

# TECHNIQUES FOR OFFLOADING LTE EVOLVED PACKET CORE TRAFFIC USING OPENFLOW: A COMPARATIVE SURVEY & DESIGN REFERENCE

Kam Nasim<sup>1</sup> and Trevor James Hall<sup>2</sup>

<sup>1</sup>Department of Electrical and Computer Engineering, University of Ottawa, Canada

<sup>2</sup>Professor and Canada Research Chair in Photonic Network Technology, University of Ottawa, Canada

## ABSTRACT

*Cellular users of today have an insatiable appetite for bandwidth and data. Data-intensive applications, such as video on demand, online gaming and video conferencing, have gained prominence. This, coupled with recent innovations in the mobile network such as LTE/4G, poses a unique challenge to network operators in how to extract the most value from their deployments while reducing their Total Cost of Operations (TCO). To this end, a number of enhancements have been proposed to the “conventional” LTE mobile network. Most of these recognize the monolithic and non-elastic nature of the mobile backend and propose complementing core functionality with concepts borrowed from Software Defined Networking (SDN). In this paper, we will attempt to explore some existing options within the LTE standard to address traffic challenges. We then survey some SDN-enabled alternatives and comment on their merits and drawbacks.*

## KEYWORDS

*SDN, LTE, EPC, openflow, mobile networks, deep packet inspection, network applications.*

## 1. INTRODUCTION

With recent, significant strides being made in the realm of virtualization, cloud services, and mobile computing, network planners and Mobile Network Operators (MNOs) are being led to re-examine and reconsider traditional networking architectures. More than anything else, this re-focus is driven by the changing face of network applications. It is estimated that mobile gaming is growing at an annual rate of 19% grossing \$13.9 billion worth of revenue annually [2]. In their annual Visual Networking Index report on Global Mobile data forecasts, Cisco stipulates that globally, mobile IP video traffic will hold a mammoth 79% share of all consumer Internet traffic by 2018, up from 66% in 2013 [11]. They revised this metric in 2017 to forecast that by 2021 up to 86% of total Internet traffic would be video and multimedia content [12]. Innovation within the mobile network itself is considered as a key driver of this next generation of network applications. Compared to the previous 3G standard, LTE is:

- A faster radio access technology. Providing a peak downlink of 100Mb/s and Peak Uplink of 50Mb/s with speeds being even higher with LTE-Advanced;
- A flat, all IP based backhaul network (AIPN) which reduces total cost of ownership and provides less complex network and node deployments; and,
- 3 to 5 times lower latency which facilitates a better user experience and enables more time-sensitive and mission critical applications.

Another key facilitator is the global proliferation of data centers and cloud computing. The Cisco Global Cloud Index [20] stipulates that by 2019, more than 86% of workloads will be processed by cloud data centers. Global cloud IP traffic will more than quadruple over the next 5 years, reaching 719 Exabytes per month.

However despite all of these facilitators, MNOs are stuck between a rock and a hard place: they must continue to meet expectations of the “vanilla” consumer, who simply wants to use his or her device to make phone calls and send messages, and yet also provide a network that can accommodate this new breed of data hungry consumers.

The original LTE network was not designed for this level of service differentiation. Yes we have QoS and Evolved Packet System (EPS) bearer configurations but they have more to do with resource allocation (such as in the case of GBR bearers) and prioritization on shared channels - the packet still traverses the same network. Given the centralized and non-elastic nature of nodes in the LTE core, it is not convenient for MNOs to reroute traffic on-demand to address network congestion issues.

Existing research argues that Software-Defined Networks (SDN) may be able to overcome some of the drawbacks of LTE/EPC mobile networks by allowing network designers to have a logically centralized *controller* that can maintain a global view of the network topology and provide control elasticity. OpenFlow rules can be applied to redirect traffic to special nodes and thereby alleviate contention in different parts of the mobile network. The remainder of this paper will focus on discussing this synergy in detail.

This paper is structured as follows: In order to set the stage for an in-depth discussion on Software Defined Mobile Network (SDMN) architectures, this section highlights a key technological facilitator pertinent to our discussion. Section two outlines the existing traffic offloading techniques in the LTE core network. The next section surveys similar techniques of offloading available in SDN empowered mobile networks. We will follow this up with a reference implementation of an Agnostic Traffic Offloading Service that we refer to as AETOS. The section that follows outlines AETOS’s performance and offloading fairness constraints. Section seven implements the AETOS reference design in a simulation test bench which offers some key insights into its operational benefits. Finally, we conclude with a wish list of possible future enhancements to AETOS.

## 1.1 NETWORK FUNCTION VIRTUALIZATION

Network Function Virtualization (NFV) as it is more commonly known, has been heralded as the next big thing in the networking world. Simply put, NFV is the principle of moving host-based network functions and applications from dedicated hardware onto virtual instances running on commercial off-the-shelf (COTS) servers either on-premise or in the cloud [9]. The key driving force behind NFV is its ability to scale capacity on-demand. Setting up a new network node is as simple as spinning a new virtual instance without going through lengthy qualification cycles or steep learning curves as is characteristic of new hardware deployments. MNOs have wholly embraced this concept since it allows them to scale their infrastructure during peak demands or certain seasonal trends, without permanently investing in hardware and rack space on-premise. Another major leg-up in NFV’s adoption is vendor neutrality: by moving away from proprietary network hardware, MNOs are no longer locked in with a given vendor and can build an eclectic network; one that maximizes their shareholder value without compromising their customer experience. Given that most data centers hosting cloud networks are multi-tenant; infrastructure and administrative staff and utility costs are split and shared, thereby reducing operational expenditures (OPEX) for each of these MNOs, in comparison

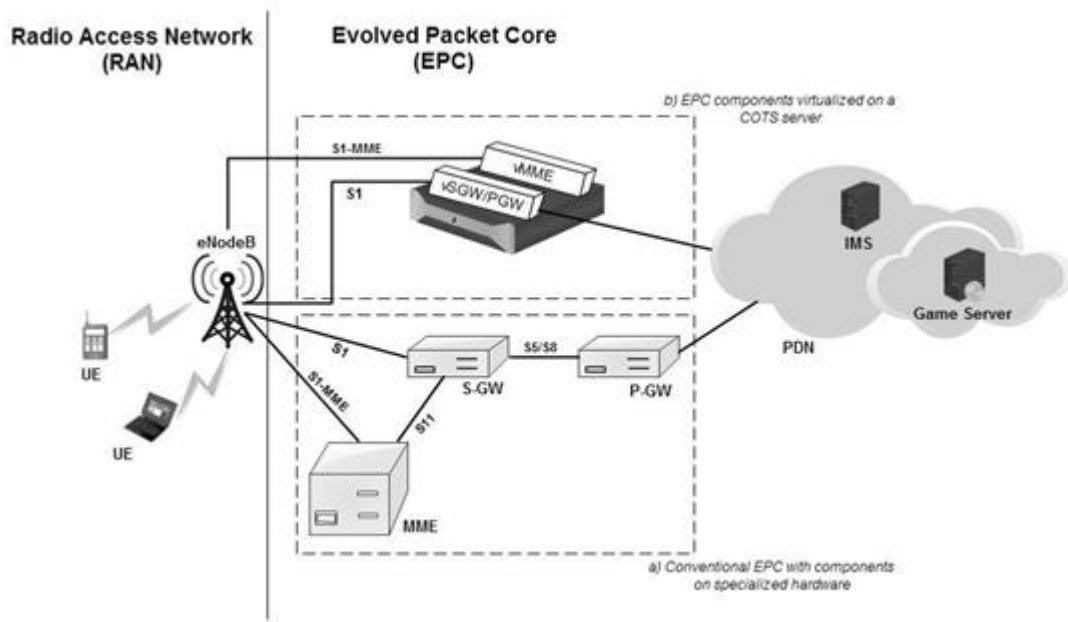


Figure 1: Author depiction of an LTE network topology illustrating both conventional Evolved Packet Core (EPC) and a Network Function Virtualization (NFV) based EPC with the OPEX incurred, had they hosted these network functions on-premise.

Figure 1 illustrates a vanilla LTE-EPC architecture, with core components deployed over dedicated hardware, and how it would look like if said components were virtualized on a commercial off-the-shelf (COTS) server.

For the remainder of this article, our overarching goal is to expose designs that leverage SDN to dynamically steer and offload EPC traffic to either a fixed IP network or to a cloud based Gateway, either as means for the MNO to address surges in user data traffic, or as a facility available to the end user (the UE) to pay a premium for this service.

## 2. EXISTING OFFLOADING TECHNIQUES IN EPC

The concept of traffic offloading has its proven benefits and is well noted within the LTE standard. In an earlier study was done in 2015 it was estimated that by 2017, more than 30% of core network traffic may be offloaded to fixed IP networks [14]. In Cisco's annual Global Data Report [12], it was indicated that 60% of total mobile data traffic was offloaded onto the fixed network through Wi-Fi or otherwise, and a total 10.7 exabytes of mobile data traffic was offloaded onto the fixed network each month. Indeed both 3GPP Release 10 and Release 11 specifications have proposed various offloading techniques, to address contention in both EUTRAN and the EPC core network. Suffice to say, and without embarking into too much detail, Wi-Fi based offloading is the de facto technique for congestion alleviation with the radio access network. To this end, Access Network Discovery and Selection Function (ANDSF) was defined in recent 3GPP standards as a framework for issuing policies to a device where traffic routing decisions between a cellular and WiFi access network are being made [3].

However, for traffic offloading within the EPC core network, the following are some standardized techniques available to MNOs [22]:

## 2.1 LOCAL IP ACCESS (LIPA)

LIPA based offloading is only usable on LTE femtocells, or Home EnodeBs (HeNBs) as they are also known. These are limited bandwidth, low power transmitters meant for home or campus deployments. A local gateway (L-GW), is connected to the HeNB and allows UEs, attached to this HeNB, the ability to bypass the core network and access certain services over the local fixed IP network.

LIPA makes sense where the MNO provides its subscribers certain geo-location based or intranet services, for which traversing the core network is inefficient. LIPA posed an interesting business case for Over-the-top (OTT) players such as advertisers and content providers to offer specialized value-added content to subscribers using their femtocells.

Promising as it might sound, LIPA has its limitations. For one thing, the facility is not available on regular enodeBs, i.e. macro cells. Additionally, LIPA provides user congestion alleviation at the expense of increased signaling in the EPC core network. The initial UE message to the Mobility Management Entity (MME) must now include an information element to indicate the IP address of the L-GW. Also all paging requests from the L-GW to the S-GW are forwarded to the MME and must be serviced there. This additional signaling must be done for every L-GW, which might be as numerous as the number of HeNBs in the MNO's Radio Access Network (RAN), in the event that L-GWs are co-located within the femtocells.

We posit that this is one area that would benefit from a logically centralized SDN control, that takes away compute complexity from the HeNBs in setting up the L-S5 tunnels and L-GW IP Addresses and does so globally or at the very least for a cluster of geographically close femtocells.

## 2.2 SELECTED IP TRAFFIC OFFLOAD (SIPTO)

Unlike LIPA, SIPTO can be used on both LTE macrocells and femtocells. SIPTO involves the MME selecting a (S-GW, P-GW) pair geographically close to the UE's point of attachment in the RAN thereby minimizing the UE's round-trip-time (RTT). We would like to point out that here lies a key distinction between LIPA and SIPTO. In the case of the former, the decision to offload is made by the HeNB which may or may not be owned by the same operator as the core network. Whereas for the latter, this decision is centralized at the MME. SIPTO leverages this centralized selection to provide service continuity to its UE subscribers during mobility and handover. However, similar to LIPA, relief in the user plane comes at the cost of increased signaling within the core network. In addition to the initial gateway selection in SIPTO, the MME needs to monitor UE tracking area updates and initiate explicit detach followed by a reattach procedure for gateways that end up being closer to the UE.

It is our perception that SIPTO is one area which would also benefit from existing SDNMN research and solutions that propose moving MME functions within a logically centralized control. The additional signaling involved in SIPTO is further exacerbated by increased east-west communication between MME nodes within an MME pool when a UE context (SIPTO subscribed) is transferred from one MME to another. Having a global state machine to issue SIPTO policies and make decisions would cut through some of these east-west exchanges.

Table 1: LIPA vs. SIPTO vs. S1 flex comparison table

Characteristics	LIPA	SIPTO	S1 Flex
<i>RAN limitations</i>	not applicable for LTE macrocells	no limitations	no limitations
<i>UE mobility support</i>	no mobility support	supports mobility	supports mobility
<i>Traffic offloading traffic</i>	Home EnodeB	MME	MME/EnodeB
<i>Type of traffic of- loaded</i>	user-plane	user-plane	control-plane
<i>Perceived Benefits through SDN</i>	logically centralizing L5S tunnel creation and L-GW IP address allocation will reduce processing complexity at the HeNB and promote mobility	improves MME's SGW/P-GW selection function and improve UE context transfers across MMEs	reduces the number of concurrent S1MME connections and reduces process complexity at the enodeB

### 2.3 S1-FLEX

S1-flex is an LTE standard concept that allows MNOs to create a pool of MMEs. Individual eNodeBs can then have multiple concurrent S1 connections to the MMEs within the pool as opposed to a single S1-MME link as is the conventional case. Although S1-flex is not a “traffic” offloading technique per se, and some would even argue that it is akin to failover redundancy, it does provide operators the ability to offload EPC signaling traffic to a bypass core network which offers up its MME within the common MME pool. It is also for this reason that S1-flex is a popular deployment option for multi-operator RAN sharing where MNOs contribute to an MME pool and have macro cells multi-homed onto this pool.

The technology offers a number of merits to MNOs. Sudden imbalances in control traffic or over-commitment of an MME node can be addressed without detaching the enodeB or losing UE context. S1-flex allows for seamless migration and load equalization with the MME pools. Secondly, since UE context is shared between the MME nodes of a pool, it reduces the number of tracking area update exchanges between the MME and Home Subscriber Service (HSS). If a UE moves within its pool area, call processing can still be done without registering the new location update with the HSS.

It is worth mentioning that while LIPA and SIPTO focused on user traffic offloading, S1flex is a pure control plane offering. However, the mechanism does have some perceived shortcomings. For one, the flat distributed architecture of LTE would imply a M:N mesh of S1-MME connections. The guarantee of failover protection and control redundancy comes at the expense of significantly more chatter on the wire and within the operator’s backhaul network.

Table 1 summarizes some of the characteristics of these three offloading techniques. All mechanisms discussed above have their merits and are prevalent in the industry. However they are not mutually exclusive. MNOs may opt to have all three schemes present in their networks. The next section, however, will survey some SDN based offloading solutions that truly challenge the status quo.

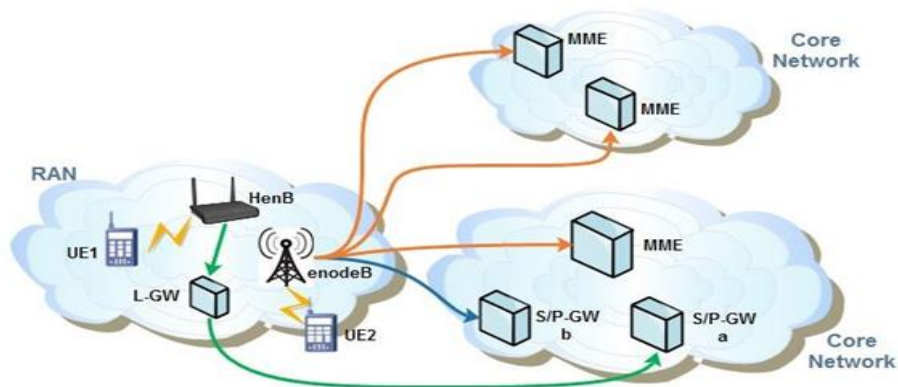


Figure 2: Deployment depicting LIPA, SIPTO and S1-flex traffic offloading approaches. The arrows in Green represent LIPA where-in HenB interfaces with L-GW. The workflow in Blue represent the fact that UE2 is SIPTO subscribed and while normally its enodeB should have connected to S/P-GW<sub>a</sub>, in this case the MME selects the closer S/P-GW<sub>b</sub>. The flow in Orange represents S1 Flex where the enodeB connects to an MME pool.

### 3. EPC OFFLOADING TECHNIQUES THAT LEVERAGE SDN CONCEPTS

In general, traffic offloading in the LTE network is not a novel concept. However, as is usually the case with performance optimizations, offloading in the EPC core network is an after-thought and suffers from some of the same ailments as the LTE network. For one, performance gains in the user plane come at the expense of an increase in signaling complexity. Owing to the centralized nature of gateway nodes within the core network, and usually, at the network edge, UE flows have to traverse large distances to reach web servers that may even be located geographically close to them. Having a controller with a global network view can improve this suboptimal routing by improving coordination between MME nodes. In addition, with NFV technology, virtualized gateway nodes can be spawned gratuitously nearer to the RAN.

The following subsections will review some state-of-the-art works within the confines of EPC based offloading and highlight some common themes amongst these works. These sections have been laid out in a progressive manner, where one theme transitions into the next.

#### 3.1 FULLY DECOMPOSED VS. PARTIALLY DECOMPOSED NFV

We begin our SDNMN survey by mentioning the study done by Basta et al in [5] which is one of the earlier works in this area. The authors propose four possible deployment scenarios for full or partial decomposition of S-GW/P-GW functionality as virtualized network functions (VNFs) in the cloud. These designs propose replacing all switches and forwarding elements within the operator's backhaul network with Open Flow enabled switches that use a logically centralized SDN controller to steer the traffic from the core EPC network to the cloud network. Functions such as policy control, packet filtering and charging are kept within the Open Flow switches. The authors propose using Open Flow counters and statistics messages, as defined in the specification, to realize offline and online tariff charging.

It is our opinion that a fully decomposed S-GW/P-GW design may not be feasible owing to stringent latency and bandwidth requirements in the data plane. Since both EPC control and data plane traffic is encapsulated as GPRS Tunneling Protocol (GTP) packets, either the forwarding Openflow switches must parse these GTP headers or these switches must send those packets north-bound to the controller which in turn must have in-built GTP parsing logic. Given that

Open Flow presently does not have any provisions for GTP header parsing, such processing will need to happen in the controller. We argue that doing so will incur a substantial performance penalty for the data plane. In contrast, control plane traffic is fewer and further in between and has more relaxed timing requirements. GTP parsing for such traffic therefore seems a more feasible option.

### **3.2 PARTIALLY DECOMPOSED S-GW**

Nevertheless, we cannot discount the fact that Basta's study paved the way for a number of interesting follow-on proposals in this area. One of which was [6] in which Said et al focus on resiliency, and load balancing within the Mobile network. The authors recognized the benefits of partially decomposing EPC functionality within the cloud and propose an architecture where the entire intelligence within the SGW (SGW-C) and MME is centralized as applications running on top of an Open Flow controller.

To this end, the authors replace the GTP based control protocols that run on the S1-MME (between MME and eNodeB) and S11 (between MME and SGW) interfaces by the OF protocol. All aspects of SGW responsible for user plane and data forwarding (SGW-D) are still kept within the MNO's core network, as well as all other EPC components such as PGW and HSS. The key benefit of consolidating SGW-C functionality above the SDN controller, is that a user session may be load balanced between SGW-D sessions without having to reallocate a Tunnel Id (TEID) and re-provision sessions. The TEID value is therefore allocated once per session within the SGW-C and remains invariant during UE session mobility across SGW-Ds. In some aspects, we feel that Said's design can borrow concepts from SIPTO and make the decision to offload a UE session to a new SGW-D based also on the UE's geo-proximity to a SGW-D (and not just load statistics)

### **3.3 UE AGNOSTIC TRAFFIC REDIRECTIONS AND MIDDLEBOXES**

All things considered, the concept of full decomposition of an EPC component in the cloud is still an interesting viewpoint and one which is exploited by Banerjee et al in their MOCA architecture [4]. In MOCA, UE uplink data is offloading to an on-premise cloud network that contains virtualized SGW and PGW nodes that receive UE traffic and transmit the corresponding Downlink traffic. This offload is meant for specific UE sessions only. The authors propose online gaming as being the business case that will drive adoption of this offloading strategy; although the architecture is general enough to be applicable to a number of applications. MOCA does require a number of changes to the MME in order to drive dynamic offloading of specific mobile traffic.

In order to determine the UE session that needs to be offloaded, and the enodeB - SGW endpoints for this session, the MME will need to extract UE IP address, the SGW TEID, the enodeB IP address and the enodeB TEID from the message primitives involved in a UE attach signaling procedure. The MME will need to maintain a mapping between the UE IMEI and these extracted parameters, for use when offloading is required. It must be noted here that inspection of EPC control traffic to extract these offloading primitives is not a requirement uniquely imposed by MOCA. Government legislations require MNOs to provide network traffic information to law enforcement agencies. Lawful interception is therefore a main feature of most EPC networks. Amongst other types of information, MNOs must provide Intercept-Related Information(IRI) for UE attachment and detachment procedures [16] which is what MOCA needs to extract the offloading attributes.

Offloading is initiated via an external trigger to the MME, from an offloading front-end which in the author's view is the game server front-end. If the offloading front-end is an MNO provided

service then we may contend with using one of the existing interfaces such as the S6a interface, however in the case that this is a 3rd party service, such as the gaming server use-case that MOCA's authors have stipulated, then we do have some interoperability and security concerns, since operators would not be too keen to open up a critical control component to external service providers and OTT players.

The authors of MOCA realize that Open Flow switches in their current stage do not support GTP header parsing, and therefore propose using Middleboxes that parse the GTP header for packets in Uplink and redirect the flow to the offload cloud network if it belongs to a UE that has been marked for offloading. Similarly, the Middlebox will receive downlink packets from the cloud based S-GW but must rewrite their outer source IP address to that of the original S-GW. This is because unlike some of the previously discussed techniques which require explicit signaling to register the new S-GW, both the UE and eNodeB are agnostic to the offloading and assume the packet comes from the registered S-GW. Finally, although not explicitly indicated by the authors, it is our understanding that MOCA is meant only for UDP based sessions as TCP requires explicit handshaking between the UE and the cloud based gateway which will not work as the flow redirection is being planned to be transparent to the UE.

After a close examination of MOCA's design, one may question the need for an offload S-GW and P-GW if the intended use-case is simply to connect to a local Packet Data Network (PDN). Banarjee et al indicate clearly that the new S-GW will not respond to paging requests as its presence is oblivious to the UE; nor can we expect it to be a mobility anchor for inter-3GPP communication. The new P-GW is also a fairly thinned out version of the original thing - it will not be allocating UE IP addresses and likely not providing inter-connect with non-3GPP technologies. The offloading network is simply a fixed IP network, very similar in kind to the local network provided to UEs in LIPA. In the usability context that MOCA has stipulated i.e. online gaming and other delay-sensitive applications, providing fully decomposed gateways is a bit of an overkill.

Such a perspective resonated with the authors of SMORE [10]. Cho et al, whose work is based on MOCA, recognized that any modifications to the standard MME including opening it up for 3rd party interaction, is a path of maximum MNO resistance and thereby proposed an alternative design that is identical to MOCA with the exception of the offload network not containing full-fledged S-GW and P-GW nodes. Accordingly, the MME does not need to be modified to send overloaded Create Session Request messages. Also, the external offload trigger now comes to a separate component within SMORE's architecture called the SMORE controller which may be an application running on top of an SDN controller. SMORE seems to be a more pragmatic approach to the issue of UE agnostic traffic offloading.

One common theme in both MOCA and SMORE is that offloading is caused by an external input. As we mentioned previously, there are certain interoperability challenges in providing non-standard interfaces within EPCs. Given the high value nature of networks today and sensitive information they carry, operators are apprehensive about modifications not dictated by the LTE standard. Most operators would argue that the trigger to offload must come from within the network itself.

The authors of ProCel had a similar idea in [18]. Nagaraj et al lay claim that upto 70% of traffic in today's networks is delay tolerant and is either data or video-demand traffic, which does not require stringent QoS guarantees, and seamless L3 mobility since it is less tolerant and neither does it normally require inter-networking with other 3GPP technologies. Traversing such traffic across the core EPC network overburdens the network mainly because of the large signalling overheads associated with setting up GTP tunnels for these sessions that may not necessarily



benefit from it. To this end, ProCel attempts to proactively classify all flows in the core network into either core flows, that will continue through the EPC as usual, and non-core flows that will be offloaded to a fixed IP network. The authors indicate in their work that signalling is growing 50% faster than data traffic within LTE networks and adopting such bypass techniques will help to alleviate some of these loads.

Concerning GTP parsing, all these designers recognize the absence of inline matching on the Open Flow switch and propose using middle boxes as a remedial measure. However the exact ramifications of doing so are not discussed. As we mentioned previously, the lowest granularity afforded to use by an Open Flow rule is the IP five-tuple. For UE traffic this becomes:

**In the case of Downlink:** source IP address is the S-GW IP and destination IP address is the enodeB IP.

**In the case of Uplink:** source IP address is the enodeB IP and source IP address is the S-GW IP address

Since Open Flow rules do not support bidirectional at this time, two rules will need to be installed for both the downlink and uplink directions but with different flow actions. What is implied here is that the flow will match all packets between the same endpoints and therefore be steered towards the middle-box which will then decapsulate the GTP header and inspect the inner header to determine if this packet is indeed for a UE session that was marked for offloading. If it turns out that this was indeed the case then the effort was not in vain and the packet can be forwarded to the offload network. However, as will be the general case, most packets will belong to sessions that should not be offloaded. This introduces an additional latency which may be highly undesirable for delay-sensitive sessions such as GBR bearers used for conversational voice, VoIP and video conferencing.

In our understanding, building GTP parsing logic inside the OpenFlow controller (or on an application running on top of it) is a more pragmatic approach since the OF switch does not need to send the entire packet to the controller but only the header and a buffer Id in case the controller requests the packet payload. The OF controller can offset some of the latency introduced for non-offloading flows by rewriting the output port to that of the destination or the next hop switch, thereby saving some cycles had the non matching packet been processed through the switches L2 forwarding table.

The implications of this realization are severe: *if an enodeB has N UEs attached to it and only 1 of them is marked for offloading then all remaining N-1 UEs will incur a performance hit.*

Grounded in this reality, we do believe that the problem is slightly less worse for ProCel when compared to MOCA and SMORE. Since ProCel impartially tries to classify all flows, and not just those traversing a certain enodeB, the additional delay introduced is a network constant and is deterministic, allowing MNOs to factor that apriori.

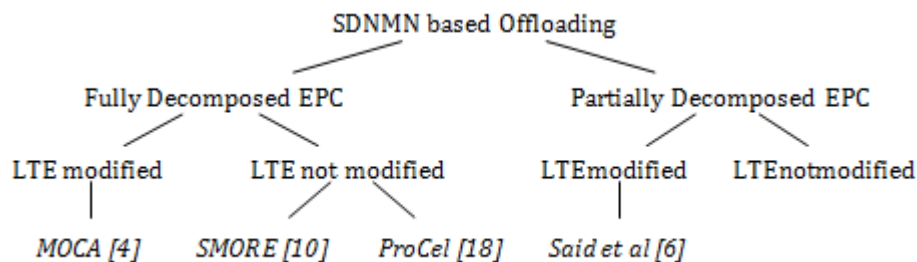


Figure 3: Taxonomy of surveyed SDNMN offloading techniques for core network traffic

The taxonomy tree presented in Figure 3 depicts a summary view and classified the works surveyed in this paper across two tiers a) whether full or partial virtualization of EPC core components is required, and b) whether a change in the LTE - EPC standard will be required

#### 4. AGNOSTIC EPC TRAFFIC OFFLOADING THROUGH SDN (AETOS)

We propose a reference design for traffic offloading based on some of the designs evaluated in the previous sections. AETOS, as we call it, is based on an established set of *design non-negotiables*, influenced by some of the outlined related works:

- The design should discourage modifications to standard LTE components, interfaces, and behavior. Any such changes would require buy-in from the 3GPP standardization body.
- Benefits in the user plane, via offloading, should not be achieved on the back of increased complexity in the control plane. Existing offloading techniques such as LIPA, SIPTO and S1-Flex require additional signaling to set up the offloading pathways.
- In the absence of inline GTP parsing on the OF switch, GTP header encapsulation/decapsulation is best served via purpose-defined middleboxes located at proximity to the RAN and not at the gateways at the core network edge. This is to minimize processing delays experienced by sessions which are sent to the GTP parsing middlebox, only to be sent back since they were not marked for offloading.

One observation we eluded to earlier was the fact that existing LTE offloading mechanisms such as SIPTO and S1-Flex provide data plane alleviations at the expense of additional signaling required to establish these offloading pathways. To this end, UE agnostic offloading solutions such as MOCA, SMORE and ProCel, are attractive design inspirations for our framework.

Finally for the design to make practical sense, from an implementation perspective, it will require modifications within the provider's backhaul network, i.e. the series of switches and Ethernet links that connect EPC components with each other in the same switching domain and/or the routers involved in the EPC components span multiple subnets or routing domains. Above all, the backhaul network must now be comprised of *Openflowhybrid*[13] switches that can talk upstream with an SDN controller. The Open flow switch must be able to steer traffic away from the core network (either the MME or the S-GW) to a fixed IP gateway that represents the offloading server. The backhaul network will also need to be comprised of aggregation points, located near the RAN edge, and which terminate multiple enodeB's within that geographic region and connect them to the EPC network.

Fortunately, such aggregation points already exist within the mobile backhaul and are referred to as Mobile Telephone Switching Offices (MTSO's) [8]. In our design, an OF switch, with middle-boxes co-located at the switch, would be a perfect embodiment of an MTSO and represent an optimal placement for our offloading framework. Therefore these MTSOs shall act as traffic funnels either to and from the EPC or to other MTSOs.

Figure 4 illustrates the logical breakdown of an MTSO as it relates to our system architecture. Each MTSO will be comprised of an Open Flow Switch, with a co-located middlebox that will be responsible for GTP header parsing and extraction. As can be seen in the figure, MTSO A serves as a middle-man between the EnodeB's in the radio access network, and the MME and S-P/GW nodes in the EPC. Such deployments will be hierarchical with MTSOs interconnected with each other, contingent on the size and complexity of the mobile operator's backhaul network.

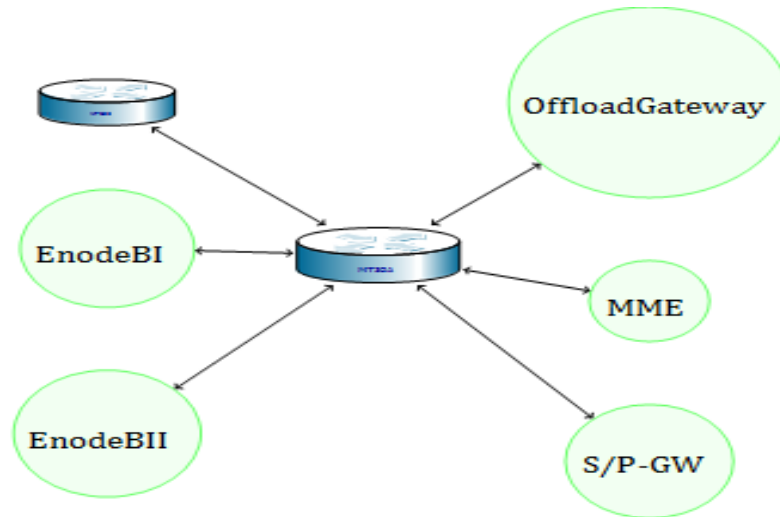


Figure 4: MTSO Deployment

AETOS shall have multiple layers of interactions between a) the EPC core network and MTSO, b) the SDN controller and the MTSO, and c) an orchestration agent to receive offloading rules and policies from external stimuli. These rules may indicate specific UE IMEIs to offload or more generic QoS session parameters, such as offload all non-GBR traffic.

Analysis of each of these different layers of interactions and the message primitives involve having been outlined in detail in [19]. Once offloading rules have been installed, all incoming UE flows need to be matched against those rules. Consequently, the MTSO shall forward traffic from UEs marked for offloading to the fixed IP network. The open Flow rules that steer said traffic are formulated in the co-located Middle-box. The processing that happens at the Middle-box can be summarized by the algorithm given below:

---

**Algorithm 1 Middlebox Packet Processing**

---

**Data:** Packet In message

**Result:** Packet Out message

```

1      begin
2      strip outer most IP header decapsulate GTP header parse source and destination
      TEID and
      UE IMSI if UE IMSI tagged for Offloading then
3      |   if matched Uplink rule then
4      |   |   if offload gateway is S-GW then
5      |   |   |   modify GTP header → destination TEID = offloadServer TEID
6      |   |   |   action1 → Rewrite Destination IP(offload gateway IP) action2 → Output(offload
7      |   |   |   gateway endpoint) add action1, action2 to action list
8      |   |   |   else if matched Downlink rule then
9      |   |   |   modify GTP header → source TEID = S-GW TEID action1 → Rewrite Source IP(S-
10      |   |   |   GW IP) add action1 to action list
11      |   |   else
12      |   |   |   action1 → Output(PacketIn out port) add action1 to action list
13      |   |   end if
14      |   end if
15      end if
16      Create Packet Out msg with action → action list

```

---

As a consequence, in addition to the forward-to-controller rules in the OF switch, the middle box will now receive *Packet-In* messages. We benefit from an optimization that exists within the Open flow standard where the payload of the received Packet-In message is only the IP header and not the full original packet. A Buffer Id is provided within the message in case the controller requires the full packet.

## 5. AETOS PERFORMANCE CONSTRAINTS

AETOS must remain cognizant to the fact that improved operating efficiency in the LTE core network, via traffic offloading will incur a performance hit in order to steer flows towards the offloading network. It should also be aware of the practical limitations of switches vis-a-vis inline GTP parsing and the penalty incurred in sending all its flows to the middle box.

We, therefore, stipulate the following measurement goals for any prototype deployment that aims to qualify the AETOS architecture:

**The Round-Trip-Time(RTT) must not exceed the assigned packet delay budget of the UE session:**

Intrinsic to a UE session's QoS Class Identifier (QCI) profile is the acceptable RTT. For instance, for interactive video and certain online gaming applications, this is typically set at 100 - 150ms [17].

Measurements will need to be taken to ensure that both the offloaded sessions, and the ones that were sent to the middlebox but returned (since they were not marked for offloading) both satisfy their RTT requirements.

**Based on the cardinal principles of seamless LTE handovers, offloaded UE sessions should not be disrupted during X2 handovers:**

Measurements will need to be taken to ensure that AETOS can recognize a handover and update the UE extraction table and forward these updates across the X2 interface, all the while ensuring that the session is not disrupted.

## 6. AETOS FAIRNESS CONSTRAINTS

The measurement goals outlined in the previous section serve as basic sanity for qualifying AETOS as being field acceptable. However, a design being practical does not imply that it is *fair*. User sessions that had to be diverted to the co-located middlebox, only to be returned back can be deemed to have been treated *unfairly* if we contrast this to offloading sessions that benefited from a net improvement in RTT after being transferred to the Offload Network.

On the subject of network engineering, a number of fairness measures are available in academia of which Jain's Fairness Index (JFI) [15] is the most prevalent. In the classical sense of the term, a design is deemed fair if every user determines that it has received the same equal portion of a shared resource [15]. In the context of this discussion, let us now assume that the "shared resource" is RTT. Each user, by virtue of the network, is provided an RTT. This RTT value is dependent on the QCI value of the session or the network path is undertaken. We also argue that for MNOs to justify offloading, the RTT of offloaded sessions needs to be lower than that of core sessions. MNOs can achieve lower RTTs in the offload network(s), simply by virtue of it being a fixed IP network, being closer to the RAN edge, and having a smaller user base. Therefore right off the bat, we know that any measure of fairness would be tilted in favour of the offloading sessions. However, the bigger question is what degree of unfairness MNOs can consider

acceptable in lieu of improved operational efficiency. Let us now compute a fairness relationship between the user sessions in an AETOS enriched core network.

Let us assume that  $n$  UE sessions are active through an MTSO, let  $x$  be a measure of RTT with  $x_i$  being the RTT of the  $i^{th}$  session. Out of these  $n$  sessions, let us assume that  $r$  sessions are regular and not marked for offloading. All others have a constant lower RTT value of  $x_{offload}$ . The JFI can then be computed as:

$$f(x) = \frac{(\sum_{i=1}^r x_i + \sum_{i=r}^n x_{offload})^2}{n(\sum_{i=1}^r (x_i)^2 + \sum_{i=r}^n (x_{offload})^2)}$$

The constant summation of the offload RTT can be simplified to:

$$f(x) = \frac{(\sum_{i=1}^r x_i + (n-r)(x_{offload}))^2}{n(\sum_{i=1}^r (x_i)^2 + (n-r)(x_{offload})^2)}$$

The quantity  $(n-r)$  represents the number of sessions that will be offloaded. Practicality dictates that MNOs will not maintain different counts of offloaded sessions per MTSO, and would rather have a global *offload factor* that can apply to a number of MTSOs. Ergo, let us now re-write the  $(n-r)$  quantity as a factor of  $r$ , represented below by  $\gamma$ :

$$f(x) = \frac{(\sum_{i=1}^r x_i + (\gamma r)(x_{offload}))^2}{n(\sum_{i=1}^r (x_i)^2 + (\gamma r)(x_{offload})^2)} \quad (1)$$

To simplify this relation further, let us now assume that all the regular UE sessions have the same constant RTT value (a more advanced analysis would model this RTT as a distribution), referred to as  $x_{regular}$ . The JFI value now becomes:

$$f(x) = \frac{(rx_{regular} + (\gamma r)(x_{offload}))^2}{n(r(x_{regular})^2 + (\gamma r)(x_{offload})^2)} \quad (2)$$

By expressing  $n$ , the total number of sessions as a factor of  $r$ , this relation can be further simplified to:

$$f(x) = 1/(\gamma + 1) * \frac{(x_{regular} + \gamma x_{offload})^2}{x_{regular}^2 + \gamma x_{offload}^2} \quad (3)$$

Based on the principals of polynomial expansion, we know that the numerator term will be larger than the denominator and their ratio will be greater than unity. We also know that the JFI is bounded by a maximum value of unity. Therefore the possible range of JFI for the above set of constraints and assumptions now becomes:

$$\boxed{1/(\gamma + 1) \leq f(x) \leq 1; \forall x \geq 0} \quad (4)$$

MNOs will be more concerned with the lower range as it represents the worst-case scenario. As an example, consider a deployment where 50% of the sessions are offloaded. Based on the relationship above, the system has a worst case fairness value of 0.67 meaning that 33% of sessions have been subject to discrimination. Consider now the scenario outlined in ProCel [18] which postulates that 75% of traffic is “delay-tolerant” and therefore subject to offloading. The minimum JFI value now becomes 0.57 with 43% of UE sessions being treated unfairly by the system.

It is our hope that the relationship, derived above, serve as a useful measure for MNOs in deciding the degree of offloading in their networks, and to what extent a system bias can be tolerated in their deployments.

## 7. AETOS PERFORMANCE EVALUATION

A performance evaluation simulation for the AETOS framework was implemented using Network Simulator 3 (NS3) [7, 21]. NS3 is a discrete event simulator whose modular and open source nature makes it fairly popular in the research community. LTE EPC modules are available out-of-box and were modified to interact with the SDN controller and the Open Flow protocol. The simulation considers an equal number of guaranteed and non-guaranteed bearers (GBR vs. non-GBR) with all UEs belonging to an EnodeB placed equidistant from it, to minimize variable RAN effects and have UEs experience similar propagation path losses. In reality, however, we should expect much more complex deployment and traffic patterns.

A baseline run is done over an increasing number of UE sessions, followed by a re-iteration that offloads all the non-GBR sessions. A worst-case scenario is played out wherein traffic pertaining to only a single UE is offloaded amongst all the active UE sessions. A number of insights were drawn related to the trend of the performance metrics and how their values varied with the scale of the network:

1. The transmit throughput (as seen at the enodeBs) increases in direct proportion to the number of UE sessions. However, the receive throughput eventually plateaus since the switch serves as a major bottleneck. Beyond this point any further increase in transmit throughput translates into an increase in packet delay
2. The average packet delay rises sharply when scaling UE sessions. Such a rise may be observed in a real implementation as well and constrains the number of GBR sessions that an MTSO deployment can handle.
3. Maximum jitter grows with the number of GBR sessions. It represents the maximum bias between GBR and non-GBR sessions. Since the P-GW favors the former, nonGBR sessions are penalised by longer packet queuing delays. It is for this reason also that a majority of the packet drops observed in this deployment are for non-GBR sessions.

Figure 5 illustrates the variation of mean packet delay with an increasing number of UEs per EnodeB. A logarithmic curve of best fit is used to indicate the underlying trend as well as facilitate easier comparison between the three test schemes.




Figure 5: Variation in Mean Delay with UE scale over 3 different deployment scenarios




Figure 6: Variation in Mean Jitter with UE scale over 3 different deployment scenarios

Figure 6 outlines the variation in mean packet jitter against number of UEs/EnodeB. Given the cyclic nature of this trend, a fourth order polynomial is used as curve of best fit.

**When only 1 UE is offloaded:** all UE traffic between those EnodeB and the S-GW endpoint(s) is forwarded to the middlebox controller for GTP parsing. Most of these will be sent back to the switch, as-is, via *Packet-Out* messages. Those that do match the offloading criteria will experience much lower RTTs (owing to the simplistic nature of the UDP Echo Server that acts as the offloading gateway). It must be noted that this resonates with real-world expectation as the offloading gateways will be providing MNOs significantly lower packet RTTs than the core network.

The mean delay drops for small UE scale but rises sharply beyond 4 UE sessions, so much so that at 150 UEs/EnodeB, the observed mean packet delay is many folds higher than the baseline delay. Conversely, the jitter value between two consecutive packets is maximized if one of them goes through the conventional core network pathway, while the other goes through the much faster offloading pathway. The mean jitter is therefore higher than the baseline at small UE scales only. This is because, at smaller UEs scales, the proportion of offloaded sessions is high (50% for 2 UEs/EnodeB; 25% for 4 UEs/EnodeB) enough to cause the average jitter to rise.

**When all non-GBR sessions are offloaded:** there is a net lower contention at the EnodeB MAC scheduling layer (for downlink) or the P-GW (for uplink). These layers can now focus better on prioritizing GBR sessions. It is for this reason that the mean packet delay improves significantly in this scheme and the benefit grows further with the number of active UEs, regardless of the increasing performance bottleneck at the controller.

Consequently, the mean jitter in this scheme is much higher than the baseline curve. As more and more sessions are offloaded, the net difference between those that go through the conventional core network and those flows that go through the faster offloading pathway, grows and causes the mean jitter to swing in that direction. This serves as further testament to the fact that at larger scales, a dedicated and service-centric offloading server is better suited to performance than a centralized S-GW/P-GW service located far from the RAN edge.

## 8. CONCLUSION

In this work, we have investigated novel traffic offloading schemes and compared them to existing offloading mechanisms such as LIPA, SIPTO, and S1 Flex. When analyzing some of these works, we remain in favor of scenarios that do not provide data plane offloading at the expense of increased signalling in the control plane. Our reference design, therefore, focuses on the genre of *agnostic* traffic offloading, where the endpoints remain oblivious to the presence of the middleware fabric that offloads its traffic.

We then derived a Jain's Fairness Index relation for our AETOS reference with the aim of providing MNOs the ability to gauge the fairness of their deployment scenario in real-time based on variations in their network RTTs and user sessions. It is hoped that this mathematical relationship will serve further to incentivize offloading by allowing MNOs to justify their choices from a quantitative point of view, i.e. *how many sessions to offload given my fairness threshold?*, as opposed to a purely qualitative point of view, i.e. *should I offload or not?*. By asking the right questions and providing the right answers, mobile operators may gain easier buy-in from their stakeholders in rolling out such schemes.

Our work then explores some of these cardinal principles by prototyping a simple AETOS on the NS3 simulator. Our findings also resonate with what we postulated earlier: *agnostic offloading in the face of non-inline GTP parsing only makes sense if a broad group of sessions are offloaded simultaneously*. This serves to bolster ProCel's delay-tolerant traffic offloading scheme while reminding us that the on-demand gaming server application devised by SMORE and MOCA will suffer from tremendous practical barriers to deployment.

## 9. FUTURE WORK

Our work may be expanded upon in a number of interesting ways. A future work might investigate supporting agnostic offloading over a TCP based transport. This will be a greater challenge given that a TCP session is negotiated end-to-end between the P-GW packet gateway and the UE. Packets injected into the session from the offloading gateway in the downlink direction run the risk of TCP checksum failures at the UE. Support for TCP session offloading in AETOS will facilitate wider adoption of this framework within the mobile community.

Another subject to ponder is the fact that the GTP parsing middle-box is statically provisioned which has the disadvantage of permanently reserving MNO resources even during low traffic periods where offloading may not necessarily be required. Most modern controllers allow dynamic Service Function Chaining (SFC). Open Daylight (ODL) [1] is one such Controller platform which prides itself in being carrier-grade and is widely adopted by a number of industry partners. SFC may be used to dynamically install the GTP parsing middle-box as driven by a trigger from the Orchestration Agent. Future research may be interested in further investigating these service chains.

## ACKNOWLEDGMENT

The authors would like to thank the authors and contributors whose works were included in this study.

## REFERENCES

- [1] Opendaylight platform. <https://www.opendaylight.org/>. Accessed: 2016-01-30.
- [2] Global Games Market Grows to \$86.1bn in 2016. <http://www.newzoo.com/pressreleases/global-games-market-grows-to-86-1bn-in-2016>, October 2013.



- [3] Mojdeh Amani, Toktam Mahmoodi, Mallikarjun Tatipamula, and Hamid Aghvami. Sdn based data offloading for sdn based data offloading for 5g mobile networks g mobile networks. ZTECOMMUNICATIONS, page 34, 2014.
- [4] Arijit Banerjee, Xu Chen, Jeffrey Erman, Vijay Gopalakrishnan, Seungjoon Lee, and Jacobus Van Der Merwe. Moca: a lightweight mobile cloud offloading architecture. In Proceedings of the eighth ACM international workshop on Mobility in the evolving internet architecture, pages 11–16. ACM, 2013.
- [5] ArsanyBasta, Wolfgang Kellerer, Marco Hoffmann, Karel Hoffmann, and E-D Schmidt. A virtual sdn-enabled lte-epc architecture: a case study for s-/p-gateways functions. In Future Networks and Services (SDN4FNS), 2013 IEEE SDN for, pages 1–7. IEEE, 2013.
- [6] Siwar Ben Hadj Said, Malla Reddy Sama, KarineGuillouard, Lucian Suciu, Gael Simon, Xavier Lagrange, and Jean-Marie Bonnin. New control plane in 3gpp lte/epc architecture for on-demand connectivity service. In Cloud Networking (CloudNet), 2013 IEEE 2nd International Conference on, pages 205–209. IEEE, 2013.
- [7] Gustavo Carneiro. Ns-3: Network simulator 3. In UTM Lab Meeting April, volume 20, 2010.
- [8] KJS Chadha, CF Hunnicutt, SR Peck, and J Tebes. Advanced mobile phone service: Mobile telephone switching office. Bell System Technical Journal, 58(1):71–95, 1979.
- [9] Margaret Chiosi, Don Clarke, P Willis, A Reid, J Feger, M Bugenhagen, W Khan, M Fargano, C Cui, H Deng, et al. Network functions virtualisation introductory white paper. In SDN and OpenFlow World Congress, 2012.
- [10] Junguk Cho, Binh Nguyen, Arijit Banerjee, Robert Ricci, Jacobus Van der Merwe, and Kirk Webb. Smore: software-defined networking mobile offloading architecture. In Proceedings of the 4th workshop on All things cellular: operations, applications, & challenges, pages 21–26. ACM, 2014.
- [11] Cisco Visual Networking Index Cisco. Global mobile data traffic forecast update, 2013–2018. white paper, 2014.
- [12] Cisco Visual Networking Index Cisco. Global mobile data traffic forecast update, 2016–2021. white paper, 2017.
- [13] OpenFlow Switch Consortium et al. Openflow switch specification version 1.1. 0 implemented (wire protocol 0x02), february 2011.
- [14] SouheirEido and Annie Gravey. How much lte traffic can be offloaded? In Advances in Communication Networking, pages 48–58. Springer, 2014.
- [15] Raj Jain, Dah-Ming Chiu, and William Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. 1998.
- [16] Amit Kumar, Jyotsna Sengupta, and Yun-fei Liu. 3gpp lte: The future of mobile broadband. Wireless Personal Communications, 62(3):671–686, 2012.
- [17] Alcatel Lucent. The lte network architecture a comprehensive tutorial. Strategic Whitepaper, 2009.
- [18] Kanthi Nagaraj and Sachin Katti. Procel: Smart traffic handling for a scalable software epc. In Proceedings of the third workshop on Hot topics in software defined networking, pages 43–48. ACM, 2014.
- [19] Kam Nasim. AETOS: An Architecture for Offloading Core Software Defined Networking Concepts. Master’s thesis, University of Ottawa, 2016.
- [20] CISCO Global Visual Networking and Cloud Index. Forecast and methodology, 2011-2016.
- [21] George F. Riley and Thomas R. Henderson. Modeling and Tools for Network Simulation, chapter The ns-3 Network Simulator, pages 15–34. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [22] CB Sankaran. Data offloading techniques in 3gpp rel-10 networks: A tutorial. Communications Magazine, IEEE, 50(6):46–53, 2012.