include the network gateway (NGW) and requesting node (RN), which are described as follows.

Network Gateway — This is an SDN-enabled device offering gateway services through which end users can connect to the network infrastructure (e.g., the Internet). It can be an end-user device, where the user is willing to share connectivity, or an SDN-enabled RAP (WiFi APs, WiMAX routers, etc.). NGWs run UCN-CS' NGW services as depicted in Fig. 2b. Note that NGWs rely on an SDN-enabled switch in order to forward traffic accordingly. In terms of their implementation, SDN-enabled switches may, for example, comprise an instance of an OpenFlow client and act as an SDN forwarding device (e.g., using an OpenFlow software switch such as Open vSwitch [19]).

Requesting Node — Typically, this is an end-user device that can use an NGW as a provider of connectivity and networking services. RNs run UCN-CS's RN services as illustrated in Fig. 2c. Note that RNs do not need to be SDN-enabled since they do not forward traffic. For incremental deployment purposes, when communicating with legacy devices, NGWs will fall back and provide compatible connectivity services (e.g., WiFi, WiMAX).

As previously discussed, the cost of provisioning the current network infrastructure by increasing and upgrading radio coverage can be prohibitive. Furthermore, it may also be inefficient in terms of network resource utilization, especially in the case of overprovisioning to

meet peak demand. Our proposed network capacity sharing architecture broadens the scope of access network infrastructure and provides ubiquitous connectivity in a scalable and resource-efficient fashion; it does so by relying on an already deployed network of wireless (mobile) end users. For example, in the case of the scenario depicted in Fig. 2, RNs connect to the infrastructure via other end-user SDN-enabled devices acting as NGWs. The RNs can be directly connected to the NGW, connected via multiple NGWs, or reach an NGW via multihop routing in a MANET. The NGWs are controlled by an SDN controller and execute forwarding rules at the controller's command.

In the case where RNs are part of a MANET, they can reach NGWs through the MANET routing protocol run by the MANET, or by letting the SDN controller itself define the routes and install them in the MANET nodes. The second case presents many interesting opportunities and benefits, but also many challenges and open research issues. For example, relying on the global view of the network at the SDN controller enables a number of services, such as resilience to link failure, fast route (re)computation, and load balancing. In this context, the trade-offs between these added services and impact of the extra traffic overhead and delay due to SDN control operations remain to be evaluated against routing protocol overhead, performance, and services enabled by legacy MANET routing mechanisms.

MANETs are inherently decentralized, and in some cases may be intermittently connected to the network infrastructure due a variety of factors such as wireless channel impairments, power limitations, and mobility of the participating nodes, to list a few. Consequently, relying on a centralized SDN controller may not be viable. We argue that a decentralized SDN control plane approach is more adequate in such inherently distributed scenarios. If we consider the MANET example, any SDN-enabled ad hoc node could assume the role of SDN controller for the MANET when needed (e.g., in the case of network partitioning). In this case, eligible MANET nodes run an election protocol among them in order to define the most suitable candidate to take over the control of the MANET when needed. Decentralization and distribution of control in challenged networks remains an under-explored field and is one of the targets of our ongoing research efforts.

One distinguishing feature of the proposed framework is that security is an integral part of its capacity sharing services. For instance, in order to control access to network resources, a new RN must be authenticated beforehand and have its subscription verified; data confidentiality and authenticity can also be offered. Currently the basic operations provided by our framework,[5] which are illustrated in Fig. 3, are as follows.

Gateway Discovery — NGW nodes, via their UCN-CS layer (Fig. 2b), send periodic messages announcing their gateway services. Upon receiving these messages, an RN will choose an NGW by sending a Request message enabled by the RN's UCN-CS layer (Fig. 2c) to the selected NGWs. Such requests will be forwarded to the SDN controller, which will then choose the best NGW (e.g., based on the NAC's output) and assign it to the RN.

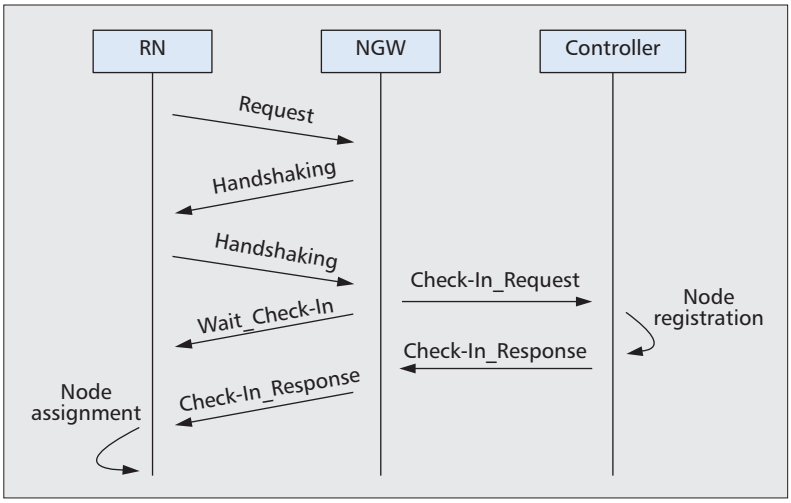

Figure 3. Basic capacity sharing operations: gateway discovery, handshaking, and RN check-in.

Handshaking — An NGW node chosen by the NAC mechanism responds to a user request and initiates a handshaking procedure for node authentication.

RN Check-in — The NGW requests authorization from the SDN controller, which queries its database in order to approve allocation of resources to the designated RN.

When an RN is authorized, the SDN controller adds the proper new entries in the flow table of the selected NGW as well as the flow table(s) in the forwarding devices on the RN's data path toward the Internet.

Node Admission Control

In order to provide efficient resource sharing and utilization, we propose the NAC mechanism illustrated in Fig. 4. A NAC works as a service running on the SDN controller and uses input from the user and forwarding devices. More specifically, the SDN-enabled NAC will receive information from SDN-capable devices (e.g., RAP devices that are able to communicate with the SDN controller, and capable of implementing and executing forwarding actions and rules) and the SDN controller's knowledge base, which may include policy rules and topology information. The NAC's decision engine would then be able to make decisions about whether to admit a new node, migrate nodes among RAPs, and so on. NAC's decision engine communicates with the OpenFlow engine, which then sets the appropriate forwarding rules at the new elected RAP device and potential forwarding devices in the data path between the new end-user device and the Internet. At this point, topology information should also be updated to maintain consistency.

The trade-offs between the amount of overhead incurred (e.g., information exchanged, stored, and processed) and the resulting accuracy and responsiveness need to be considered when designing a NAC. There are interesting

[5] Such functionalities have already been implemented and are available for download at http://inrg.cse.ucsc.edu/community/Software.

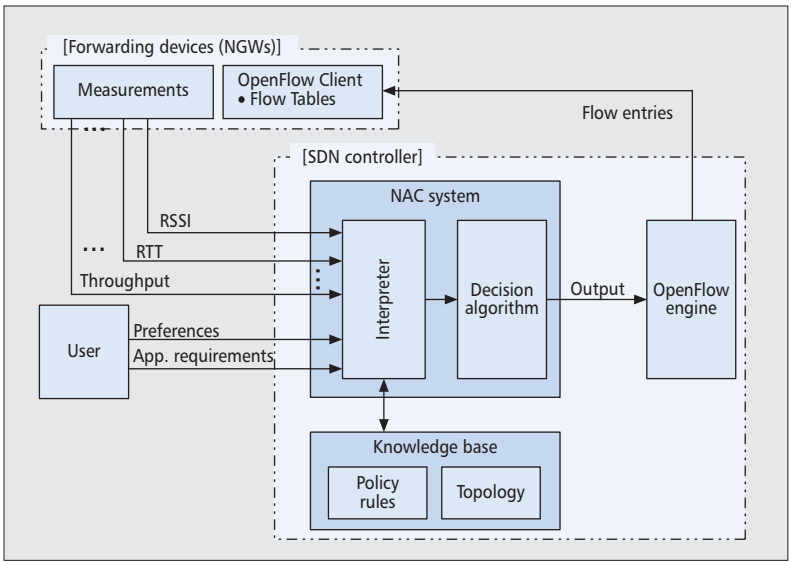

Figure 4. Node admission control (NAC).

research opportunities in addressing these tradeoffs in order to design an efficient NAC.

Resource Discovery and Network Measurements

Future internetworks will likely grow increasingly heterogeneous in terms of the devices they interconnect, and the networks and links used to interconnect them. Therefore, a variety of factors should be considered when choosing an end device as an NGW. Such factors range from battery lifetime to network connectivity and trust.

An SDN controller that collects this information periodically is able to make informed decisions about when and where to admit new nodes (i.e., performing NAC, as illustrated in Fig. 4), or to which RAP to hand over already connected RNs (e.g., performing load balancing).

Information about current network conditions is key to support decisions such as admission control, QoS, and mobility management. In the case of current SDN standards like OpenFlow, basic network measurements such as port and queue statistics can be used to estimate link bandwidth availability of RAP devices. Network statistics can be collected by NGWs and provided to controllers, which in turn can then use available bandwidth and queue statistics to decide when and where to hand over RNs. For instance, a RAP experiencing high load can have some of its RNs offloaded to another gateway device, which will likely benefit the RNs as well since they are likely to experience better service.

SDN-enabled devices can be queried periodically by the controller in order to collect relevant statistics, as depicted in Fig. 4. This figure shows a few examples of metrics/statistics relevant to the NAC's decision making process, including received signal strength indication (RSSI) and round-trip time (RTT). However, it is worth noting that in heterogeneous network technologies, link quality can be assessed differently, and different measurement mechanisms could be considered. For instance, Open vSwitch provides an interface for measuring bandwidth. However, delay-related statistics cannot be obtained using current standard SDN implementations.

Security

In order to provide security services such as confidentiality, integrity, and mutual authentication, several cryptographic schemes could be used. It is important, then, to discuss some of the tradeoffs related to choosing a particular cryptographic method, addressing efficiency, and increasing resilience to attacks such as impersonation, unauthorized data access, or data modification.

An SDN domain is composed of an SDN controller, SDN-capable devices that may be acting as NGWs, and general end-user devices as RNs. We argue that identity-based cryptography (IBC) [20] is well suited to provide simple but efficient security services in an administrative domain, which is the case of ISPs and their customers, and companies and their employees.

IBC allows a user to calculate a public key from an arbitrary string. Choosing the user's identity as a public key has advantages such as:

- There is no need to verify the public key using an online certification authority (CA).
- A user only needs the recipients' identities in order to calculate public keys (i.e., there is no need to ask for public keys). Thus, IBC-based cryptographic protocols are simple and efficient since they eliminate the need for generating and managing users' certificates. A user's public key is used as the user's identity, and the question then becomes how to obtain the corresponding secret key.

In IBC schemes, a trusted third party (TTP) is responsible for secret key generation, which is performed using the TTP's secret key, also known as master secret key, and the public key of the target user. Note that all secret keys can be computed by the TTP. Fortunately, in the scenario explored here, there is a synergy present between SDN controllers and TTPs. Controllers can be regarded as trusted entities, since they provide interfaces to applications that perform management tasks. Thus, in the context of IBC, a controller could be responsible for generating (and possibly distributing) private keys to users in its domain.

However, when inter-domain functions and services are considered (e.g., cooperation between different ISPs), IBC may not be adequate. Since the TTP can impersonate all users on the network, and access and modify transmitted data, it is strongly recommended that the TTP be managed by a single owner. In this case, public key distribution schemes need to be implemented. Additionally, public key validation are needed to ensure that a public key received from the network is indeed the correct one. This is usually implemented using a public key infrastructure (PKI), which then requires a trusted entity, the CA, to be online to respond to certification requests (i.e., a certificate for the requested public key and identity).

Mobility Management and Handover

With the proliferation of wireless mobile devices connecting to the Internet, there is a need for efficient and scalable mobility management in order to guarantee uninterrupted network ser-

vices while maintaining desired levels of QoS. In this context, SDN enables fast and transparent switching of RNs between NGWs. This can be achieved by instantiating an SDN software switch in the RN, which is then able to operate as a bridge with virtual interfaces, each one associated with a single NGW [21]. The "best" NGW could be used as a primary access device to the infrastructure network.

In a densely deployed access network, the choice of the NGW can be made by the user device or the network infrastructure. This also defines which one will be in charge of mobility management. Should the RNs be in charge, they need to be provided with information for deciding where and when to switch (e.g., signal strength, link quality statistics). On the other hand, the SDN controller's global network knowledge allows it to provide centralized NAC, in which users are associated with NGWs automatically based on users preferences and availability of resources. Alternatively, the controller could simply only expose to each RN the particular NGW with which it wants the RN to associate [22].

Incentives and Trust Models

In an SDN-enabled capacity sharing framework, the SDN controller can serve as the entity that provides incentives to end users to share their connectivity with others. The SDN controller is also responsible for authenticating nodes that are willing to serve as RAPs as well as end-user nodes. Unlike the current SDN paradigm, which relies on (logically) centralized control, in a distributed SDN model, the main SDN controller for a particular domain (e.g., ISP) could delegate certain decisions (e.g., whether to agree to forward traffic for a particular node or set of nodes) to local controllers (e.g., RAPs).

Furthermore, it is not enough that control messages successfully and securely reach their destination; both endpoints must be able to trust each other to act properly. Forwarding nodes need to be able to trust that the discovered controller is not malicious before accepting control. Likewise, the controller must be able to trust that forwarding nodes that have accepted control acting as NGWs are correctly following instructions. For this trust to exist, mechanisms must be in place to ensure the legitimacy of nodes/controllers and the authenticity of the control traffic, and verify that devices act as expected in response to instructions. Additionally, with a global view of the network, an SDN controller can decide whether to delegate forwarding capabilities to potential forwarding nodes or even to permit or prevent access from RNs and other devices based on an eligibility function implementing needed trust and reputation models.

Load Balancing

Detecting overloaded NGWs is important to guarantee adequate network performance. This can be achieved by using centralized NAC to switch an RN to a new NGW. As previously pointed out, the SDN controller, with its ability to obtain global knowledge of the network, can decide to migrate RNs among NGWs in order to balance network load and thus offer adequate service to end users. An interesting challenge to consider is how to prevent possible oscillation (or "ping-ponging") of an RN between NGWs. Since the network is densely deployed, redundancy can be used to avoid excessive user device migration situations. For example, an already migrated RN that further experiences low throughput can be provided with services by more than one NGW simultaneously. Moreover, when moving RNs from one NGW to another, based on available resources and load conditions, the NAC system may be in charge of such decisions, and a hysteresis approach must be taken in order to efficiently reduce the handover initiation delay, avoid ping-pong behavior (i.e., multiple and consecutive associations and disassociations between a group of NGWs), and decrease the number of unnecessary handovers.

Fault Tolerance

The goal of fault tolerance at the data plane is to detect link failures and take recovery actions. In the case of OpenFlow, there are no topology monitoring specifications that can be implemented by leveraging Link-Layer Discovery Protocol (LLDP) messages or customizing switch functionalities [23].

In the context of UCNs, node failures should also be considered, which are somewhat related to load balancing. Keeping track of nodes not only allows to detect an overload situation, but also to detect a node failure. In the latter case, SDNs can be useful for acquiring the measurements from the forwarding devices, as discussed previously, and to set up the flows from the Internet to a new NGW via alternative paths, as needed. A centralized NAC system can maintain a location database so that the flow affected by a NGW failure would be redirected accordingly.

Controller faults can be dealt with by a variety of methods. SDN implementations such as OpenFlow consider a fail-safe mode for SDN-enabled devices so that packets can be forwarded by using the same method as traditional layer 2 (L2) switches. A rather efficient design choice is to distribute the control plane using a single controller in charge while the others operate as replicas [24].

As previously pointed out, we are also exploring the idea of logically distributing the control plane using a hierarchy of controllers. According to the control hierarchy, controllers at different hierarchical levels will be responsible for different control functions. The control hierarchy tries to match the internetwork's hierarchical structure, where the main controller resides at the backbone level, and secondary controllers are responsible for regionals, stubs, and so on. This decentralized control model is also well suited for internetworks consisting of infrastructure-based as well as self-organizing networks that may be frequently disconnected from the infrastructure.

The goal of fault tolerance at the data plane is to detect link failures and take recovery actions. In the case of OpenFlow, there are no topology monitoring specifications that can be implemented by leveraging Link-Layer Discovery Protocol (LLDP) messages or customizing switch functionalities.

Concluding Remarks

In this article, we explore SDN as a promising approach to address the challenges raised by UCNs, in particular providing efficient network capacity sharing services. We contend that, based on its knowledge and control of the network infrastructure, the SDN controller will be able to efficiently orchestrate capacity sharing

We contend that, based on its knowledge and control of the network infrastructure, the SDN controller will be able to efficiently orchestrate capacity sharing efforts involving end-user devices as well as network access elements such as access points and base stations.

efforts involving end-user devices as well as network access elements such as access points and base stations. In exploring SDN-enabled capacity sharing in UCNs, we discuss requirements and challenges raised by sharing network resources in a scalable and efficient manner, and propose a simple but general framework that includes functions such as mobility management, node admission control, fault tolerance, and load balancing. One distinguishing feature of our SDN-based capacity sharing approach is that security is an integral part of the proposed framework.

Acknowledgments

The authors would like to thank the reviewers for their careful examination of the manuscript and valuable comments, which helped to considerably improve the quality of the article. This work is partly funded by the Community Associated Team between INRIA and UCSC and the French ANR under the DISCO "ANR-13-INFR-013" project, and by NSF grant CNS 1150704. It was also partly funded by Brazil CNPq under the grant 245588/2012-4, and São Paulo Research Foundation (FAPESP), Brazil, under grant 2013/15417-4.

References

[1] M. Mendonca *et al.*, "Software Defined Networking for Heterogeneous Networks," *IEEE COMSOC MMTC E-Letter*, 2013.
[2] R. Sofia and P. Mendes, "User-Provided Networks: Consumer as Provider," *IEEE Commun. Mag.*, vol. 46, no. 12, Dec. 2008, pp. 86–91.
[3] I. Chlamtac, M. Conti, and J. J.-N. Liu, "Mobile Ad Hoc Networking: Imperatives and Challenges," *Ad Hoc Networks*, vol. 1, no. 1, 2003, pp. 13–64.
[4] Z. Zhang and Q. Zhang, "Delay/Disruption Tolerant Mobile Ad Hoc Networks: Latest Developments," *Wireless Commun. and Mobile Computing*, vol. 7, no. 10, 2007, pp. 1219–32.
[5] EC IST Program, "HAGGLE — An innovative Paradigm for Autonomic Opportunistic Communication," 2006–2010.
[6] P. Hui *et al.*, "Pocket Switched Networks and Human Mobility in Conference Environments," *Proc. 2005 ACM SIGCOMM Workshop on Delay-Tolerant Networking*, 2005, pp. 244–51.
[7] EC IST Program, "SOCIALNETS — Social Networking for Pervasive Adaptation," 2008–2011.
[8] EC IST Program FP6-027748, "Biologically Inspired Network and Services — BIONETS," 2006–2010.
[9] EC IST Program FP7 257418, "ULOOP project," 2014.
[10] EC IST Program FP7 224024, "PERIMETER — User-Centric Paradigm for Seamless Mobility in Future Internet," 2008–2011.
[11] A. Doria *et al.*, "Forwarding and Control Element Separation (ForCES) Protocol Specification," RFC 5810 (proposed standard), Mar. 2010.
[12] N. McKeown *et al.*, "Openflow: Enabling Innovation in Campus Networks," *ACM SIGCOMM Computer Commun. Review*, vol. 38, no. 2, 2008, pp. 69–74.
[13] B. Nunes *et al.*, "A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks," 2014.
[14] M. A. S. Santos *et al.*, "Software-Defined Networking Based Capacity Sharing in Hybrid Networks," *Proc. 21st IEEE Int'l Conf. Network Protocols*, 2013, Oct. 2013, pp. 1–6.
[15] K. Leung *et al.*, "Proxy Mobile IPv6," RFC 5213, Aug. 2008.
[16] D. Johnson, C. Perkins, and J. Arkko, "Mobility Support in IPv6," RFC 3775, June 2004.
[17] G. Giaretta, "Interactions between Proxy Mobile IPv6 (PMIPv6) and Mobile IPv6 (MIPv6): Scenarios and Related Issues," RFC 6612, May 2012.
[18] P. Dely *et al.*, "Best-Ap: Non-Intrusive Estimation of Available Bandwidth and Its Application for Dynamic Access Point Selection," *Computer Commun.*, vol. 39, no. 0, 2014.
[19] B. Pfaff *et al.*, "Extending Networking into the Virtualization Layer," *Proc. ACM SIGCOMM Workshops*, 2009.
[20] D. Boneh and M. Franklin, "Identity-Based Encryption from the Weil Pairing," *CRYPTO 2001*, J. Kilian, Ed., vol. 2139 of LNCS, Springer, 2001, pp. 213–29.
[21] P. Dely *et al.*, "Best-Ap: Non-Intrusive Estimation of Available Bandwidth and Its Application for Dynamic Access Point Selection," *Computer Commun.*, 2013.
[22] R. Murty *et al.*, "Designing High Performance Enterprise Wi-Fi Networks," *NSDI*, vol. 8, 2008, pp. 73–88.
[23] J. Kempf *et al.*, "Scalable Fault Management for Openflow," *Proc.IEEE ICC*, June 2012, pp. 6606–10.
[24] S. Jain *et al.*, "B4: Experience with a Globally Deployed Software-Defined WAN," *SIGCOMM Comp. Commun. Rev.*, vol. 43, no. 4, Aug. 2013, pp. 3–14.

Biographies

Bruno Astuto A. Nunes is a post-doctoral researcher at INRIA Sophia Antipolis, France. He received his B.Sc. in electronic engineering at the Federal University of Rio de Janeiro (UFRJ), Brazil, where he also completed his M.Sc. degree in computer science. He received his Ph.D. degree in computer engineering from the University of California (UC), Santa Cruz.

Mateus A. S. Santos (mateus@larc.usp.br) received his Bachelor's degree from Universidade Federal de Itajuba in 2001 and his Master's egree from Instituto de Matematica e Estatistica, Universidade de Sao Paulo (IME/USP) in 2009. He is a doctoral student at the Department of Computer Engineering and Digital Systems (PCS), Escola Politecnica da Universidade de Sao Paulo (EPUSP) since 2010. From 2013 to 2014 he was also a research scholar with the Computer Engineering Department at UC Santa Cruz. He has industry experience with companies such as Hewlett-Packard. His primary research interests are in the area of software-defined networking, network security, and sensor networks.

Katia Obraczka [F] is a professor of computer engineering at UC Santa Cruz. Before joining UCSC, she was a esearch scientist at the University of Southern California's (USC's) Information Sciences Institute (ISI) and had a joint appointment at USC's Computer Science Department. Her research interests span the areas of computer networks, distributed systems, and Internet information systems. Her lab, the Internetwork Research Group (i-NRG) at UCSC, conducts research on designing and developing protocol architectures motivated by the internets of the future. She has been a PI and a co-PI in a number of projects sponsored by government agencies (e.g., NSF, DARPA, NASA, ARO, DoE, AFOSR) as well as industry (e.g., Cisco, Google, Nokia).

Thierry Turletti is a senior research scientist in the DIANA team at INRIA. He received M.Sc. and Ph.D. degrees in computer science from the University of Nice-Sophia Antipolis, France. His current research interests include software defined networking, trustable network evaluation platforms and wireless networking. He is serving on the editorial boards of the *Wireless Networks* and *Advances in Multimedia* journals.

Bruno T. de Oliveira is a doctoral student in the Department of Computer Engineering and Digital Systems (PCS), Escola Politecnica da Universidade de Sao Paulo (EPUSP). His main research interests include wireless sensor networks, software-defined networking, and applied cryptography. He received his Master's degree from EPUSP in 2012 and his Bachelor's degree in information systems from Escola de Artes, Ciencias e Humanidades da Universidade de Sao Paulo (EACH-USP) in 2009.

Cíntia Borges Margi has been an assistant professor at PCS of EPUSP since June/2010. From February 2007 until May 2010, she worked as an assitant professor in EACH/USP in the information systems area. She received her B.S. and M.Sc. in electrical engineering from the University of Sao Paulo in December 1997 and December 2000, respectively. She finished her Ph.D. in the Computer Engineering Department at UC Santa Cruz in June 2006, where she received a four-year fellowship from CNPq-Brazil. Her research interests include computer nNetworks and architecture, specifically wireless sensor networks (protocols, systems, security, energy consumption, and management) and software-defined networking.