Policy Driven Multi-band Spectrum Aggregation for Ultra-broadband Wireless Networks

Lance Hartung
University of Wisconsin-Madison
hartung@cs.wisc.edu

Milind M. Buddhikot
Alcatel-Lucent Bell Labs
milind.buddhikot@bell-labs.com

***Abstract*—Sustaining the impressive mobile Internet revolution that has transformed our lives requires addressing the key challenge of designing future ultra-broadband wireless networks that can keep up with the data traffic growth. One of the key techniques to increase per-user as well as system capacity is to increase the amount of spectrum used in the network. However, given cost-effective radios that can radiate in 100s MHz of spectrum are still not practical, alternate solutions that concurrently exploit multiple radios and multi-band support in current devices and networks need to be explored. The carrier aggregation technology in LTE-Advanced networks represents a step in this direction. However, such an intra-network, single standard baseband specific approach does not scale well as it does not effectively exploit multi-network. multi-band and multi-operator connectivity and therefore, to rectify these shortcomings we envision an aggregation solution at transport or network-layer. We propose an end-to-end architecture that uses Multi-path TCP (MPTCP) as the key building block and supports concept of multi-path policies that exposes to end-users and networks unprecedented control on use of network resources on various spectrum bands and paths. We identify that uncontrolled widespread MPTCP use can lead to unfair resource allocation and reduced spectral efficiency. We solve this problem using a centralized network-resident service for managing multi-path usage throughout an administrative domain such as an operator network, or enterprise. This service makes central decisions based on network-wide information and operator policies. We describe our end-to-end prototype and report on experiments that demonstrate optimized aggregate network capacity in a variety of interesting scenarios.**

I. Introduction

As the mobile data traffic growth continues unabated, dramatic increase in wireless network capacity is necessary. There are three main ways to increase capacity: (1) Increase spectral efficiency measured in bits/sec/Hz to better use a given amount of spectrum, (2) increase area spectral efficiency measured in bits/sec/Hz/m^2 by dramatically shrinking the cell size, and (3) increase the amount of spectrum used in the network.

The second principle above has led to introduction of small cells - indoor as well as outdoor in the cellular architecture, which now is transformed from existing macro cellular networks into a Heterogeneous Network, often termed simply as *hetnet*. In the first wave of deployment, termed shared carrier, small cells will share the channels with the the macrocells they embed in. In this type of deployment, the co-channel interference limits deployment flexibility, total capacity and efficiency of traffic offload. The unconstrained capacity scaling requires that small cells use channels from dedicated spectrum not used in macro-cells.

As for the third principle of increasing the amount of spectrum used in the network, technology trends indicate that future small cells will likely support a mix of licensed (e.g.: FDD/TDD 3G/4G LTE bands), unlicensed (e.g.: 2.4, 5.1, 5.4, 5.8 GHz bands) and new shared bands (e.g.: 3.5 GHz Naval Radar band in USA, 470-700 MHz DTV band in some parts of the world and 2.3 GHz band in Europe). Also, there has been a strong push to use LTE technology across even unlicensed and shared bands for many reasons: (1) offloading traffic from licensed bands greatly improves the throughput and user experience, (2) LTE solutions are more spectrally efficient than Wi-Fi [2], (3) the LTE carrier aggregation (CA) framework can enable seamless integration of multiple spectrum bands, and (4) a common LTE architecture across spectrum bands can simplify the network by exploiting already deployed backhaul and core network. Significant industry efforts in this direction have produced three technology variants: (1) *LTE-Unlicensed (LTE-U)* based on 3GPP Rel 10/11/12 and defined by LTE-U forum[5], (2) *Licensed Assisted Access (LAA)* which is being investigated in 3GPP Release 13, and (3) *muLTEfire* proposed and championed Qualcomm[3].

The LTE-U and LAA focus on solutions that combine unlicensed secondary carriers with a primary licensed band carrier anchor using carrier aggregation to achieve higher throughput and enhanced user QoE. LTE-U applies only to countries that do not mandate Listen-Before-Talk (LBT) protocols (e.g.: U.S.A and Korea) and will support only Supplemental Downlink (SDL) mode of operation[6]. It relies on channel selection and LTE-U ON/OFF duty cycle mechanisms to share the unlicensed channel in Time Division Multiple access (TDM) manner. LAA is a more general technology that will enhance LTE baseband with Listen-Before-Talk (LBT) co-existence mechanism to better share an unlicensed channel with the incumbent Wi-Fi and any other networks including LAA. Unlike LAA and LTE-U which require primary licensed band anchor, and therefore can be used only by cellular operators, muLTEfire can operate “stand alone” and can be deployed much like Wi-Fi.

LTE deployment in unlicensed spectrum is primarily a small cell technology due to lower transmit power requirements of the unlicensed spectrum. Clearly, these trends point to likely evolution to multi-band, multi-operator small cells supporting diverse basebands such as LTE, LTE-U/LAA, and Wi-Fi.

Current user equipment (UE) (e.g.: phones, tablets and laptops) already support multiple radio interfaces such as a cellular interface capable of operation in multiple 3G (CDMA, WCDMA) and 4G (FDD/TDD LTE) bands and a Wi-Fi interface capable of 11n/11ac operation in 2.4/5 GHz unlicensed bands. Future UEs will add at least one more interface supporting LTE-U/LAA or muLTEfire baseband in unlicensed and shared bands. Current UE software allows multiple interfaces to be used in an either-or fashion, i.e. only one interface is active at any given time. This limits the user-experienced throughput to capacity in the corresponding network and spectrum band – typically in the range of 10s of Mbps. To achieve high sustained per user/device or application throughput, it is essential to enable concurrent use of multiple radio interfaces.

Given a network of multi-band, multi-radio small cells and client devices, realizing ultra-broadband capacity – characterized by multi-gigabit/sec access capacity and 100s of Mbps of per-user throughput is an important research challenge. Specifically, we need a flexible end-to-end solution that can (a) seamlessly aggregate (mobile) traffic across 100s of MHz of non-contiguous spectrum and across multiple small cells and (b) provide network and users control on such aggregation. This paper reports design and implementation of such a solution.

A. Our approach

Given cost-effective radios that can radiate in very wide channels (e.g.: 100+ MHz) are still not practical, alternate solutions are needed to concurrently exploit multiple radios and multi-band support in current devices and networks. The carrier aggregation (CA) technology implemented in LTE-Advanced networks represents a step in this direction. However, this intra-network, single standard, layer-2 approach at baseband requires new standards and baseband chips for UEs and basestations and therefore, does not scale to exploit multi-network. multi-band and multi-operator connectivity. To rectify these shortcomings, we envision an aggregation solution at higher layers of the protocol stack.

Some of the fundamental behaviors built into the network stack, historically designed for single network connection, hinder effective use of multiple wireless interfaces. Although routers within the network have flexibility to exploit multipath, connections between mobile devices and a remote server are bound to an IP address assigned to one of the interfaces on mobile device. Therefore, we focus on multipath aggregation at the transport layer. We propose use of a promising new transport protocol, Multipath TCP (MPTCP). Multipath TCP (MPTCP) is an extension to TCP that enables hosts to create additional subflows that spread data for a single connection across multiple network paths. MPTCP has been used for a variety of specific applications such as in data center networks for dispersing traffic away from congested paths[20], and in mobile devices for aggregation and mobility support. In particular, it has been demonstrated in prior work to aggregate cellular and Wi-Fi paths [21].

B. Our Contributions

We address the end-to-end interactions that occur when multipath functionality—specifically using MPTCP—is deployed on network of multi-radio wireless devices. Our main contributions are as follows:

A new aggregation architecture: We propose an end-to-end architecture that uses Multi-path TCP (MPTCP) as the key building block and supports concept of multi-path policies that exposes to end-users and networks unprecedented control to optimize multipath usage for configurable objectives such as aggregate throughput and user priority.

Centralized policy management: We identify that uncontrolled widespread MPTCP use can lead to unfair resource allocation and reduced spectral efficiency. We solve this problem using a centralized network-resident service for managing multipath usage throughout an administrative domain. This service makes decisions based on network-wide information and operator policies.

Prototype demonstration: We describe our end-to-end prototype that uses (1) a modified MPTCP implementation, (2) a transparent, high capacity, network resident aggregation proxy in the data path, (3) a network resident policy controller with our subflow control algorithm, and (4) client devices with support for receiving policy commands and controlling MPTCP subflows. Using our prototype, we evaluate multi-band link aggregation under challenging scenarios in which clients interact and demonstrate improvements in aggregate throughput possible through network-assisted coordination of multipath clients.

C. Outline

The rest of the paper is organized as follows: Section II provides overview of policy driven aggregation solution. In Section III, we discuss in detail our design, providing rationale behind use of MPTCP and centralized policy control on aggregation. Section IV describes the Policy Responsive Aggregated Multiband (PRAM) architecture and GLOBAL policy optimization algorithm. Section V presents experimental evaluation of our architecture and algorithm. Section VI provides a brief overview of the related work on capacity aggregation. Section VII is devoted to discussion of various topics relevant to our aggregation technology, such as application to cable/DSL access, LTE/Wi-Fi integration, comparison with dual connectivity (DC) and carrier aggregation (CA) and relevance to net neutrality debate.. Finally, Section VIII presents the conclusions and future work.

II. Overview of Our Solution

Our solution relies on aggregating spectrum bands at the Layer-4 i.e. transport layer of the network protocol stack. In doing so, our approach is completely agnostic to (a) number of interfaces, (b) frequency bands of operation, (c) licensed, unlicensed or shared nature of the band, and (d) type of baseband (e.g. : LTE, Wi-Fi, LTE-U, Wi-Gig, Bluetooth etc.) of each interface. We leverage the IETF standardized Multipath TCP protocol as the basis for our design. However, we introduce several new enhancements on how it is used:

- **Managed MPTCP in a controlled domain:** Our research suggests that if a large network of mobile devices were to enable MPTCP without explicit coordination between the devices, the aggregate network capacity can suffer. This results from individual user devices having no information about the impact that their traffic has on other users. Additionally, capacity needs of users vary based on time-of-day, location, applications in use as well as willingness-to-pay. It is therefore essential that a network operator take into account these factors to provide preferential service to particular users or applications that need it the most. Such differentiation is also essential because even if a device has a low priority traffic flow, it can potentially take a large share of the network capacity by enabling multipath TCP.
- **Multipath policies:** We introduce the notion of multipath policies, which are rules that dictate in a potentially fine-grained manner whether a user, application, or individual network connection should be allowed or denied the capability of using a network interface for aggregation of a particular traffic flow.
- **Policy manager:** We introduce (i) a policy manager in the network and (ii) a policy agent that runs on each multi-band capable device. This new infrastructure has the purpose of coordinating usage of multipath. The policy manager determines when and how devices (or apps or users) should make use of multipath in order to optimize objectives including but not limited to aggregate network capacity and differential service to devices. In order for the policy manager to optimize network capacity, it may need to collect performance measurements from the policy agents or from other network entities. These measurements may be instantaneous or time-averaged values of signal strength, transmission rates, or loss rates. This infrastructure is also useful for enabling differential service by adjusting multipath usage in a way that is sensitive to users service classes. In addition, the policy manager can integrate with network billing systems to tie insertion of policies allowing greater per-user/app capacity with payment to monetize higher grade of service.
- **Use of MPTCP aggregation in small cell networks:** Any multipath aggregation will perform well when the component paths have similar delay characteristics. Therefore, it is critical to select paths with similar characteristics to achieve high aggregation efficiency. For example, if we aggregate a network path over a LTE connection to a large macro-cell with a nearby Wi-Fi connection, the delay characteristics are often very different with link delay on LTE path order of magnitude more than on Wi-Fi path. We actively eliminate such scenarios: first by using the aggregation mechanism only in a small cell environment over cluster of small cells in a well-defined small region, we limit delay diversity across paths. Second, using the policy mechanism and path knowledge in the policy manager, we actively eliminate aggregation across paths with disparate characteristics.

III. Design Discussion

A. End-host Support for Multipath

Modern mobile devices support multiple radios and in principle, it is possible to make use of multiple links simultaneously by routing different network connections to different interfaces. For example, Delphi [15] uses collaboration and learning to choose between cellular and Wi-Fi networks for the specific needs of each application. In order to support the more ambitious goal of aggregating capacity for a single network connection, we will require some additional support on the device and elsewhere in the network.

We use multipath TCP since it has been proposed as an IETF standard and has a stable Linux implementation. We suggest that deploying a new OS-level feature on mobile devices is feasible due to consumer behavior of frequent device upgrades as well as the possibility for over-the-air updates. Of greater concern is the lack of support for MPTCP in the Internet at large. With MPTCP being an end-to-end protocol, the client cannot take advantage of it if the other end of the connection does not also support it. In the absence of widespread support for MPTCP, we recommend that the operator introduce a TCP translation proxy[4]. The translation proxy splits a TCP connection into two halves, one connected to the MPTCP-enabled device and one connected to the legacy TCP service, and passes data between the two. Researchers have used split TCP approaches to enable the use of different congestion control algorithms on sections of a path with very different characteristics from the wired path such as high delay satellite links [14] and lossy wireless links [13], so the details and trade-offs of splitting TCP connections are well known [17]. In particular, this approach is compatible with all TCP-based applications unless network layer encryption is enabled as with IPsec. Splitting TCP introduces a slight risk to data integrity by acknowledging data before it has been received by the other end host, which makes the session vulnerable in the case that the proxy fails. However, in accordance with the end-to-end principle, applications already incorporate integrity checks, acknowledgments, and mechanisms for verifying the completion of a transaction.

Why use multipath TCP? Our work focuses on the effects of multi-band aggregation on network capacity and on tools for users and network operators to control the use of multipath. Out of various multipath protocols, MPTCP is a serious contender for gaining widespread deployment on the Internet, so we focus on MPTCP for our prototype and evaluation. The changes required to enable multipath are confined to the transport layer, so it can be enabled for any application that uses TCP. Although not every application uses TCP as its transport, a large fraction of Internet traffic does. It may even be possible to enhance certain non-TCP applications by

TABLE I
PREDICTED CLIENT THROUGHPUT (MBPS) IN TWO-DEVICE MULTIPATH SCENARIO.

	Random Access (Multipath)			Proportional Fair (Multipath)			Single Best		
	BS 1	BS 2	Client Total	BS 1	BS 2	Client Total	BS 1	BS 2	Client Total
Client A	5.4	5.4	10.8	27	3	30	54	0	54
Client B	5.4	5.4	10.8	3	27	30	0	54	54
Network Total			21.6			60			108

sending their traffic though a MPTCP tunnel via VTun[10] or OpenVPN[8]. We have verified that such an approach is feasible but do not characterize its performance here as such an evaluation would depend on the nature of the application in question. Delay-sensitive applications are likely to suffer from jitter introduced by TCP's congestion avoidance and reliable delivery. Finally, we expect our work to be compatible with other multipath protocols, whether they reside at the link layer, network layer, or transport layer, as long as they implement similar mechanisms for enabling and disabling subflows.

Among the alternatives that we considered are implementations at the link layer such as MultiFacet [22]. A link layer solution has the advantage that it will work transparently with any transport protocol without the need for a translation proxy, but it requires software support at base stations and on mobile devices alike. We find no apparent technical reason to prefer one over the other but feel that a transport layer solution would be easier to deploy in networks today.

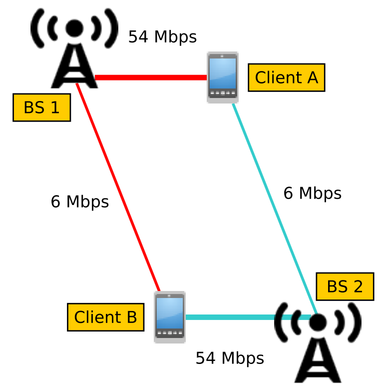

Fig. 1. Example scenario with two client devices simultaneously connected to two base stations.

B. Implications of interaction between MPTCP enabled devices

Suppose a number of network devices have enabled MPTCP and collectively contend for access to wireless resources. Prior research on MPTCP found that in some scenarios using the Linked Increases Algorithm (LIA) for congestion control could lower aggregate network throughput without any benefit to the MPTCP user [16]. Motivated by their work, we wonder if the complex dynamics of wireless communication could interact under MPTCP to produce inefficient network utilization. In particular, we hypothesize that any negative effects on aggregate capacity due to contending devices having different channel conditions will be compounded when they enable multipath.

We start with an example scenario in which two dual-radio client devices are communicating with two base stations, as shown in Figure 1. We assume the base stations use completely orthogonal channels, such that the client devices are able to communicate on both channels simultaneously without self-interference. However, due to spatial proximity, Client A has a stronger channel with BS 1, and Client B has a stronger channel with BS 2, which affects the maximum transmission rate they can use. For the sake of concreteness, we assume 802.11g rates of 6 Mbps and 54 Mbps, as labeled in the figure.

Using two well-known channel access models, we predict the effective data rates of the clients when using multipath versus single path in Table I. For the Random Access model, which is typical of Wi-Fi systems, we assume the client devices receive access to the channel with equal probability at each transmission opportunity, but the client farther from its access point consumes a much larger portion of the airtime. As a result, both clients are reduced to a very low effective data rate—worse than the weakest link. This type of problem motivated a line of research on proportional fair scheduling at base stations [23]. For the Proportional Fair model, we assume that the base stations implement a proportional fair scheduler that allocates equal airtime to each of the clients. Generally, proportional fair schedulers will avoid starving any particular user, so this allocation still gives some airtime to clients with weak channel conditions. Finally, for the Single Best model, we assume each client receives exclusive access to its stronger link. In this scenario, the optimal allocation makes both users better off and improves the overall spectral efficiency of the network.

Consider simple example shown below for a network using Wi-Fi random access. Here a naive multipath connection setup will allow clients A and B to associate with both access points AP1 and AP2. However, in case of AP1, allowing Client B to associate at 6 Mbps brings down the aggregate system throughput and similarly for AP2, allowing client A to associate will reduce its throughput. Ideal assignment policy would require Client A to use only 1 interface and associate with AP1 and Client B to associate with AP2 using 1 interface. Such optimization guarantees high system capacity and high per user throughput. The Policy Server that has view of all bands and small cells in its aggregation scope aims to achieve such system capacity optimization.

This scenario seems to suggest that the best solution is to employ a method for choosing the optimal single path for a network flow. However, keeping in mind that mobility and varying traffic loads make wireless networks highly dynamic, there are several ways in which the scenario could change to become more favorable for multipath. For example, if one client stops transmitting or moves out of range, then the other should opportunistically make use of both links, regardless of their relative quality. In another case, the administrator may prefer to give one client but not the other access to multiple links for higher throughput. Finally, in situations where the two clients have similar channel conditions, this inefficiency is less likely to exist. A depiction of these various modes of operation is shown in Figure 2.

C. Centralized subflow Policy Management

Individual devices are unlikely to have the information necessary to predict how multipath flows will interact in wireless networks and when it would be beneficial to enable or disable certain wireless links. As a solution, we introduce a management server that tracks client channel conditions and dynamically optimizes how multipath is used. In congruence with the principles of Software-Defined Networking (SDN), the management server is highly programmable in order to support the objectives of network operators and users.

Aside from using control over multipath in order to optimize throughput, there are other potential objectives that network operators and users may wish to express. For example, network operators can use fine-grained control to improve the performance of sponsored applications. On the other hand, users themselves may appreciate the ability to make informed decisions about when to use multipath—and thus additional wireless resources which may be billed differently—to accelerate specific applications such as video streams. We incorporate an example of such functionality into our optimization algorithm in the form of user priority classes, e.g. higher priority users receive preferential access to multipath.

An agent running on the client device periodically communicates with the management server in order to coordinate which wireless links should be used for network connections. If a protocol is available for obtaining channel quality and usage information directly from the base stations, then the central manager can obtain such information for computing the optimal link assignments. For our design we do not assume that this information is available from base stations, so clients collect local measurements of channel conditions and provide them during their communication with the management server.

In our design, the management server would be responsible for a single zone such as an enterprise or university campus. Although the optimal link assignment is computationally complex, it is entirely feasible for server-class machines to handle the management and proxy needs for zones of this size.

IV. Policy Responsive Aggregated Multi-band (PRAM) architecture for Capacity Expansion

Figure 3 above illustrates our candidate aggregation architecture. It shows a small cell network consisting of indoor, outdoor as well as large venue based small cells. Each of these small cells supports multiple spectrum bands (indicated by dotted lines of different colors). The end-user device such as a tablet also supports all or subset of these bands and therefore can connect in multiple bands to same or different small cells.

We introduce four new logical entities:

Aggregation Proxy: This network resident server is a transparent proxy that splits the connection from the client devices to the internet hosts/servers (e.g.: web server, YouTube, Hulu, Netflix etc.) into two parts: (1) First part is a multipath connection (e.g.: MPTCP sub flows) from client device – over multiple spectrum bands and passing through a single or multiple small cells to the proxy and (2) Second part is a legacy TCP or MPTCP connection from the proxy to the internet server. The Aggregation Proxy controls a large cluster of small cells (potentially 100s of small cells in a neighborhood or a large venue (e.g.: enterprise building, airport). We term the cluster as the aggregation scope of this proxy. It is essential that the capacity of the link between the small cells and the proxy exceed the maximum capacity of aggregated wireless links.

Policy Server: This is a network resident server that manages the network, device, app and user specific policies aimed at achieving desired objectives such as maximum system capacity, best user experience for select user(s) or application classes and monetization of providing better capacity. Each Policy Server has an associated aggregation proxy and therefore corresponding *aggregation scope* over which it controls the policies. It can leverage measurements from the clients and other information such as network conditions, user profile, payment confirmation etc. to drive a computational algorithm called *GLobal Optimization of Band Aggregation Levels (GLOBAL)* that decides a set of small cells that a given client device should use for its traffic.

Policy Agent: This is a piece of software that runs on the client device and interfaces to the network policy server. The two entities maintain a bidirectional signaling channel. The server can use this channel to: (1) convey policy on which interfaces to use, (2) which apps are allowed to use multiple interfaces, (3) price or similar cost of using multiple interfaces, and (4) request periodic measurements. The client can use it for (1) registration message containing capability description, (2) performance measurements such as instantaneous or time-averaged values of signal strength, transmission rates, or loss rates.

Aggregation controller with policy support: This is an operating system specific software in the client device that translates the policies into actual manipulation of traffic flows over multiple interfaces in the device. In our current Linux/Android implementation this capability is realized using the Netlink socket API that can be used by the policy agent to signal to the modified MPTCP protocol stack to turn on or off subflows for each active MPTCP connection corresponding to an APP, user or entire device.

Note that the Aggregation Proxy and Policy Server can be implemented as physically separate servers or integrated in a single server. Also, the Policy agent and Aggregation

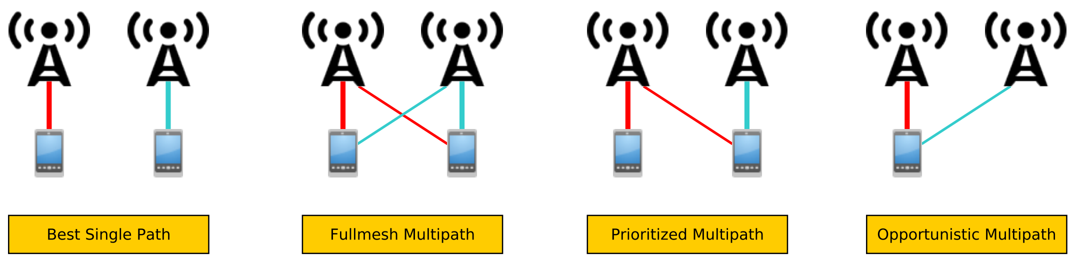

Fig. 2. Four examples of the modes of operation that are possible with multipath-capable devices.

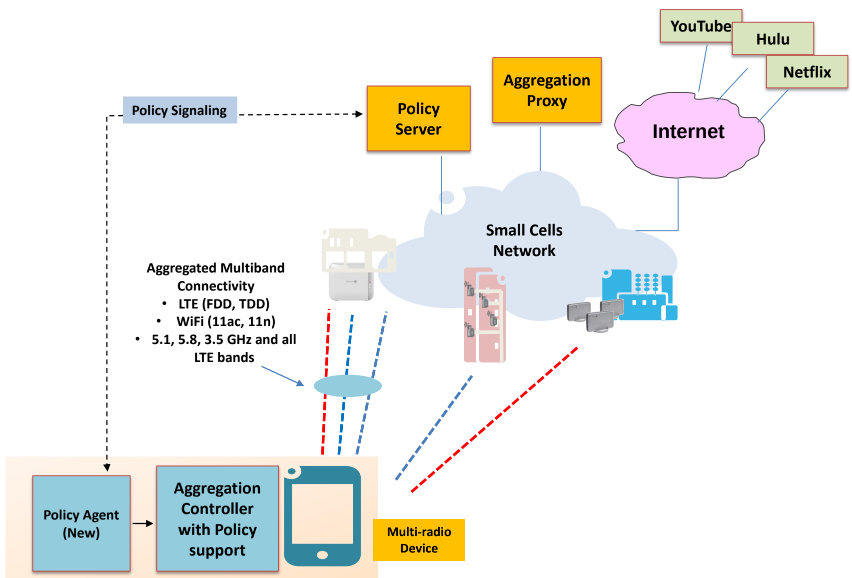

Fig. 3. Policy Responsive Aggregated Multi-band (PRAM) architecture

Controller, can be implemented in a single piece of software with an API that can be used by apps (e.g.: YouTube, Hulu, Netflix apps).

Policy signaling methods: Note that the policy signaling between the Policy Server and Policy Agent can be accomplished in many ways. One possible implementation, illustrated in Figure 4 is using XMPP protocol[11].

A. *Prototype Implementation*

In our experimental prototype, the Aggregation Proxy is implemented as a *Transparent Split Proxy*. Whenever, MPTCP enabled client or server side enables first TCP subflow, the routing is setup such that the TCP connection setup is routed via the Aggregation Proxy. Using packet filtering mechanisms (e.g.: iptables in Linux), the aggregation proxy detects such connection attempt and setups two independent TCP sockets: a client side socket is used to terminate the connection to the client device whereas the other socket – the server socket is used to terminate an independent TCP connection to the internet server/host. Any further subflows between the client and the aggregation proxy are setup under the client socket under the control of the policy manager. Similarly, if the internet host supports MPTCP, the aggregation proxy can initiate multiple MPTCP subflows under that server socket. The policy server can also control these subflows on the aggregation proxy end. Figure 5 illustrates the two sockets and the corresponding MPTCP subflows.

Multipath Policy Programming Support

Our solution relies on new software deployed on mobile devices and within the operator's network. We base our design off the Linux implementation of MPTCP [21]. The current implementation allows the user to configure interfaces to be allowed, disallowed, or backup-only for multipath purposes. This setting is applied system-wide, so an interface may only be enabled or disabled for all MPTCP flows simultaneously. During an MPTCP connection, the endpoints establish subflows on all allowed and backup paths, while a bit in the kernel memory specifies the backup status. Therefore, with some modifications to the code, it is possible to enable/disable

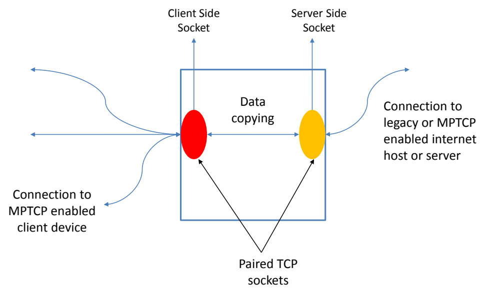

(a) MPTCP sockets towards client and server

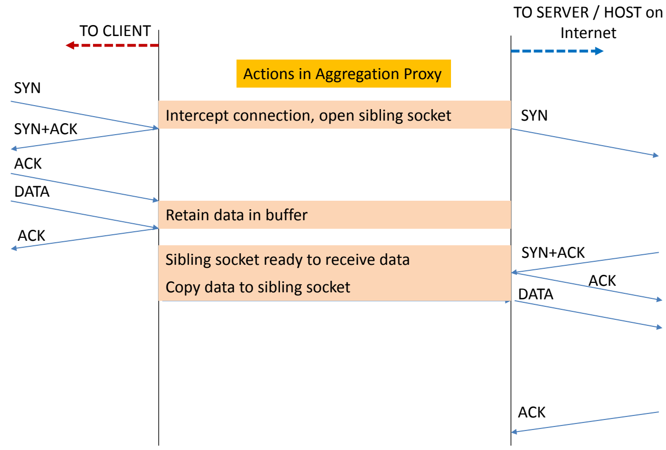

(b) Message exchange for transparent proxying

Fig. 5. Transparent proxy: sockets and message exchange

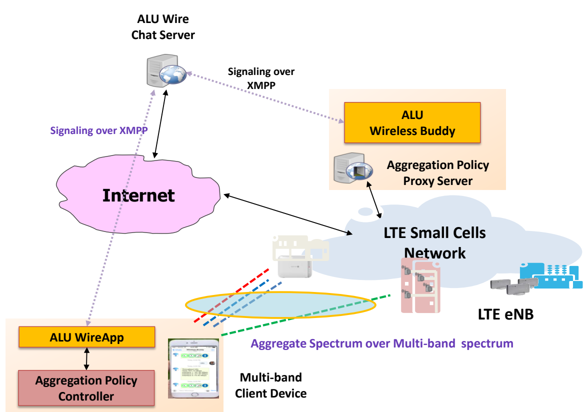

Fig. 4. Chat based policy signaling between network and user device

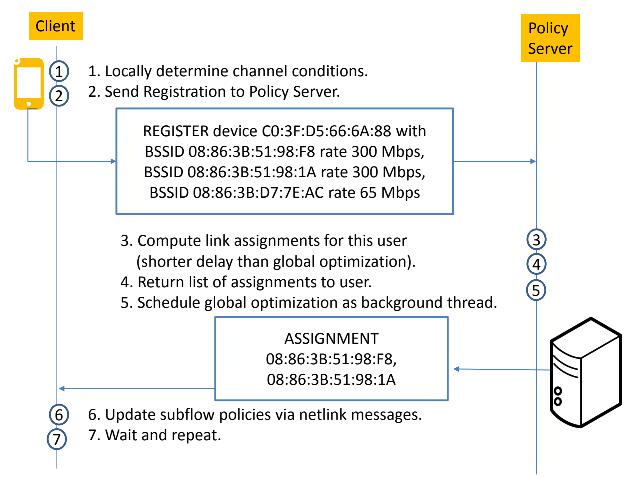

Fig. 6. Message exchange with policy server

a given subflow simply by changing its backup status. We implement an API that allows a user-space tool or an app to manipulate a table of rules for enabling or disabling subflows.

Policy Signaling between the network and the client

Figure 6 illustrates an example of message exchange between the policy agent in the client device and the policy server.

The agent running on the client device tracks the local wireless channel conditions and periodically notifies the management server. The management server uses its knowledge of the network to compute the optimal link assignments for all clients. We implement this as a request-response protocol, which we call *Registration*. The client periodically sends a Registration message to the server containing a unique identifier and list of wireless interfaces, the base stations to which the device is associated, and the average transmit rate.

We implement Registrations using an HTTP POST operation. The immediate response from the server uses just a local version of the optimization algorithm to incorporate the new link quality information into its assignment for the user, which keeps the delay in responding to the user very small. The management server also schedules an iteration of the global optimization algorithm to run in a background thread if it is not already running.

How quickly the management server can respond to changing network conditions depends on several factors.

- Clients that are leaving the system should ideally notify the manager of this fact. In the absence of a notification, the manager must rely on a timeout.
- If updates are distributed solely by clients querying the manager, then the frequency of queries is an important consideration.

- Optionally, the system can be designed such that the manager pushes updates to client devices, which imposes the additional implementation requirement that the client must listen for such messages.

Additionally, the coherence time of measurements plays an important role in finding an optimal solution. Since, we anticipate our solution to be most applicable to enterprise networks with predominantly low-speed mobility, we suggest an update interval on the order of 10s of seconds (e.g.:30 secs) would be appropriate.

B. GLOBAL Optimization Algorithm

The *GLobal Optimization of Band Aggregation Levels* (GLOBAL) algorithm is run at the policy server in the steps 3 and 5 in Figure 6. The GLOBAL formulates flow assignment as an optimization problem, specifically, maximization of total capacity obtained by all users. The goal of the optimization module is to adjust multipath usage to maximize a global utility function given knowledge of wireless channel conditions at the client devices. We choose proportional fairness with priority to be the utility function, since this is known to result in high aggregate throughput without starving any users[12].

Our optimization function is shown below.

$$\max \sum_{i \in \text{users}} \gamma_i \log t_i$$

$$\text{where } t_i = \sum_{j \in \text{links}} C_{ij} a_{ij}$$

$\gamma_i > 0$ is the priority of user i, t_i is the predicted aggregate throughput experienced by user i, C_{ij} is the predicted capacity user can receive on link j, and a_{ij} is a binary variable set to 1 if user i is assigned link j. The predicted capacity C_{ij} is computed dynamically as a function of the channel conditions of all users sharing the link under the context of a channel sharing model. Two models we consider here are the random access and proportional fair models.

$$C_{ij} = \frac{1}{\sum_{i \in \text{users}} \frac{a_{ij}}{r_{ij}}} \qquad \text{(Random Access)}$$

$$C_{ij} = \frac{r_{ij}}{\sum_{i \in \text{users}} a_{ij}} \qquad \text{(Proportional Fair)}$$

Here $r_i j > 0$ is the average transmit rate for user i on link j. For simplicity we only consider uplink capacity here, but in general, both uplink and downlink should be considered particularly in relation to the needs of the application. To be completely general, we need to handle the case in which some links are unavailable to some users, for example because they are out of range or mutually exclusive. In the implementation, we set $r_{ij} = C_{ij} = a_{ij} = 0$ for links j unavailable to user i. Likewise, access points with no associated users or users with no available links are special cases to be handled appropriately.

The use of binary decision variables and a nonlinear objective function make finding the global maximum intractable for a large number of users. Since, the problem is NP-complete and computationally intractable for large networks, we solve it using an approximate technique called *Hill Climbing*. Alternate methods such as simulated annealing can also be used.

We start from the initial state of all usable links assigned ($a_{ij} = 1$ if $r_{ij} > 0, 0$ otherwise). One iteration of hill climbing walks through the list of user and link combinations in random order and adds or removes an assignment if the change will improve the global utility function.

V. EXPERIMENTAL EVALUATION

We begin with an experiment to evaluate how MPTCP performs under the scenario depicted in Figure 1. In this scenario, two client devices each have access to two access points; however their channel conditions with the two access points are very different. We hypothesize that enabling multipath in this scenario will lead to performance degradation for both clients unless somehow they can discover a more optimal allocation. This scenario resembles the famous Prisoner's Dilemma[24]. Both clients, seeing that they have two wireless links available to them will independently determine that it is best to make use of both links for bandwidth aggregation. Given the known performance issues with random access wireless communication, the client with a weaker signal will consume a large amount of airtime by transmitting at a low data rate and thus leave the other client unable to make full use of his more favorable channel. We test this scenario in hardware, which then gives us an experimental basis for using our central manager to coordinate client access to multipath.

Having demonstrated the possibility for the central manager to improve network efficiency in a static scenario, we next explore how it responds in dynamic scenarios. Suppose in the previous scenario, one of the clients leaves the system by powering off or physically moving. Then we would like the remaining client to gain access to the resources that have been freed. In particular, it should begin using the access point that was previously reserved for the client that has now left the system.

Finally, we consider what happens when the clients are assigned to different service classes and when the asymmetry in channel conditions is more or less severe. Under this more realistic scenario is when the central manager reveals its full benefit. By considering the channel conditions of all clients in the system in addition to their service classes, the central manager is able to strike a good balance between the competing goals of maximizing system capacity and providing a fair but differentiated level of service to each client.

A. Platform

Here we describe the experimental platform. We use 802.11n links with 20 MHz channels for our experiments, although there is little reason to expect our results not to generalize to other MAC layers. A single multi-band (2.4 GHz and 5.7 GHz) access point is used to serve two multi-radio clients. The clients are connected by RF cables and splitters to reduce impact from external Wi-Fi devices in the vicinity while still allowing the two client devices to contend with each other for access to either channel. The clients are configured with different levels of attenuation to create relatively strong and weak channel conditions. We change the amount of attenuation between two scenarios, which are detailed in Table II. In particular, the difference between the strong and weak channels is more pronounced in the *challenged* scenario as compared to the *favorable* scenario. We induce these different channel conditions in order to create scenarios in which multipath clients may contend with other clients with stronger or weaker links than their own. It represents a physical implementation of the scenario envisioned in Figure 1.

TABLE II
CLIENT CHANNEL QUALITIES IN THE CHALLENGED AND FAVORABLE CONDITIONS AVERAGED OVER 1000 MEASUREMENTS.

	Challenged Mean RSSI (SD)		Favorable Mean RSSI (SD)	
	2.4 GHz	5.7 GHz	2.4 GHz	5.7 GHz
Client A	-84.6 (3.2)	-36.1 (2.2)	-55.0 (1.2)	-37.2 (2.5)
Client B	-33.8 (3.1)	-83.4 (1.0)	-35.9 (1.7)	-58.7 (1.3)

B. Results

Each iteration of the experiment takes 60 seconds during which both clients simultaneously run MPTCP-enabled `iperf` upload tests. Client A runs `iperf` for the entire 60 seconds, while Client B runs `iperf` for the first 30 seconds only. We record the traffic at the receiver side in one-second intervals.

- **Single Best:** Each client uses only its strongest link during the test, which means that the clients never contend with each other for channel access.
- **Unmodified MPTCP:** Each client uses MPTCP with the fullmesh path manager to create an active subflow on each link during the test. Decisions about whether to use the stronger or weaker link are therefore handled by the MPTCP packet scheduler.
- **Optimized / Equal Priority:** Each client's subflows are dynamically optimized according to global network conditions. During the first 30 seconds, the behavior mimics the **Single Best** strategy, but after Client B leaves, both subflows are activated for Client A.
- **Optimized / B Prioritized:** Once again the clients' subflows are optimized for global conditions, but in this case, Client B is assumed to be a higher priority user. For the duration of Client B's transfer, both subflows are activated, whereas Client A only uses its stronger link. After Client B leaves, Client A is given access to both links.

We run 20 iterations under each condition and average the measurements for each one-second interval to obtain the results shown in Figure 7. The results of our experiment support our idea to selectively enable multipath when it would be beneficial to do so. In the challenged wireless scenario, each client has one very good link and one very poor link. Under those conditions, it is only beneficial to aggregate the two links if one device has exclusive access to the poor link, e.g.

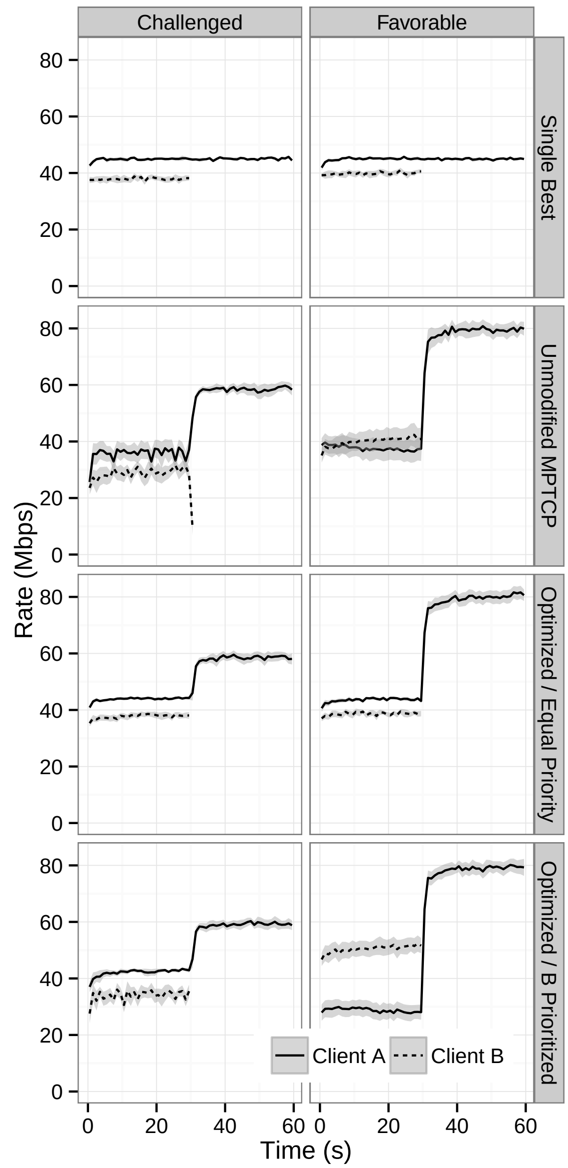

Fig. 7. Throughput in iperf tests averaged over 20 tests for each of eight different conditions.

Client A after Client B has left. Our results show that our optimization approach produces high aggregate throughput than both Single Best (9% higher) and Unmodified MPTCP (13% higher) and has a small negative effect when prioritization is enabled. In the favorable wireless situation, each client has two good links, so we find larger gains from aggregation. Our optimization approach produces higher aggregate throughput than Single Best (24% higher) and is effectively similar to Unmodified MPTCP, and our prioritization scheme successfully boosts the targeted flow's performance without harming aggregate throughput. Particularly notable findings are that unmodified MPTCP is detrimental to aggregate throughput in the challenged scenario, and prioritization is only beneficial in the favorable scenario. This behavior is what necessitates

TABLE III
AVERAGE AGGREGATE THROUGHPUT FOR TWO CLIENTS UNDER DIFFERENT CONDITIONS WITH 95% CONFIDENCE INTERVALS.

	Challenged Rate (Mbps)	Favorable Rate (Mbps)
Single Best	63.9 ± 0.3	64.9 ± 0.2
Unmodified MPTCP	61.4 ± 1.0	78.3 ± 1.3
Optimized / Equal Priority	69.8 ± 0.7	80.7 ± 1.1
Optimized / B Prioritized	67.5 ± 0.9	78.9 ± 1.2

consideration of channel conditions by the policy server.

VI. RELATED WORK

Using multiple network links for bandwidth aggregation is a well-known problem. Proposed solutions generally work at either the transport layer, e.g. MPTCP [21], or the link layer, e.g. MultiFacet [22]. To the best of our knowledge, no prior work has explored how multiple wireless devices interact when using multipath.

Other related work includes projects that try to optimize network selection, primarily between Wi-Fi and cellular on smart phones, e.g. Deng, et al. [15] and MultiNets [19]. Although these works are based on the best single link paradigm, they are complementary to our work. By relaxing the assumption that only a single wireless link will be used at any time, we may apply the insights from network selection works to guide selection of the best subset of wireless links for a given type of user traffic.

Work by Ling et al. [18] proposes integration of Wi-Fi and LTE networks, at different layers – Transport, RAN and MAC - with the flexibility of choosing the access that is best suited depending on the direction (uplink or downlink), type of traffic. Their proposed schemes blend Wi-Fi and LTE systems in a way that leverages bandwidth advantage of Wi-Fi and LTE's advantage of predictable QoE due to scheduled access, to overcome the limitations of each access when used individually. Our policy based framework can be extended to complement and enhance performance of these schemes.

VII. FURTHER DISCUSSION

We feel that the time is right for mobile devices to begin exploiting aggregation of their multiple wireless interfaces. With the availability of an open source and stable MPTCP implementation along with the assistance of a network-resident translation proxy, this is entirely feasible in the near term. However, our work suggests that large scale deployment of MPTCP in wireless networks will result in complex interactions between devices, so we propose the simultaneous deployment of a network-resident management server responsible for ensuring that users receive the benefits of multipath within the context of operator-defined objectives.

Our proposed architecture leaves room for much interesting future work. We advocate for an expressive policy language that allows management of subflows at the granularity of applications or network connections rather than at the system level. This leaves the possibility for the network operator to accelerate certain preferred applications. We also envision

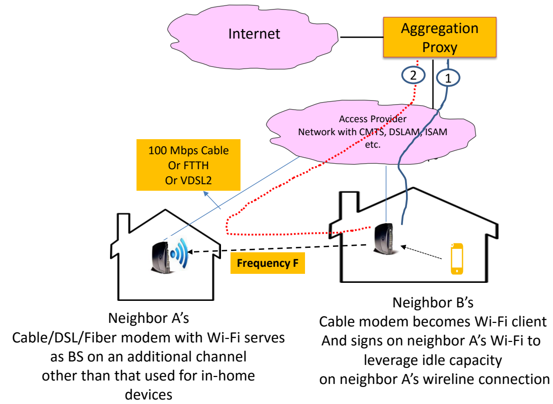

Fig. 8. Policy driven link aggregation for ultra-broadband fixed wireless access enhanced wireline access

user interaction with the system, perhaps prompting the user to make an informed cost-benefit decision with regards to accelerating a poorly-performing application.

In the traditional model of deploying small cell network, the aggregation proxy and policy server would be owned and managed within an operator network. However, since these servers are completely independent of the radio access network, they can be owned and operated in a cloud by a third party. This can open up new business models for "managed aggregated multi-band" connectivity.

Aggregation in Fixed Wireless Access (FWA): Enhancing Cable /DSL/FTTH Internet Access

Figure 8 illustrates how PRAM can be used to augment wireline access such as cable, DSL and even Fiber-to-Home networks. Typically, a cable/DSL/FTTH modem either has a co-located or integrated Wi-Fi access point that provides in-home Wi-Fi access. The newer installations of such modems even support dual band Wi-Fi connectivity in 2.4 GHz as well 5.8 GHz Wi-Fi.

When a neighbor is under utilizing its cable / DSL / Fiber internet link, it can serve as a backhaul for neighbors cable / DSL/ FTTH modem. In our example above, Neighbor Lbs cable modem is experiencing high load and needs to augment its internet connection to a speed higher than the peak capacity offered by its wireline cable connection. To do so, it turns on its Wi-Fi interface in client mode and associates with neighbor A's Wi-Fi access point. The traffic from the Neighbor B is now aggregated over two network paths: (1) *Path 1*: which is a direct path over wireline connection from Neighbor Lbs modem to the cable access network and (2) *Path 2*: where the traffic goes over a wireless hop to Neighbor A's access point and then sent over Neighbor A's wireline internet connection. The aggregation of the traffic happens in the aggregation proxy installed in the network much the same as in our discussion before.

Note that in the example above, we can upgrade Wi-Fi

APs / small cells to support additional shared bands or LTE in licensed or shared bands. Also, we can employ a policy mechanism to allow such aggregation to turn on (1) only for households that have signed up for a *boost* plan, (2) when demand for capacity in the household increases (e.g.: increase in the number of Netflix connections), and (3) spare capacity is available in neighbors internet connection.

LTE/Wi-Fi Integration, Alcatel-Lucent BOOST and AggregatedMultiband (AM)) Currently, topic of LTE and Wi-Fi integration is of great interest as Wi-Fi and LTE networks are becoming ubiquitous and a large number of mobile devices support both network accesses. With such integrated access, also termed Dual Connectivity (DC), client device connected to LTE network can seamlessly switch over to Wi-Fi networks when they become available and vice versa. 3GPP has standardized a well designed framework called Access network discovery and selection function (ANDSF) [9]. This function implemented within the LTE evolved packet core (EPC), assists user equipment (UE) to discover and use non-3GPP access networks – such as Wi-Fi or WIMAX – for data communications in addition to 3GPP access networks (such as HSPA or LTE) and to provide the UE with rules policing the connection to these networks. Our architecture supports functionality similar to ANDSF to enable seamless roaming across LTE and Wi-Fi as a simple matter of policy insertion. Specifically, when a client device reports visible LTE and Wi-Fi networks it can connect to, the Policy Server can insert a policy in the client devices to mark one of the network connections as active and other as backup. In the event the device loses active network connection, it changes the status of backup connection to active and toggles current active connection to backup. Of course, PRAM also supports a mode, where traffic can be concurrently sent and aggregated across both interfaces.

Alcatel-Lucent recently announced enhanced LTE and Wi-Fi integration strategy called 'Wireless Unified Networks' (WUN) [1] which aims to blend best capabilities of Wi-Fi® and cellular technologies. Central to this strategy is the concept called BOOST [1] which is an IP layer integration of LTE and Wi-Fi that relies on simultaneous and complementary use of the two wireless technologies. BOOST overcomes limitations of Wi-Fi in Uplink (UL) and Cellular Downlink (DL) to provide wire-line like experience. It has two variants – the *Wi-Fi boost* which combines cellular uplink with wi-fi downlink and the Cellular boost which uses unlicensed spectrum downlink to enhance the performance of cellular. The technology has wide applicability since it requires no changes to the incumbent Wi-Fi infrastructure and can be delivered onto end-user devices as an OS/App upgrade. Our PRAM architecture can realize both versions of BOOST and any of its variation using appropriate policy insertion. For example, Wi-Fi BOOST can be implemented using policy installed at the client device to send uplink ACKs for all MPTCP connections only on LTE (cellular) interface and relaying to other end-point that the LTE downlink is in "backup" mode.

Comparing Dual Connectivity (DC), Carrier Aggregation (CA) and Aggregated Multiband (AM) The DC, CA and AM techniques require multiple radio interfaces to be concurrently powered on, which has implications on the battery life of UEs.

DC allows a RAN layer integration of Wi-Fi and LTE networks. While it allows control of bearer distribution across the LTE and Wi-Fi in a way that is responsive to RAN conditions, it creates dependency between LTE and Wi-Fi RAN infrastructure (control and data plane interfaces) and requires changes at the protocol stack. Note that Wi-Fi interface can use Power Save (PS) sleep feature available in some (especially enterprise class) 802.11 access points (APs) to improve battery life. This impacts session handoff and recovery latency when LTE connection disconnects (e.g.: due to mobility) and Wi-Fi interface needs to be reactivated.

In the layer-1/layer-2 Carrier Aggregation (CA) LTE technology, the base station scheduler dynamically maps traffic at different times to group of OFDMA sub-carriers (commonly termed as Physical Resource Blocks (PRB)) in various portions of the aggregated spectrum bands. Though radio frontend for each of the bands need to be powered on, the power consumption in each RF chain increases above baseline only when PRBs are used for transmission. Since, the scheduler controls PRB allocation at a very fine time scale (of msecs), CA is very responsive in accurately tracking traffic variation and connection state of each interface/band.

On the contrary, being a transport layer technique, AM relies on end-to-end MPTCP acks to discover on/off state of each interface and as such takes minimum one round trip time to respond to changes in connectivity. In well engineered small cell environments with sufficient backhaul capacity, and low or portable mobility, round trip times are of the order of tens of milliseconds and as such AM technique can perform quite well but is slower than CA. Since, AM is aimed at higher capacity, it consumes battery much faster than RAN layer mechanisms such as DC, CA due to its slow response to path conditions.

The most attractive feature of AM is that unlike CA, it is a software intensive, future proof technology that can exploit any spectrum band and baseband available in the device. If the UE memory and CPU systems are not bottleneck, it can reflect capacity increases from future advances in layer-1/layer-2, such as LTE-Advanced, Massive MIMO, higher channel widths and CA into sustained peak capacity increase for applications, without any changes to software and aggregation architecture.

Policy controlled prioritization and net neutrality The notions of preferential service where certain class of devices, applications and traffic types are given higher aggregated capacity may appear to contradict principles of net neutrality[7]. We note that our mechanisms can be used in controlled domains such as private enterprises, verticals (e.g.: hospitals, public safety, emergency response, oil and gas) and private domains (such as large venues, stadiums, malls, hotels etc.) without any concerns of net neutrality. Also, net neutrality rules can be interpreted to strictly apply to individual physical links and not necessarily apply to virtual links that are aggregation of multiple physical links. Our policy mechanisms allows creation of sustained peak performance levels much the same as different levels of peak data rate plans offered by cable and fiber access providers (e.g.: 25/25 Mbps, 50/50 Mbps etc.) which are priced differently. As such, policy controlled AM technology can be viewed as largely agnostic to net neutrality debate.

VIII. Conclusion

We presented a multi-band spectrum aggregation technology that enables small cell networks with ultra-high capacity by transparently exploiting current and future spectrum bands and leveraging advances in baseband such as LTE-A, carrier aggregation, shared spectrum band and mmWave links as and when they appear.

By introducing central management of multipath usage via the proxy manager, our solution to the problem of creating high access capacity can ensure better utilization of scarce spectrum resources and enable a novel mechanism for prioritization of different traffic classes which can be monetized.

Our solution is potentially valuable to any organization that operates a large number of small cells and wishes to enable its users to exploit multipath in a controlled way for higher performance. We believe our technologies have potential to open a new way of offering aggregated high capacity connections and "always best connected" service and can enable new business models.

References

[1] Alcatel-lucent wireless unified networks (wun). https://www.alcatel-lucent.com/solutions/wireless-unified-networks.

[2] Benefits of lte in unlicensed spectrum, various presentations. http://www.3gpp.org/news-events/3gpp-news/1603-lte_in_unlicensed.

[3] Introducing multefire: Lte-like performance with wi-fi-like simplicity. https://www.qualcomm.com/news/onq/2015/06/11/introducing-multefire-lte-performance-wi-fi-simplicity.

[4] Linux kernel sources. tcproxy. http://www.kernel.org/doc/Documentation/networking/tcproxy.tx.

[5] LTE-U Forum. http://www.lteuforum.org/.

[6] LTE-U Forum Documents. http://www.lteuforum.org/documents.html.

[7] Net neurtality. https://en.wikipedia.org/wiki/Net_neutrality.

[8] Open vpn. htp://openvpn.net.

[9] Ts 23.402 architecture enhancements for non-3gpp accesses. http://www.3gpp.org/DynaReport/23402.htm.

[10] Virtual tunnel (vtun). http://vtun.sourceforge.net.

[11] P. S. andre ed. Extensible messaging and presence protocol (xmpp): Core. *IETF RFC 3920*, Oct. 2004.

[12] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, P. Whiting, and R. Vijayakumar. Providing quality of service over a shared wireless link. *Communications Magazine, IEEE*, 39(2):150–154, Feb 2001.

[13] A. V. Bakre and B. R. Badrinath. I-TCP: indirect TCP for mobile hosts. In *Proceedings of the 15th International Conference on Distributed Computing Systems, Vancouver, British Columbia, Canada, May 30 - June 2, 1995*, pages 136–143, 1995.

[14] C. Caini, R. Firrincieli, and D. Lacamera. Pepsal: a performance enhancing proxy designed for TCP satellite connections. In *Proceedings of the 63rd IEEE Vehicular Technology Conference, VTC Spring 2006, 7-10 May 2006, Melbourne, Australia*, pages 2607–2611, 2006.

[15] S. Deng, A. Sivaraman, and H. Balakrishnan. All your network are belong to us: a transport framework for mobile network selection. In *15th Workshop on Mobile Computing Systems and Applications, HotMobile '14, Santa Barbara, CA, USA, February 26-27, 2014*, page 19, 2014.

[16] R. Khalili, N. Gast, M. Popovic, U. Upadhyay, and J.-Y. Le Boudec. Mptcp is not pareto-optimal: Performance issues and a possible solution. In *Proceedings of the 8th International Conference on Emerging Networking Experiments and Technologies*, CoNEXT '12, pages 1–12, New York, NY, USA, 2012. ACM.

[17] M. Kojo, J. Griner, and Z. Shelby. Performance enhancing proxies intended to mitigate link-related degradations", rfc 3135, 2001.

[18] J. Ling, S. Kanugovi, V. Subramanian, and K. A. Parmod. Enhanced capacity and coverage by wi-fi lte integration. *IEEE Communications Magazine*, pages 165–171, Mar. 2015.

[19] S. Nirjon, A. Nicoara, C.-H. Hsu, J. P. Singh, and J. A. Stankovic. Multinets: A system for real-time switching between multiple network interfaces on mobile devices. *ACM Trans. Embed. Comput. Syst.*, 13(4s):121:1–121:25, Apr. 2014.

[20] C. Raiciu, S. Barre, C. Pluntke, A. Greenhalgh, D. Wischik, and M. Handley. Improving datacenter performance and robustness with multipath tcp. In *Proceedings of the ACM SIGCOMM 2011 Conference*, SIGCOMM '11, pages 266–277, New York, NY, USA, 2011. ACM.

[21] C. Raiciu, C. Paasch, S. Barré, A. Ford, M. Honda, F. Duchene, O. Bonaventure, and M. Handley. How hard can it be? designing and implementing a deployable multipath TCP. In *Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2012, San Jose, CA, USA, April 25-27, 2012*, pages 399–412, 2012.

[22] S. Sen, M. Griepentrog, J. Yoon, and S. Banerjee. A case for enhancing dual radio repeater performance through striping, aggregation, and channel sharing. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, MobiCom '14, pages 569–580, New York, NY, USA, 2014. ACM.

[23] D. Tse. Multiuser diversity in wireless networks, 2001.

[24] Wikipedia. Prisoner's dilemma. https://en.wikipedia.org/wiki/Prisoner%27s_dilemma.