SoftRAN: Software Defined Radio Access Network

Aditya Gudipati*, Daniel Perry*, Li Erran Li†, Sachin Katti*
*Stanford University
†Bell Labs, Alcatel-Lucent

ABSTRACT

An important piece of the cellular network infrastructure is the radio access network (RAN) that provides wide-area wireless connectivity to mobile devices. The fundamental problem the RAN solves is figuring out how best to use and manage limited spectrum to achieve this connectivity. In a dense wireless deployment with mobile nodes and limited spectrum, it becomes a difficult task to allocate radio resources, implement handovers, manage interference, balance load between cells, etc.

We argue that LTE's current distributed control plane is suboptimal in achieving the above objective. We propose SoftRAN, a fundamental rethink of the radio access layer. SoftRAN is a software defined centralized control plane for radio access networks that abstracts all base stations in a local geographical area as a virtual big-base station comprised of a central controller and radio elements (individual physical base stations). In defining such an architecture, we create a framework through which a local geographical network can effectively perform load balancing and interference management, as well as maximize throughput, global utility, or any other objective.

Categories and Subject Descriptors

C.2[**Network Architecture and Design**]: Wireless Communication; C.2[**Network Operations**]: Network Management;

General Terms

Design, Management

Keywords

Radio Access Networks, Software Defined Networking

1. INTRODUCTION

Wireless infrastructure is becoming increasingly chaotic and dense [4]. This is driven by the need to support exponentially increasing mobile traffic and the fact that spectrum is limited. Consequently one main mechanism to increase capacity is making networks dense (otherwise referred to as cell splitting). By bringing the infrastructure closer to the client, networks can theoretically improve link quality to each user and reduce the number of users that each base station needs to support. However, spectrum is a scarce resource. For example, in the US, AT&T and Verizon both have less than 100MHz of spectrum nationwide that they can use for LTE. Due to the lack of spectrum, neighboring base stations in a dense deployment have to operate on the same channel. This is referred to as deploying networks with a frequency reuse factor of one.

Managing dense networks with frequency reuse one is significantly more complex due to a tight coupling in control plane decision making at neighboring base stations. That is, the radio resource management decisions (i.e. deciding what spectrum to use to transmit at what power to which client) made at one base station have substantial impact on neighboring base stations and vice versa. This coupling in control plane decision making manifests itself in two ways. First, the high frequency reuse and broadcast nature of wireless communication lead to clients of one base station experiencing significant interference from neighboring base stations. Left unmanaged, this interference significantly degrades capacity. Second, due to their smaller coverage areas, load fluctuates more rapidly due to user mobility. Consequently handovers, cell association, and resource (spectrum) allocation have to be managed at each base station in concert with its neighbors to maximize the network capacity by carrying out tasks such as interference management, load balancing, etc.

Traditionally, the radio access network has been treated as a collection of base stations, each largely making independent control plane decisions on the radio layer with some loose distributed coordination via mechanisms such as SON (self organizing networks), ICIC, etc. However, due to their small cell sizes, in dense networks coordinated control plane decisions have to be made across several neighboring base stations simultaneously–often with as low of a latency as possible. Distributed coordination algorithms do not scale well as they need to work with larger number of base stations, especially in terms of latency. This leads to poor performance, reducing capacity significantly due to the inability to manage interference and balance load. Further, distributed coordination algorithms tend to become more complex, since often they require iterative and periodic adjustment of radio layer allocation decisions that are hard to get right at scale.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HotSDN'13, August 16, 2013, Hong Kong, China.

Copyright 2013 ACM 978-1-4503-2178-5/13/08 ...$15.00.

In this paper, we propose a fundamental rethink of the radio access layer. Instead of looking at the radio access layer as a collection of independent base stations, we argue that all base stations deployed in a geographical area should be abstracted as a virtual big-base station which is made up of radio elements (the individual physical base stations). A logically centralized control plane makes all decisions regarding handovers and interference management, while the radio elements are simpler devices with minimal control logic. Instead of naively assuming that each base station has its own set of resources, we acknowledge that all neighboring base stations are allocating from a fixed set of shared resources. Thus, we abstract the radio resources as a three dimensional grid of space, time, and frequency slots; and program them in a software defined fashion through a logically centralized radio access control plane. We define APIs between the radio elements and the control plane which allow radio elements to update the global view of the network at the control plane and also allow the control plane to communicate radio resource management decisions back to the radio elements.

The biggest challenge in creating such a software defined radio access network is the inherent delay between any centralized controller and the individual radio elements. Typically, the individual radio elements are connected to a centralized controller by backhaul links (wired or wireless) that have a latency of 5-10ms. This inherent delay between the controller and the radio element implies that the radio element has a more updated view of the local state and can, in certain scenarios, better manage its resources locally. We cannot simply expect the controller to perfectly allocate resources over long time scales because of rapidly varying channel conditions at the radio elements themselves. To address this challenge and ensure that the radio element is given this opportunity to adjust to rapidly varying channel conditions, SoftRAN *refactors* the control plane functionalities between the radio element and the controller. The centralized controller makes decisions that affect the global network state, while, each individual radio element handles local control decisions which do not affect other neighboring radio elements.

The cellular industry has recognized the difficulty in managing interference, load, etc., and as a result, 3GPP (the LTE standards body) is working towards defining coordination mechanisms such as e-ICIC (enhanced inter cell interference coordination). In SoftRAN, we present a systematic architecture for realizing such mechanisms in a modular fashion. Instead of having to rely on and wait for standardization each time a new protocol is developed, we envision being able to rapidly implement such protocols through our software-defined architecture.

We present a preliminary design and architecture of SoftRAN, use cases, and a feasibility analysis. We show that controller scalability is not an issue; a single node can easily handle all the base stations in a geographical area.

2. DESIGN

The main functionality of the radio access network is to manage the radio resources at base stations in order to provide wide-area wireless connectivity to the mobile clients. Note that within this scope the RAN can work towards various objectives including but not limited to maximizing throughput, minimizing delay, ensuring fairness between flows, etc. Because of this, we will remain general when referring to the exact 'RAN objective' that the RAN is trying to achieve. The RAN seeks to meet this objective by taking the following actions with regard to the base stations' data plane:

- Perform handovers
- Allocate a group of resource blocks to each client (more specifically, each flow). LTE uses OFDM, where radio resources can be assumed to be comprised of time and frequency slots, called resource blocks. The channel quality and the interference for each client can vary across these resource blocks; hence, they must be smartly shuffled between the contending clients
- Assign transmit power values to each resource block at each base station while conforming to overall transmit power constraints. This decision needs to be made while keeping in mind the interference caused at neighboring cells.

These aspects can be termed 'control decisions' and together they form the 'control plane' of the radio access network. Due to the broadcast nature of wireless, the control decisions taken at one base station affect the decision making and performance at neighboring base stations. That is, the control decisions across neighboring base stations are coupled with each other.

2.1 Coupled control plane in dense networks

Current Leta's distributed control plane is designed with sparse deployments in mind. Allowing each base station to make its own radio resource management decisions is intuitive when the decisions have no effect on adjacent base stations (since they are not close enough). This is not the case, however, for dense deployments. With users and base stations collocated in a small geographical area, interference and client mobility motivate the need for coordination.

Let us first take a simple scenario of two clients, each being served by a unique base station. Each base station will independently decide a transmit power to use based on the quality of service required. However, because the base stations and clients are all close to each other, the transmit powers chosen by a base station will determine the interference seen by the client of the neighboring base station. As a result, the transmit powers used by the base stations are dependent on each other and must be arrived at in coordination with each other. With a higher density of base stations, this interference matrix becomes increasingly complicated and ill-suited for distributed management. Assuming that each base station needs to communicate its control decisions to all its neighboring base stations, the amount of control signaling increases quadratically.

Also, in dense deployments, clients will spend more time at and near cell boundaries. As a result, handing over clients to neighboring base stations becomes critical in ensuring good link quality for all clients. Handovers are also critical in balancing load across neighboring base stations, by handing over clients from over-loaded base stations to neighboring under-loaded base stations. While handover decisions in sparse deployments can be straightforward given the lack of candidate base stations, this number of candidate base stations grows quickly in dense deployments. Additionally, the signal strengths that a client observes to multiple base stations will be more comparable in a dense deployment. Hence, the traditional technique of simply handing a client over to the base station providing the highest signal strength fails to grasp the nuances of dynamic load in this chaotic environment. Coordination is necessary to manage handovers in such dense and dynamic environments.

Along with the scalability issues in control signaling, distributed coordination also makes the job of a network operator more difficult, as he or she now needs to implement and debug control algorithms distributed across multiple base stations.

We believe that we must move beyond distributed algorithms to handle control plane decisions effectively. The chaotic nature of mobile communications necessitates that base stations consider their local environment and work together to achieve greater total network utility.

2.2 Big Base Station Abstraction

SoftRAN proposes a centralized architecture as an alternative to the distributed control plane currently implemented in LTE networks. It abstracts out all the base stations deployed in a geographical area as a virtual big-base station while considering all the physical base stations as just radio elements with minimal control logic.

These radio elements are then managed by a logically centralized entity which makes control plane decisions for all the radio elements in the geographical area. We call this logically centralized entity, the *controller* of the big base station. The controller maintains a global view of the radio access network and provides a framework on which control algorithms can be implemented.

The radio resources in the network are abstracted out as a *3D resource grid*: base station index, time and frequency. That is, a decision needs to be made about each time-frequency slot at each base station. For each block on this 3D grid, the controller needs to assign the transmit power used and the flow that would be served.

Thus, from the network operator's perspective, all the radio elements in a geographical area can be conceptually thought of as a single big base station with a 3 dimensional resource grid. We believe, this abstraction would greatly simplify design and implementation of control algorithms in a dense network.

However, from the client's perspective, we cannot strictly achieve a 'single base station' abstraction without changes to the LTE standards. The client will continue to sense multiple base stations and at each handover, will need to carry out the traditional handshakes with both the previous and the new base station. However, we believe that centralized control would lead to smoother handovers and reduce dropped connections and ping-pong (multiple handovers between the same pair of base stations). With a smoother control on transmit powers across multiple base stations and better management of interference, the clients would also experience a more stable connection especially in mobile scenarios.

Conceptually, SoftRAN is a software-defined approach to the Radio Access Network. Analogous to traditional SDN, SoftRAN, proposes to abstract out the control plane from individual nodes in a network and instead logically centralize it. This logically centralized entity then maintains a global network view and provides a framework on which control algorithms can be implemented in a modular fashion.

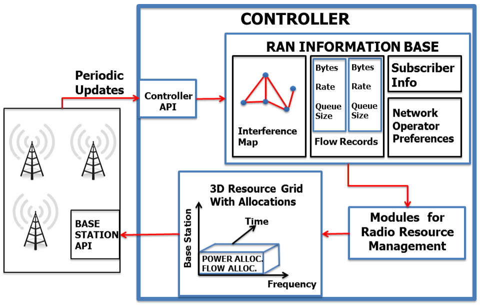

Figure 1: SoftRAN architecture

2.3 Realizing Big Base Station Abstraction

SoftRAN achieves the big base station abstraction architecturally as shown in Fig.1. Realizing such an architecture has two main challenges :

- Designing a controller which can provide a framework for different control algorithms to operate on.
- Ensuring that the delay between the controller and the radio element does not negatively impact performance.

2.3.1 Controller Architecture

As shown in Fig.1, a centralized controller is deployed, which receives periodic updates of local network state from all the radio elements in a 'local geographical area'. Given these updates, the controller updates and maintains the global network state in the form of a database, which we call the 'RAN Information Base'. The RIB (short for RAN Information Base) conceptually consists of the following elements:

- Interference map: A weighted graph, where each node represents a radio element or an active client in a geographical area and the weight of the edges represent the channel strength between the two nodes.
- Flow Records: A record of the relevant parameters of an ongoing flow, e.g. number of bytes transmitted, average transmission rate, number of packets queued, etc.
- Network Operator Preferences: In case, the network operator needs to prioritize certain flows over others, he can enter his preferences into the RIB.

The RIB is accessed by the various control modules (deployed by the network operator), which take the decisions needed for radio resource management. That is, they assign groups of resource blocks to clients, while simultaneously specifying transmit powers to be used by the radio elements (physical base stations) at each resource block. To get a sense of the scale involved, one can assume that the 'geographical area' refers to the range of a macrocell (range of few kilometers) and encapsulates about 20 microcells.

2.3.2 Refactoring the Control Plane

The inherent delay between the controller and the radio elements implies that the radio element has a more updated view of the local state. Thus, in spite of the coordination that a centralized controller provides, the control decisions which depend on rapidly varying network parameters can only be optimized at the radio element. Hence, there is a need to refactor the control functionality between the centralized controller and the radio elements.

There are two main principles guiding the refactoring of the control plane:

- All control decisions that influence the decision making at neighboring radio elements must be made at the controller, since such decisions need to be coordinated across radio elements.
- All decisions that are based on frequently varying parameters should preferably be made at the radio element, since the inherent delay between the radio element and the controller increases the response time to these frequently varying parameters.

Handovers: Handovers clearly fall under our first principles: ‘influence the decision making at neighboring radio elements’, as handing over a client to a neighboring radio element increases the load of the neighbor.

Transmit Powers: Transmit power being used by a particular radio element influences the interference perceived by neighboring radio elements. Hence, transmit power allocation per resource block also need to be handled at the controller.

Resource Block Allocation: On the down-link, resource block allocation does not have an impact on the decision making at neighboring cells. As long as the transmit powers used by the radio element are known, the neighboring radio elements do not need to know which clients are being served with each resource block. Moreover, the resource block allocation among contending clients will be influenced by the channel measurements reported by these clients which can be as frequent as 2 milliseconds.

Hence, the radio element has a more updated view of the channel and this decision falls under our second principle: ‘decisions based on frequently varying parameters.’ Thus, resource block allocation on the downlink can be done by the radio element. However, the controller will still suggest a downlink resource block allocation which should be followed unless the radio element observes an updated view of the wireless environment.

Notice however, that on the up-link, the scenario is exactly the opposite. The up-link resource block allocation decides the client which would be transmitting on each particular resource block. This in turn, will affect the up-link interference as seen by the neighboring radio elements. Hence, according to our first principle, the up-link resource block allocation is made at the controller.

To summarize, the radio element is only responsible for updating the downlink resource block allocation, if and when it receives updated information about the wireless channel strength from the clients. All other control decisions are handled by the logically centralized control plane.

3. USE CASES

We now present a few example scenarios that illustrate some of the advantages of having a logically centralized control plane. The global view of interference and load allows better management of radio resources and helps achieving the RAN objectives.

3.1 Load Balancing

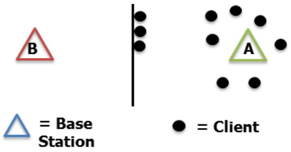

Figure 2: Load balancing

First, we examine a case of straightforward load balancing to increase throughput. Fig.2 depicts a scenario in which base station A (BS_A) is overloaded with clients while base station B (BS_B) is not serving any clients. Additionally, BS_A is serving three edge users who observe a lower signal quality than the other clients connected to BS_A. In some current LTE deployments, handovers are only initiated when a client observes a higher received power from a neighboring base station compared to its serving base station. For the distributed case, no handovers will be made and the edge users will continue to receive poor service.

If we now apply SoftRAN, the controller, with the knowledge of the entire network state, will easily recognize that the edge users will receive immense benefit from being handed over to BS_B. Although the signal strengths they observe from BS_B will be slightly inferior to those observed from BS_A, BS_B will be able to serve them with far more resource blocks. In performing these handovers, more resources (could be both transmit power or resource blocks) are simultaneously opened up to the other clients at BS_A. Such load balancing in a more complicated scenario requires extensive communication between the base stations, whereas a central controller can simply determine the allocation that balances the load and issue instructions to the base stations accordingly.

3.2 Utility Optimization

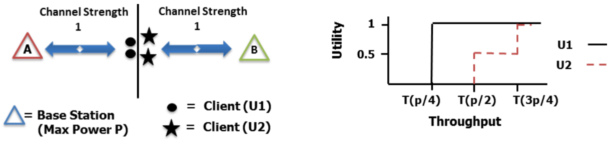

(a) Network Topology (b) Utility Functions

Figure 3: Utility Optimization: Setup

In this example, we examine quality of service (QoS) requirements in LTE; our objective is to maximize the global utility of the flows in the network. Let us assume the scenario shown in Fig.3(a), in which two base stations A and B (BS_A and BS_B) each have 2 clients at their respective cell edges. We also assume that the RAN has 4 resource blocks ($RB_{1...4}$) available. BS_A initially serves two clients with utility function U_1, while BS_B initially serves two clients with utility function U_2. Each base station also has a max transmit power P. Additionally, all four clients observe unit channel strength to both base stations and a noise power of 1. U_1 is a step function, where the step occurs at a fairly low throughput; a VoIP flow for example might have such a utility function. U_2 sees a step increase in utility at two throughput values; a video stream that achieves a quality of 480p for a given throughput and 720p for a larger throughput could have such a utility function. $U1$ and $U2$ are plotted in Fig.3(b) for reference. Note that $T(S)$ in the figure, signifies the throughput for a given SINR, S.

For the distributed case, we will assume that each base station allocates its resources to maximize the utility of its clients. For this scenario, lets also assume that the two base stations have settled on a resource allocation which ensures that all of the clients observe channels without interference; that is, all clients are allocated distinct resource blocks. Table 1 depicts the power allocation on each of the resource blocks, which results in a total utility of 3 for the distributed case.

If we now apply SoftRAN to the same scenario, we are able to better maximize the total utility. With a global view of the network, the central controller would note that the VoIP clients are using more resources than they need to achieve their maximum utility. In light of this, the controller would issue an instruction to BS_A to hand one of the VoIP clients over to BS_B and an instruction to BS_B to hand one of the video streaming clients over to BS_A. It would also specify the power allocation on the resource blocks as shown in Table 2 for the two base stations. Given this allocation, the network achieves a total utility of 4. From a qualitative perspective, we see that utilizing a central controller allows the network to instinctively minimize wasted resources, thus freeing up resources for more demanding flows. Specifically in this example we imagine two clients being able to stream videos in higher quality without compromising the call quality of nearby VoIP clients. Ultimately, we believe that having such an architecture that is able to optimize over arbitrary utility functions will be invaluable in providing better quality of service to clients.

Table 1: Distributed Resource Allocation

Base Station A				Base Station B			
RB1	RB2	RB3	RB4	RB1	RB2	RB3	RB4
U1	U1					U2	U2
P/2	P/2					P/2	P/2

Table 2: Centralized Resource Allocation

Base Station A				Base Station B			
RB1	RB2	RB3	RB4	RB1	RB2	RB3	RB4
U1	U2					U1	U2
P/4	3P/4					P/4	3P/4

4. FEASIBILITY

We wish to now do a feasibility analysis of the proposed design. For this purpose, we lay out the API between the controller and the radio elements in detail and then analyze the backhaul bandwidth required to support N radio elements on the controller with C clients per radio element and F flows per client. Also, we assume that R resource blocks need to be scheduled by the RAN at each periodic run of the controller.

4.1 Radio Element to Controller Updates

The main components of the update API from the radio element to the controller include:

Interference Channels: Each client observes interference powers from all neighboring radio elements and reports them to its serving radio element. The serving radio element then forwards these measurements to the controller, where they are used to form an interference map of the network.

Given $N*C$ clients, the total bandwidth used for reporting interference powers scales as $N*(N*C)$ (each client sees roughly N interferers)

Flow Records: Radio elements also update the controller about the flow state of all ongoing flows.

Given $N*C*F$ flows, the bandwidth used for reporting flow records scales as $N*C*F$.

4.2 Controller to Radio Element Instructions

The main component of the instruction message includes:

Neighbor Transmit Powers: Each radio element is notified of the transmit powers being used by all its neighboring radio elements. This information allows each radio element to alter its downlink resource allocation if necessary.

The total bandwidth used for reporting neighboring transmit power information scales as $N^2 * R$.

4.3 Required Backhaul Bandwidth

For a backhaul network of capacity 500 Mbps, we can support up to 20 micro cells(N = 20) with 50 clients per micro cell (C = 50) and 10 flows per client (F = 10). Since, microcells are typically a kilometer in radius, 20 such microcells would cover a geographical area with a radius of a few tens of kilometers. Thus, with an architecture like SoftRAN, a controller could serve a similar area as a macro cell with significantly cheaper hardware (low power radio elements and a server for control decision making). Note that this estimate assumes that all updates are sent at every opportunity instead of only updating those values which have changed since the previous update.

Moreover, the actual number of microcells that can be supported by the backhaul network varies with backhaul link capacities and backhaul topologies. Our goal in creating SoftRAN is not to centralize the entire RAN. Instead we hope that the big-base station abstraction allows dense deployments to be treated as sparse deployments, thus alleviating the demands of network management.

5. DISCUSSION

SoftRAN primarily provides the architecture for coordinated radio resource management through its logically centralized control plane. Given a global view of the network, a big-base station can manage interference, load, quality of service constraints, and perform other optimizations through plug and play algorithms which simplify network management. Additionally, by performing more effective handovers, SoftRAN aims to provide a smoother user experience. Though these are the primary benefits of SoftRAN, the applications do not end there.

Coordination of L1 functions: A SoftRAN radio element may have additional capabilities such as soft decoding or beamforming. For example, in the case of beamforming, the controller can choose to notify the group of base stations that need to cooperate in beamforming to a particular client. Moreover, it can also specify the set of resource blocks to use. Other physical layer techniques like soft decoding and successive interference cancellation can also be coordinated by the controller in a similar manner.

To be able to manage such complex physical layer techniques across multiple radio elements, the SoftRAN controller can also instruct the radio element on its data plane actions. These data plane instructions can be stored at the radio element as a table of rule-action pairs which can be as diverse as:

- Match a rule on a Layer 1 header and perform a corresponding Layer 3 action. For example, for soft decoding, whenever a packet from a particular client is intercepted (Layer 1 header), perform a soft decoding and route the information (Layer 3 action) to its corresponding serving base station.

- Match a rule on a Layer 3 header and perform a corresponding Layer 1 action. For example, whenever a packet routed from a neighboring base station is received (Layer 3 header), use the information in the packet to improve the physical layer decoding (Layer 1 action).

Dynamically adapting logical RAN architecture to traffic patterns: SoftRAN's architecture also allows us to envision new possibilities in the RAN. For example, we could redefine the boundaries of a big-base station (remap the association between radio elements and big-base stations) as needed on a much longer time scale. Using this flexibility, for example, we could minimize the number of handovers occurring across big base stations and hence maximize the benefits of the logically centralized control plane.

In a similar vein, we can imagine being able to turn off radio elements dynamically so as to conserve power if the network is under-loaded at a given time.

Leveraging open FemtoAPI for incremental deployment: We have designed SoftRAN in such a way that requires minimal changes to both existing base stations and the LTE standard itself. In other words, we have taken great care to ensure that our architecture is incrementally deployable with the current infrastructure and standards. FemtoAPI [6] is an abstraction being proposed by the industry to standardize the interface between the layer 1 functionality and the scheduler. Its objective is to encourage innovation and competition between the platform hardware, platform software, and application software by providing a common API to work around. SoftRAN fits perfectly in this scheme of things and such standardized API would simplify SoftRAN's implementation.

6. RELATED WORK

Our work is closely related to 3GPP Self Organizing Networks [1], cloud RAN [3] and software-defined or programmable wireless networks [8, 2, 7].

3GPP has recognized the need for better coordination and easier management of radio access networks. SON (Self Organizing Networks) was proposed in Release 8 [1]. The goal was to make the network capable of self-configuration and self-optimization. However, while the aspect of self-configuration has been well studied and their solutions have been well documented, the aspect of self-optimization has been left unexplored. We believe that SoftRAN is a step forward in realizing the goal of self-optimization.

CloudIQ [3] centralizes all data and control plane processing. However, we believe that pushing all data plane processing to a central entity imposes huge demands on the bandwidth and latency required on the backhaul. The central entity (controller) in SoftRAN is not responsible for any data plane functionality and only requires control information, thus drastically reducing the demand on backhaul bandwidth. Moreover, to handle the latency introduced by the backhaul, SoftRAN refactors control plane functionality in such a manner that latency-sensitive decision making continues to be handled by the base station.

OpenRoads [8] is the first software-defined wireless network. It is mainly based on WiFi and offers no special support for cellular networks. CellSDN [7], similar to SDN for wired networks [5], attempts to centralize the control plane for cellular core networks. In contrast, we restructure the RAN architecture to help cellular networks better manage their scarce radio resources. OpenRadio [2] proposes a novel design for a programmable wireless data plane that provides modular and declarative programming interfaces across the entire wireless stack. OpenRadio does not provide any software defined RAN controller.

7. CONCLUSION AND FUTURE WORK

Current radio access networks utilize distributed protocols to enact handovers and manage interference. While these protocols perform well enough in sparse deployments, they will be unable to effectively handle rapidly growing mobile traffic and the densification of base station deployments. To account for this changing direction of the mobile space, we have proposed a software-defined centralized control plane for radio access networks that abstracts all base stations in a local geographical area as a virtual big-base station comprised of a central controller and radio elements. Our analysis shows that such a system is not only feasible, but opens the door for further innovation and simplified network management. Through our abstraction and architecture, we have sought to create an environment which enables efficient and dynamic management of increasingly scarce and strained radio resources.

We have implemented SoftRAN in LTE-Sim. For future work, we would like to extensively evaluate the performance and scalability of SoftRAN in both the hardware and software domains.

8. ACKNOWLEDGEMENT

Aditya Gudipati's summer internship at Bell Labs in 2012 was supported by NSF grant CNS 1018292.

9. REFERENCES

[1] 3gpp. Self-organizing networks (SON) policy network resource model (NRM) integration reference point (IRP). `http://www.3gpp.org/ftp/Specs/archive/32_series/32.521/`.

[2] M. Bansal, J. Mehlman, S. Katti, and P. Levis. Openradio: a programmable wireless dataplane. HotSDN '12, pages 109–114, New York, NY, USA, 2012. ACM.

[3] S. Bhaumik, S. P. Chandrabose, M. K. Jataprolu, G. Kumar, A. Muralidhar, P. Polakos, V. Srinivasan, and T. Woo. Cloudiq: a framework for processing base stations in a data center. Mobicom '12, pages 125–136, New York, NY, USA, 2012. ACM.

[4] J. Bounds. Small cells to play a big role in at&t build out. `http://www.bizjournals.com/dallas/blog/2012/11/small-cells-to-play-big-role-in-att.html`.

[5] M. Casado, M. J. Freedman, J. Pettit, J. Luo, N. McKeown, and S. Shenker. Ethane: taking control of the enterprise. SIGCOMM '07, pages 1–12, New York, NY, USA, 2007. ACM.

[6] S. C. Forum. Femto application platform interface. `http://www.smallcellforum.org/resources-technical-papers`.

[7] L. Li, Z. Mao, and J. Rexford. Toward software-defined cellular networks. In *European Workshop on Software Defined Networking (EWSDN)*, Oct. 2012.

[8] K.-K. Yap, R. Sherwood, M. Kobayashi, T.-Y. Huang, M. Chan, N. Handigol, N. McKeown, and G. Parulkar. Blueprint for introducing innovation into wireless mobile networks. In *VISA '10: Proceedings of the second ACM SIGCOMM workshop on Virtualized infrastructure systems and architectures*, pages 25–32, New York, NY, USA, 2010. ACM.