REVIEW PAPER

Invited Paper: The Audacity of Fiber-Wireless (FiWi) Networks: Revisited for Clouds and Cloudlets

Martin Maier , Bhaskar Prasad Rimal
Optical Zeitgeist Laboratory, INRS, Montréal, H5A 1K6, Canada

Abstract: There is a growing awareness among industry players of reaping the benefits of mobile-cloud convergence by extending today's unmodified cloud to a decentralized two-level cloud-cloudlet architecture based on emerging mobile-edge computing (MEC) capabilities. In light of future 5G mobile networks moving toward decentralization based on cloudlets, intelligent base stations, and MEC, the inherent distributed processing and storage capabilities of radio-and-fiber (R&F) networks may be exploited for new applications, e.g., cognitive assistance, augmented reality, or cloud robotics. In this paper, we first revisit fiber-wireless (FiWi) networks in the context of conventional clouds and emerging cloudlets, thereby highlighting the limitations of conventional radio-over-fiber (RoF) networks such as China Mobile's centralized cloud radio access network (C-RAN) to meet the aforementioned trends. Furthermore, we pay close attention to the specific design challenges of data center networks and revisit our switchless arrayed-waveguide grating (AWG) based network with efficient support of east-west flows and enhanced scalability.

Keywords: cloudlet; computation offloading; east-west flows; intelligent base station; mobile-cloud convergence; mobile data offloading; mobile-edge computing (MEC); scalability

I. INTRODUCTION

Back in 2008, we envisioned in our paper "The Audacity of Fiber-Wireless (FiWi) Networks" [1] that *utility-supplied* computing, e.g., Google, will continue to have an increasing impact on society and replace in-house computer facilities just like in-house power generators were replaced with electrical utilities, unless new services and applications will be developed that capitalize on them. Toward this end, we advocated that FiWi networks be built using low-cost, simple, open, and ubiquitous optical fiber and wireless Ethernet technologies that allow all end users to have broadband access in order to shift the research focus from bridging the notorious first/last mile bandwidth bottleneck to the exploitation of distributed storage and processing capabilities and thereby create unforeseen services and applications that help stimulate innovation, generate revenue, and improve the quality of our every-day lives.

Fast-forwarding to 2015, there is now

growing awareness among industry players of reaping the benefits of openness and end-to-end Internet design principles at the early stages of the mobile-cloud convergence by extending today's unmodified cloud to a second level consisting of self-managed data centers with no hard state called *cloudlets* [2]. Cloudlets are decentralized entities that are located at the edge of the Internet and are well connected to it (e.g., via fiber), just one wireless hop away from associated mobile devices, thus enabling new applications that are both compute intensive and latency sensitive, e.g., cognitive assistance, augmented reality, or face recognition and navigation for emerging cloud robotics [3]. Beside these futuristic applications, cloudlets can also enable mobile access to the huge legacy world of desktop computers and applications that are likely to remain important for many years to come, as already envisioned in [1] seven years ago. The resultant two-level cloud-cloudlet architecture leverages both centralized and distributed cloud resources and services, whereby cloudlet infrastructure is deployed much like WiFi access points today and mobile devices can gracefully degrade to a fall-back mode that involves a distant cloud, if no cloudlet is available nearby [4].

Recently, a similar trend toward decentralization can be observed for the changing role of the mobile base station and the compelling opportunity for it to become a hub of service creation in order to monetize from the provision of an enriched mobile broadband experience that is characterized by low-latency, contextualized, and highly personalized applications and services [5]. Toward this end, intelligent base stations are equipped with a range of IT-based capabilities such as localized processing and content storage (e.g., of popular or viral videos) to extend applications and services, which would normally reside within the Internet or the mobile operator's centralized core and data centers, in close proximity to the mobile subscriber. As a result, functions are redistributed to the edge of mobile networks, directly inside base stations that haven't changed much since the 1990's, with the goal to transform them into a value creation engine for mobile network operators [6].

This trend toward decentralization of mobile networks is reinforced by the recently launched *mobile-edge computing* (*MEC*) industry initiative [7]. MEC aims at uniting the telco and IT cloud worlds, providing IT and cloud computing capabilities within the radio access network (RAN) for delivering services directly from the edge of the network. MEC helps operators to meet the challenges posed by the mobile data traffic explosion. More importantly, proximity, context, agility, and speed can be translated into unique value and revenue generation and can create opportunities for mobile operators, service and content providers, over-the-top (OTT) players, and independent software vendors (ISVs). Creating a new value and a refreshed ecosystem based on innovation and business value allows all players to benefit from greater cooperation and develop new disruptive applications, vertical services, and business models. Mobile operators can play a pivotal role within the new value chain and attract OTTs, developers, and Internet players to innovate over a new cutting-edge technology. There are different ways to implement MEC depending on the deployment scenarios under consideration, e.g., base stations, small cells, aggregators, or WiFi access points. Cloud computing at the edge will not only decrease operational expenditures (OPEX), but also accelerate the overall return on investment (ROI) by creating new revenue streams [8].

Fiber optic communication technologies play a crucial role for both intra- and inter-data center communications, whereby the growing deployment of new broadband access network infrastructures such as *fiber to the home* (*FTTH*) will exacerbate the bandwidth requirements for inter-datacenter long-haul and metro networks [9]. The benefits of exploiting optical circuit switching (OCS) in conventional packet switched data centers have been studied for years. For instance,

This paper revisits the FiWi networks to elaborate the feasibility and benefits of decentralizing the cloud computing and storage resources in future mobile networks.

the c-Through prototype proposed in [10] is a hybrid packet and circuit switched data center network architecture, which augments the traditional hierarchy of electrical Ethernet packet switches with a high-speed, low-complexity, rack-to-rack optical circuit-switched network to supply high bandwidth to applications. The optical switch can be reconfigured within a few milliseconds during which the optical paths are unusable. To ensure support for latency-sensitive applications, c-Through retains the electrical packet switch. Another early example of a hybrid electrical packet/optical circuit switch architecture for data centers is the HELIOS prototype, which is able to provide significant reductions in the number of switching elements, cabling, cost, and power consumption relative to previously proposed data center network architectures [11]. More recently, a novel optical switched data center interconnect based on hybrid optical switching (HOS) was proposed in [12]. HOS integrates OCS with optical burst switching (OBS) and optical packet switching (OPS) within the same data center network such that different data center applications are mapped to the optical transport mechanism (OCS, OBS, or OPS) that best suits their respective traffic characteristics. This ensures high flexibility and efficient resource utilization.

Clearly, there is a huge potential waiting to be unleashed by exploiting various optical switching techniques for emerging and future optical interconnection networks for cloud data centers. In this paper, however, we argue that there exists a very profound understanding of the merits and shortcomings of different optical switching techniques obtained from decades of research on optical wavelength division multiplexing (WDM) networks that may be adapted to data centers in a rather straightforward fashion. Therefore, we will not describe them in greater detail in the remainder of this paper, but refer the interested reader to [13] for a comprehensive survey on the major optical switching techniques, including not only OCS, OBS, and OPS but also optical flow switching (OFS), which may be of particular importance for virtual machine migrations and Big Data applications with high volume traffic flows. The focus and contributions of this paper are twofold:

- First, we revisit FiWi networks in the context of both conventional clouds and emerging cloudlets, paying particular attention to the two different types of FiWi networks: (i) traditional radio-over-fiber (RoF) networks, and (ii) so-called radio-and-fiber (R&F) networks [1]. RoF networks have been studied for decades and were also used in China Mobile's cloud RAN (C-RAN), which relies on a centralized cloud infrastructure and moves baseband units (BBUs) away from remote radio heads (RRHs), rendering the latter ones intentionally as simple as possible without any processing and storage capabilities [14]. Conversely, R&F networks are based on decentralized (optical and wireless) Ethernet technologies and perform protocol translation at the optical-wireless interface in order to cope with the disparate optical and wireless media in a more efficient fashion. Beside medium access control (MAC) protocol translation, the distributed processing and storage capabilities built into R&F networks may be exploited for a number of additional tasks. R&F may become the FiWi network type of choice in light of the aforementioned trends of future 5G mobile networks toward decentralization based on cloudlets, intelligent base stations, and MEC.
- Second, we pay close attention to some of the particular challenges of data center networks. Beside capacity and low latency, optical interconnection networks for cloud data centers have to cope with their vast power consumption and cooling demands while satisfying given cost constraints. Potentially even more important than power consumption, however, is scalability, e.g., increase of bit rates from 10 Gb/s to

40 Gb/s without requiring any upgrade of the optical interconnection network itself. Furthermore, data center network solutions have to account for the fact that east-west flows between racks are dominating data center traffic due to replication, backup, server virtualization, and parallel processing, representing 75% or more of total data center traffic [15]. Toward this end, we will revisit our high-performance switchless WDM network that is based on a passive wavelength router with spatial wavelength reuse capability for increased network capacity and bit/protocol transparency and that uses multiple free spectral ranges (FSRs) of the underlying wavelength router for improved scalability.

The remainder of the paper is structured as follows. In Section 2, we briefly review key enabling RoF and R&F technologies and highlight their joint use in our recently proposed FiWi enhanced 4G LTE-Advanced (LTE-A) heterogeneous networks (HetNets) for an improved mobile data offloading efficiency. Section 3 describes the benefits of computation offloading in FiWi enhanced LTE-A HetNets that are empowered by clouds and cloudlets. In Section 4, we elaborate on our switchless single-hop WDM optical cloud network solution for the efficient support of east-west flows, which provides transparency for easy upgradability and a significantly increased degree of concurrency for improved scalability. Finally, Section 5 concludes the paper.

II. FiWi Networks: RoF vs. R&F

By combining the capacity of optical fiber networks with the ubiquity and mobility of wireless networks, FiWi networks form a powerful platform for the support and creation of emerging as well as future unforeseen applications and services. While RoF networks use optical fiber as an analog transmission medium between a central station and one or more remote antenna units (RAUs) with the central base station being in charge of controlling access to both optical and wireless media, in R&F networks access to the optical and wireless media is controlled separately from each other by using in general two different MAC protocols in the optical and wireless media, with protocol translation taking place at their interface. A plethora of enabling optical and wireless technologies have been emerging that can be used to build future-proof bimodal FiWi broadband access networks. In the following, we provide a brief review of key enabling RoF and R&F technologies (a more detailed description can be found in [1] and references therein). Furthermore, we elaborate on our recently proposed concept of FiWi enhanced LTE-A HetNets, which allows the joint use of RoF and R&F technologies.

2.1 RoF technologies

To avoid the electronic bottleneck, the generation of radio frequency (RF) signals is best done optically. Among others, the following optical RF generation techniques were experimentally studied and demonstrated: (*i*) four-wave mixing (FWM) in a highly nonlinear dispersion-shifted fiber (HNL-DSF), which is transparent to the bit rate and modulation format; (*ii*) cross-phase modulation (XPM) in a nonlinear optical loop mirror (NOLM) in conjunction with straight pass in HNL-DSF, which enables the all-optical up-conversion of multiple wavelength channels; (*iii*) all-optical up-conversion by means of cross-absorption modulation (XAM) in an electroabsorption modulator (EAM), which has several advantages such as low power consumption, compact size, polarization insensitivity, and easy integration with other devices; (*iv*) external intensity modulation (IM), which may deploy different modulation schemes such as double-sideband (DSB), single-sideband, and optical carrier suppression; and (*v*) external phase modulation (PM). Among these techniques, external intensity and phase modulation schemes are the most practical solutions for all-optical RF generation due to their low cost, simplicity, and long-distance transmission performance.

An interesting approach to build low-cost FiWi networks is the use of a single light source at the central office to generate a downlink wavelength that is reused at RAUs for upstream transmission by means of remote modulation, thereby avoiding the need for an additional light source at each RAU. Among others, the following remodulation schemes were experimentally studied: (*i*) differential phase-shift-keying (DPSK) for downstream and on-off-keying (OOK) for upstream; (*ii*) optical carrier suppression for downstream and reuse for upstream; and (*iii*) PM for downstream and a directly modulated semiconductor optical amplifier (SOA) for upstream. The use of a colorless (i.e., wavelength-independent) SOA as an amplifier and modulator for upstream transmission provides a promising low-cost RoF solution that is easy to maintain.

2.2 R&F technologies

The vast majority of R&F based FiWi networks consist of a cascaded time division multiplexing (TDM) IEEE 802.3ah Ethernet passive optical network (EPON) in the backhaul and an IEEE 802.11 a/b/g/n/s wireless local area network (WLAN) mesh frontend. Apart from next-generation PONs, e.g., IEEE 802.3av 10G-EPON or WDM PON, the following optical technologies play an important role in the design of a flexible and cost-effective optical backhaul for FiWi networks: (*i*) tunable lasers such as directly modulated external cavity lasers, multisection distributed feedback (DFB)/distributed Bragg reflector (DBR) lasers, and tunable vertical-cavity surface-emitting laser (VCSEL); (*ii*) tunable receivers, which may be realized by using a tunable optical filter and a broadband photodiode; (*iii*) colorless optical network units (ONUs) based on reflective SOAs (RSOAs), which remotely modulate optical signals generated by centralized light sources; (*iv*) burst-mode laser drivers, which provide fast burst on/off speed, sufficient power suppression during idle period, and stable, accurate power emission during burst transmission; and (*v*) burst-mode receivers at the central optical line terminal (OLT) of a PON, which must exhibit a high sensitivity, wide dynamic range, and fast time response to arriving bursts.

2.3 FiWi enhanced LTE-A HetNets: mobile data offloading

The aforementioned RoF and R&F technologies may be used separately from each other or jointly together. To illustrate their joint use, we consider our recently proposed concept of FiWi enhanced LTE-A HetNets, which aims at unifying coverage-centric 4G mobile networks and capacity-centric FiWi broadband access networks based on data-centric Ethernet technologies in response to the unprecedented growth of mobile data traffic [16]. Figure 1 depicts the generic architecture of FiWi enhanced LTE-A HetNets. There are three different subsets of ONUs. An ONU of the first subset serves a single or multiple attached fixed (non-mobile) wired subscribers. An ONU of the second subset connects to a cellular network base station (BS), which may be a conventional macrocell BS (e.g., BS1 in Figure 1) or a small cell BS with a smaller wireless coverage area (e.g., BS 2 in Figure 1). The collocated ONU/BS1 and ONU/BS2 may rely on centralized RoF technologies such as the aforementioned C-RAN. Conversely, the third subset of remaining ONUs is equipped with a mesh portal point (MPP) to interface with the WiFi mesh network consisting of decentralized mesh points (MPs) and mesh access points (MAPs), each serving mobile users within their limited coverage area in a decentralized fashion. The collocated ONU/MPPs are realized by using R&F technologies.

By augmenting 4G LTE-A HetNets with FiWi access networks based on decentralized Ethernet technologies such as next-generation IEEE 802.11n WLAN, large amounts of

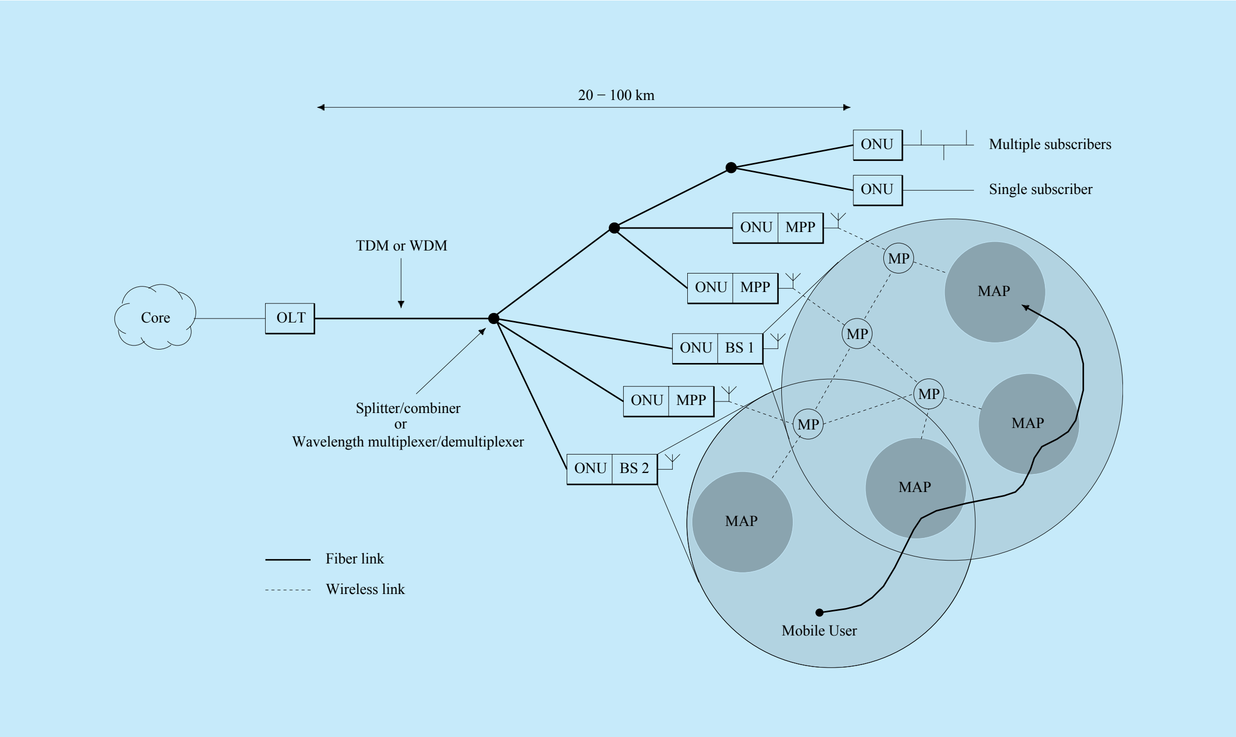

Fig. 1. *FiWi enhanced LTE-A HetNet architecture for WiFi offloading.*

mobile data traffic can be offloaded onto WiFi. For illustration, Figure 2 shows the achievable *offloading efficiency* (i.e., ratio of bytes transferred through WiFi and total number of bytes generated by mobile users) for both on-the-spot and delayed offloading. Unlike on-the-spot offloading, delayed offloading has a non-zero delay tolerance to complete the WiFi offloading within a certain offloading deadline D (set to 2, 20, and 120 minutes in Figure 2). For an offloading deadline of 120 minutes, we observe from Figure 2 that an offloading efficiency of 100% (i.e., all mobile data traffic is offloaded from LTE-A onto WiFi) can be achieved, provided that mobile users are connected to FiWi with a probability of 0.5 or higher.

Beside WiFi offloading, which represents a key aspect of the strategy of today's operators to offload mobile data traffic from their cellular networks [17], complementing fast evolving LTE-A HetNets with collocated ONU/MPPs inherently provides the opportunity to implement decentralized cloudlets, as discussed next.

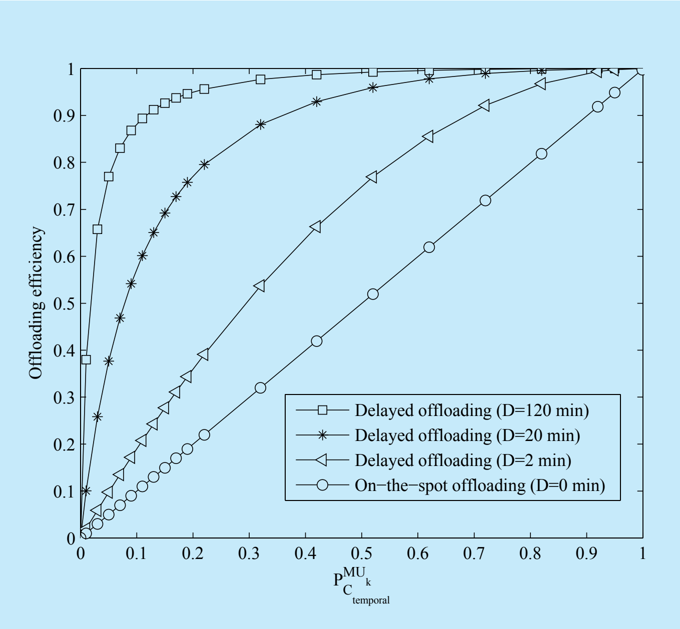

Fig.2. *Offloading efficiency vs. temporal FiWi connectivity probability of mobile users for on-the-spot and delayed offloading.*

III. Cloud and Cloudlet Empowered FiWi-HetNets: Computation Offloading

Our previous observations motivate us to

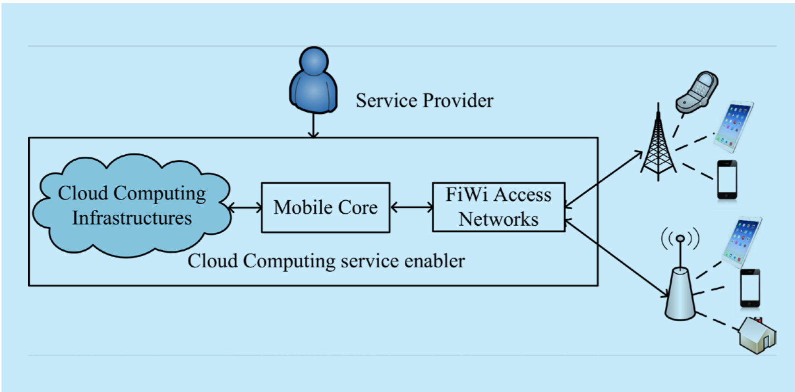

Fig. 3 *Cloud computing in FiWi enhanced LTE-A HetNets.*

rethink mobile network architectures by exploiting cloud computing and extending the concept of cloud computing beyond data centers towards mobile end-users across FiWi access networks, as illustrated in Figure 3.

Figure 4 depicts the considered architecture of cloud and cloudlet empowered FiWi-HetNets in greater detail. Cloud services are provided over a PON based FTTH network. The integrated ONU/eNBs are connected to cloudlets by optical fibers. For instance, cloudlet 1 is connected to fixed users (ONU4) as well as cellular users (ONU1/BS1), while cloudlet 2 provides cloud services to cellular users (ONU2/BS2 and ONU3/BS3). The ONU/eNBs and OLT exchange control and management messages by using the multipoint control protocol (MPCP) defined in IEEE 802.3ah. Note that the cloudlet placement may have a significant impact on the performance and user experience. Thus, the cloudlet network planning represents one of the important design issues. Proactive placement may be one promising approach among others, where the network planner analyzes the traffic and mobility history of mobile users and thereby builds a model for placing cloudlets.

As shown in Figure 4, the evolved packet core (EPC) is located at the central office. The serving gateway (S-GW) and mobility management entity (MME) are connected to the OLT via the S1 interface. The S-GW is connected to the MME through the S11 interface. The home subscriber server (HSS) and MME are connected via the S6a interface. The MME controls the high-level operation of the mobile network via signaling messages and HSS. The HSS is a database that contains information about the network operator's subscribers. The PDN-GW communicates with the outside world (Internet) using the SGi interface. The S-GW acts as a router and forwards data between the OLT and PDN-GW, whereby the OLT accesses the remote cloud infrastructures via the EPC.

The usability of mobile terminals can be expanded beyond their physical limits and their battery charging intervals can be

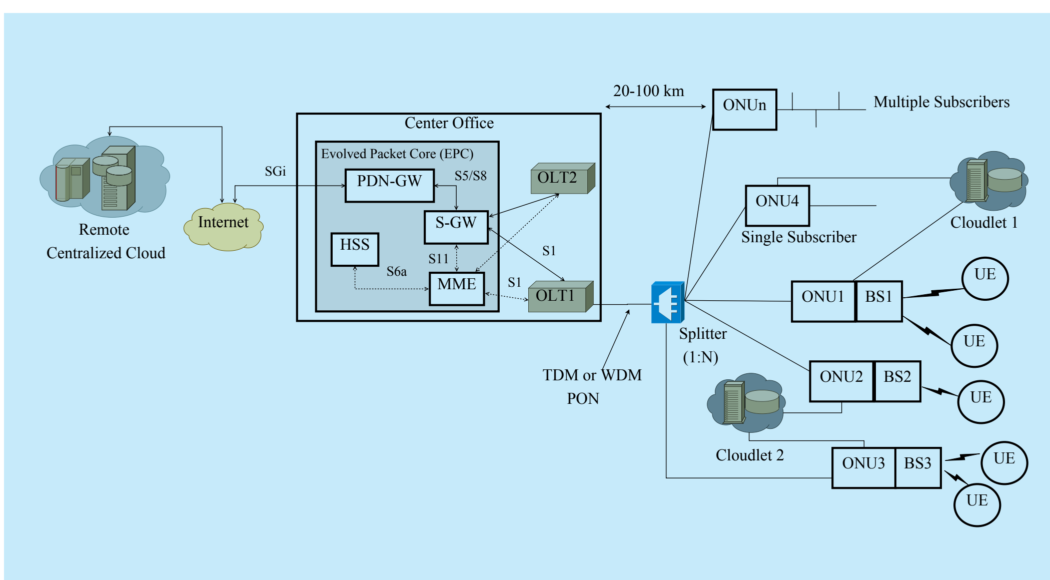

Fig. 4 *Cloud and cloudlet empowered FiWi-HetNet architecture.*

greatly prolonged by means of *computation offloading* whenever a computation-intensive task is not affordable or doable by local computing resources. Computation offloading enables the migration of computation to more resourceful servers. It is sometimes referred to as "surrogate computing" or "cyber foraging." The offloading decisions are usually made by analyzing several factors such as bandwidth, available computing resources, wireless channel conditions, etc. With computation offloading, mobile terminals access Internet connected cloud resources via WiFi or cellular networks and coordinate with application servers to locally decide to offload tasks to the cloud computing infrastructure according to given offloading strategies [18].

Computation offloading has received a significant amount of attention by both industry and academia. For instance, the authors in [18] provided an overview of computation offloading solutions to the cloud and studied how mobile applications can be enhanced with cloud services in order to achieve energy savings and enhanced performance for two different case studies, namely, cloud-assisted distributed mobile applications and mobile video compression applications. Cuervo et al. [19] proposed MAUI (Mobile Assistance Using Infrastructure) to enable fine-grained energy-aware dynamic offloading of mobile code to the remote infrastructure by combining a profiler with an integer linear programming (ILP) solver. MAUI does not support remotely executing virtualized methods called native functions (e.g., two methods that share the native state). Furthermore, MAUI requires the solver to be running on the server at runtime and programmers to annotate methods as remotable. Similarly, in [20], the authors developed the so-called CloneCloud based on offline application partitioning and migration at the method granularity. A mathematical optimizer was designed to choose the migration point that uses static analysis and optimizes mobile-device energy consumption subject to given execution time constraints. At the downside, it does not consider concurrency, trustworthiness of the clone, and access to native resources that are not yet virtualized and are not available on the clone. CloneCloud was extended in [21] to capture the existing call stack with heap objects, which reduces the transferred data size by transferring only the essential heap objects and the stack frames actually referenced in the server.

The shortcomings of MAUI and CloneCloud (e.g., offloading of only one method/thread at a time and locking issues) were addressed in ThinkAir [22]. ThinkAir supports on-demand resource allocation and exploits parallelism by dynamically creating, resuming, and destroying virtual machines (VMs) in the cloud when needed. The authors of [23] studied the feasibility and cost of off-clones (offload computation on the fly) and back-clones (clones that are used to backup user data) in terms of bandwidth and energy consumption. A logger application was developed to track the events' occurring in the devices.

The integration of cloud computing infrastructures (e.g., remote clouds and cloudlets) can be realized as Infrastructure as a Service (IaaS) in FiWi access networks. The envisioned architecture leverages local computing resources. Such distributed resources may be cloudlets. The local clouds periodically synchronize to a remote cloud for service availability. In addition, multiple local clouds may be deployed in a distributed fashion within the coverage area of FiWi access networks to improve the cloud accessibility for users. Moreover, empowering FiWi enhanced LTE-A HetNets with clouds and cloudlets may offer a broader range of benefits such as improved performance, energy saving, cost minimization, scalability, and elasticity by means of computation offloading.

IV. High-Performance AWG-Based Switchless WDM Network: Scalability and East-West Flows

Optical interconnection networks for cloud data centers can be realized by using one or more of the optical switching techniques mentioned in Section 1. Most of these optical switching techniques (except OCS) have one key challenge in common: contention resolution in the optical domain. While contention resolution can be easily realized by means of electronic random access memory (RAM), there exists no comparable buffer solution in optical networks. Instead, bulky fiber delay lines (FDLs) have to be used, which provide fixed and limited delays to resolve contention of two or more incoming packets/bursts/flows competing for the same output port of an intermediate optical switch.

This inherent shortcoming of optical switching techniques can be avoided by designing *bufferless* optical switches. A recent example of a bufferless optical switch for data center networks is the petabit optical switch architecture proposed in [24] for interconnecting top-of-rack (ToR) switches, which relies on a bufferless three-stage optical switch fabric based on interconnected arrayed-waveguide gratings (AWGs) at the core and tunable wavelength converters (TWCs) at the input modules. The AWGs are used as wavelength multiplexers and demultiplexers, while the TWCs are needed to perform wavelength conversion and thus enable dynamic configuration of the switch fabric, given that AWGs are completely passive devices that are not reconfigurable per se. At the downside, this bufferless switch does not include any splitters and thus does not support multicasting. Due to its lack of multicasting and broadcasting capabilities, the switch requires centralized control. Furthermore, the proposed petabit optical switch does not exploit the frequency-cyclic nature of AWGs for improved scalability.

Recall from Section 1 that beside capacity, low latency, low complexity and low power consumption, optical cloud network solutions have to address the following two data-center specific design challenges: (*i*) *scalability* and (*ii*) *east-west flows* between ToR switches. In the following, we will revisit our high-performance *switchless* single-hop WDM network [25], which not only inherently provides bit/protocol transparency for easy upgradability to higher data rates and/or new protocols and packet formats, but more importantly, exploits multiple FSRs of the underlying AWG to allow for spatial reuse of all WDM wavelengths at each input port for a significantly improved degree of concurrency and scalability. In addition, the switchless single-hop WDM network supports data multicasting and control broadcasting, which in turn paves the way for *distributed* MAC protocols. Despite the fact that the basic ideas behind this architecture were developed in a different context back in 2000, we believe that the above unique merits render our switchless single-hop WDM network using multiple FSRs of an AWG a promising solution for optical data center interconnection networks, as explained in greater detail next.

Similarly to [24], tunable transmitters are used to reach different destinations by simply changing the wavelength. In doing so, the switching functionality is naturally moved towards the network periphery, resulting in significantly reduced costs and complexity and simplified network management. With tunable transceivers we are able to realize switchless single-hop networks. Single-hop networks have some very desirable properties such as minimum hop distance (unity), high-channel utilization due to the lack of any traffic forwarding burden, inherent transparency, and low processing requirements at each node. The network and node architecture is depicted in Figure 5. The network is based on a $D\times D$ AWG. At each AWG input port a wavelength-insensitive $S\times 1$ combiner is attached. Similarly, at each AWG output port signals are distributed by a wavelength-insensitive $1\times S$ splitter. Each node is composed of a transmitting and receiving part. The transmitting part of a node is attached to one of the combiner ports. The receiving part of the same node is located at the opposite splitter port.

The network connects N nodes, with $N = D\times S$. For a given number of nodes N there are several possible network configurations with different values of D and S. The choice of D and S trades off spatial wavelength reuse and receiver throughput. Due to the wavelength routing characteristics of the AWG each wavelength can be used at all ports simultaneously (spatial wavelength reuse). Spatial wavelength reuse increases the degree of concurrency, resulting in an improved throughput-delay performance. Therefore, from the spectrum reuse point of view it is

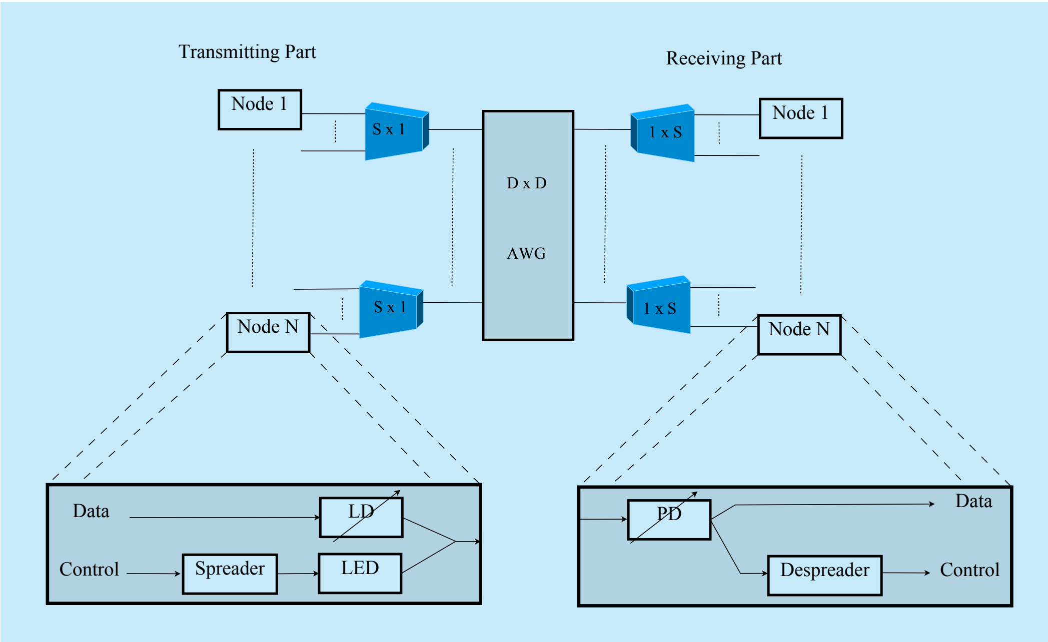

Fig. 5 *Scalable switchless single-hop WDM network using multiple FSRs of a wavelength-routing AWG for east-west flows in data centers.*

reasonable to choose a large *D* for a given *N*. On the other hand, small values of *D* imply that many receivers are attached to the same splitter, i.e., *S* becomes large. This has the advantage that each transmitted packet can be received by more nodes, translating into an increased receiver throughput. An increased receiver throughput allows for efficient multicasting since multicast packets have to be transmitted fewer times.

Let us now take a look at the node structure. Each node contains a laser diode (LD) for transmission and a photodiode (PD) for reception. Given the wavelength routing characteristics of the AWG, both transmitter and receiver have to be tunable over at least *D* wavelengths in order to provide full connectivity. In addition, each node uses a light emitting diode (LED) for broadcasting control packets. The broadband LED signal (10-100 nm) is spectrally sliced such that all receivers are able to obtain the control information. No additional receiver is required if the signaling is done in-band, i.e., LED and LD signals overlap spectrally. However, data and control information have to be distinguishable at the receiver. This can be achieved by code division multiple access (CDMA). The control information is spreaded before modulating the LED. Accordingly, at the receiving part the control information is retrieved by despreading a part of the incoming signal. Each node receives the control signals of all other *N*-1 nodes. As a consequence, all nodes have global knowledge at any time and by running a distributed MAC protocol and executing a common distributed deterministic scheduling algorithm network resources are used efficiently and packet collisions are avoided completely (see [25] for details).

Intrachannel crosstalk due to spatial wavelength reuse has to be taken into account. As a consequence, the AWG has to be realized as a free-space device or integrated AWGs with a limited physical degree can be deployed. Note, however, that for a fixed channel spacing and transceiver tuning range AWGs with a limited physical degree allow for the use of multiple FSRs, resulting in an increased number of channels between each AWG input/output pair.

Another crucial issue is the small bandwidth-distance product of LEDs even though they are used for transmitting only low-rate control information (and the distance between ToR switches in data centers is limited). Especially for a large number of nodes the splitting losses due to the combiners and splitters put severe constraints on the power budget. Those constraints can be relaxed by inserting Erbium-doped fiber amplifiers (EDFAs) between each combiner/splitter and the corresponding AWG port. Since the physical degree of the AWG is limited only a few EDFAs would be required. Alternatively, each LED signal could be preamplified or other broadband light sources such as fiber amplifiers or Fabry-Perot lasers driven into clipping could be used instead of LEDs. However, those solutions are either not very economic or support only small transmission rates. The most promising approach appears to be the use of superluminescent diodes (SLDs), which provide a significantly improved power budget.

In summary, the AWG based switchless WDM network offers a very high degree of concurrency. All wavelengths are used for data transmission and in-band signaling is deployed, i.e., no additional control channel is required. The node structure is simple and economical, while the network itself consists of passive components. The scalability of the network is significantly increased by using multiple FSRs of the AWG and spatially reusing all wavelengths at each input/output port, leading to a significantly improved throughput-delay performance of the network, as analytical results have shown in [25]. Arguably more importantly, the single-hop network architecture of Figure 5 is particularly suitable to support the typical east-west flows between racks that are dominating data center traffic.

V. Conclusion

We have highlighted our recently proposed concept of FiWi enhanced LTE-A HetNets to illustrate the joint use of RoF and R&F technologies in response to the unprecedented growth of mobile data traffic. We have shown that a mobile data offloading efficiency of 100% can be achieved for delay-tolerant traffic, provided that mobile users are connected to FiWi with a probability of 0.5 or higher. FiWi enhanced LTE-A HetNets inherently provide the opportunity to implement decentralized cloudlets. The integration of cloud computing infrastructures (e.g., remote cloudlets) can be realized as IaaS in FiWi access networks, thereby providing numerous benefits such as improved performance, energy saving, cost minimization, scalability, and elasticity by means of computation offloading.

As for intra-data center communications, we argued that there exists a very profound understanding of the merits and shortcomings of different optical switching techniques obtained from decades of research on optical WDM networks that may be adapted to data centers in a rather straightforward fashion, as witnessed by recently proposed data center network solutions based on well-known optical switching techniques such as OCS, OBS, and OPS, or a combination thereof. Beside capacity, low latency, low complexity, and low power consumption, optical cloud network solutions have to address the two data-center specific design challenges of scalability and east-west flows between ToR switches. Toward this end, we have revisited our high-performance switchless single-hop WDM network based on completely passive wavelength-splitting/routing optical components, which not only inherently provides bit/protocol transparency for easy upgradability to higher data rates and/or new protocols and packet formats, but more importantly, exploits multiple FSRs of the underlying AWG and allows for spatial reuse of all WDM wavelengths at each input port for a significantly improved degree of concurrency and throughput-delay performance as well as an enhanced scalability. Furthermore, the single-hop architecture of the network is particularly suitable to support the typical east-west flows between racks that are dominating data center traffic.

References

[1] M. Maier, N. Ghazisaidi, and M. Reisslein, "The Audacity of Fiber-Wireless (FiWi) Networks," Proc., ICST International Conference on Access Networks (AccessNets), pp. 1-10, Las Vegas, NV, USA, Oct. 2008.

[2] M. Satyanarayanan, R. Schuster, M. Ebling, G.

Fettweis, H. Flinck, K. Joshi, and K. Sabnani, “An Open Ecosystem for Mobile-Cloud Convergence,” IEEE Communications Magazine, vol. 53, no. 3, pp. 63-70, March 2015.

[3] G. Hu, W. P. Tay, and Y. Wen, “Cloud Robotics: Architecture, Challenges and Applications,” IEEE Network, vol. 26, no. 2, pp. 21-27, May/June 2012.

[4] M. Satyanarayanan, P. Bahl, R. Cáceres, and N. Davies, “The Case for VM-Based Cloudlets in Mobile Computing,” IEEE Pervasive Computing, vol. 8, no. 4, pp. 14-23, Oct.-Dec. 2009.

[5] Nokia Networks, “Intelligent base stations,” White Paper, 2014.

[6] Intel Corporation, “Increasing Mobile Operators’ Value Proposition With Edge Computing,” Technical Brief, 2013.

[7] Mobile-Edge Computing (MEC) Industry Initiative, “Mobile-Edge Computing,” Introductory Technical White Paper, Sept. 2014.

[8] Intel Corporation, “Increasing Network ROI with Cloud Computing at the Edge,” Solution Brief, 2014.

[9] C. F. Lam, H. Liu, B. Koley, X. Zhao, V. Kamalov, and V. Gill, “Fiber Optic Communication Technologies: What’s Needed for Datacenter Network Operations,” IEEE Communications Magazine, vol. 48, no. 7, pp. 32-39, July 2010.

[10] G. Wang, D. G. Andersen, M. Kaminsky, K. Papagiannaki, T. S. Eugene Ng, Michael Kozuch, and M. Ryan, “c-Through: Part-time Optics in Data Centers,” Proc., ACM SIGCOMM, pp. 327-338, New Delhi, India, Aug./Sept. 2010.

[11] N. Farrington, G. Porter, S. Radhakrishnan, H. H. Bazzaz, V. Subramanya, Y. Fainman, G. Papen, and A. Vahdat, “Helios: A Hybrid Electrical/Optical Switch Architecture for Modular Data Centers,” Proc., ACM SIGCOMM, pp. 339-350, New Delhi, India, Aug./Sept. 2010.

[12] M. Fiorani, S. Aleksic, and M. Casoni, “Hybrid Optical Switching for Data Center Networks,” Hindawi Journal of Electrical and Computer Engineering, vol. 2014, pp. 1-13, Feb. 2014.

[13] M. Maier and M. Reisslein, “Trends in Optical Switching Techniques: A Short Survey,” IEEE Network, vol. 22, no. 6, pp. 42-47, Nov./Dec. 2008.

[14] China Mobile Research Institute, “C-RAN: The Road Towards Green RAN,” White Paper, Oct. 2011.

[15] D. Tardent and N. Farrington, “Building a datacenter network with optimal performance and economy,” Lightwave Magazine, July 2014 [Online].

[16] H. Beyranvand, M. Lévesque, M. Maier, and J. A. Salehi, “FiWi Enhanced LTE-A HetNets with Unreliable Fiber Backhaul Sharing and WiFi Offloading,” Proc., IEEE INFOCOM, Hong Kong, April/May 2015.

[17] S. Singh, H. S. Dhillon, and J. G. Andrews, “Offloading in Heterogeneous Networks: Modeling, Analysis, and Design Insights,” IEEE Transactions on Wireless Communications, vol. 12, no. 5, pp. 2484-2497, May 2013.

[18] X. Ma, Y. Zhao, L. Zhang, H. Wang, and L. Peng, “When Mobile Terminals Meet the Cloud: Computation Offloading as the Bridge,” IEEE Network, vol. 27, no. 5, pp. 28-33, Sept. 2013.

[19] E. Cuervo, A. Balasubramanian, D.-K. Cho, A. Wolman, S. Saroiu, R. Chandra, and P. Bahl, “MAUI: Making Smartphones Last Longer with Code Offload,” Proc., ACM MobiSys, pp. 49-62, June 2010.

[20] B.-G. Chun, S. Ihm, P. Maniatis, M. Naik, and A. Patti, “CloneCloud: Elastic Execution between Mobile Device and Cloud,” Proc., ACM EuroSys, pp. 301-314, April 2011.

[21] S. Yang, D. Kwon, H. Yi, Y. Cho, Y. Kwon, and Y. Paek, “Techniques to Minimize State Transfer Costs for Dynamic Execution Offloading in Mobile Cloud Computing,” IEEE Transactions on Mobile Computing, vol. 13, no. 11, pp. 2648-2660, Nov. 2014.

[22] S. Kosta, A. Aucinas, P. Hui, R. Mortier, and X. Zhang, “ThinkAir: Dynamic Resource Allocation and Parallel Execution in the Cloud for Mobile Code Offloading,” Proc., IEEE INFOCOM, pp. 945-953, March 2012.

[23] M. V. Barbera, S. Kosta, A. Mei, and J. Stefa, “To Offload or Not to Offload? The Bandwidth and Energy Costs of Mobile Cloud Computing,” Proc., IEEE INFOCOM, pp. 1285-1293, April 2013.

[24] K. Xi, Y.-H. Kao, M. Yang, and H. J. Chao, “Petabit Optical Switch for Data Center Networks,” Technical Report, Polytechnic Institute of New York University, 2010.

[25] M. Maier, M. Reisslein, and A. Wolisz, “High-performance Switchless WDM Network Using Multiple Free Spectral Ranges of an Arrayed-Waveguide Grating (Best Paper Award),” Proc., SPIE Photonics East Terabit Optical Networking: Architecture, Control, and Management Issues, vol. 4213, pp. 101-112, Boston, MA, USA, Nov. 2000.

Biographies

Martin Maier is a full professor with the Institute National de la Recherche Scientifique (INRS), Montréal, Canada. He was educated at the Technical University of Berlin, Germany, and received MSc and PhD degrees (both with distinctions) in 1998 and 2003, respectively. In the summer of 2003, he was a postdoc fellow at the Massachusetts Institute of Technology (MIT), Cambridge. He was a visiting professor at Stanford University, Stanford, October 2006 through March 2007. Dr. Maier is a co-recipient of the 2009 IEEE Communications Society Best Tutorial Paper Award and Best Paper Award presented at The International Society of Optical Engineers (SPIE) Photonics East 2000-Terabit Optical Networking Conference. He is the founder and creative director of the Optical Zeitgeist Laboratory (www.zeitgeistlab.ca). He is the author of the book "Optical Switching Networks" (Cambridge University Press, 2008), which was translated into Japanese in 2009, and the lead author of the book "FiWi Access Networks" (Cambridge University Press, 2012).

Bhaskar Prasad Rimal received the M.Sc. degree in information systems from Kookmin University, Seoul, Korea in 2010. He is currently pursuing the Ph.D. degree in telecommunications at the Optical Zeitgeist Laboratory, Institut National de la Recherche Scientifique (INRS), Montréal, QC, Canada. Mr. Rimal is a recipient of the Doctoral Research Scholarship from the Quebec Merit Scholarship Program for foreign students of Fonds de Recherche du Québec-Nature et Technologies (FRQNT), the Korean Government Information Technology (IT) Fellowship, the Kookmin University IT Scholarship, and the Kookmin Excellence Award as an Excellent Role Model Fellow. His current research interests include cloud computing, mobile-edge computing, 4G/5G wireless communications, smart grids, and game theory.