VNTS: A Virtual Network Traffic Shaper For Air Time Fairness In $802.16e$ Systems

Gautam Bhanage[†], Ronak Daya[†], Ivan Seskar[†], Dipankar Raychaudhuri[†]
[†]WINLAB, Rutgers University, RT 1 South, Technology Center, NJ 08902, USA
Email: {gautamb, ronakd, seskar, ray}@winlab.rutgers.edu

***Abstract*—The $802.16e$ standard for broadband wireless access mandates the presence of QoS classes, but does not specify guidelines for the scheduler implementation or mechanisms to ensure air time fairness. Our study demonstrates the feasibility of controlling downlink airtime fairness for slices while running above a proprietary WiMAX basestation (BS) scheduler. We design and implement a virtualized infrastructure that allows users to obtain at least an allocated percentage of BS resources in the presence of saturation and link degradation. Using Kernel virtual machines for creating slices and *Click* modular router for implementing the virtual network traffic shaping engine we show that it is possible to adaptively control slice usage for downlink traffic on a Profile *A* WiMAX Basestation. The fairness index and coupling coefficient show an improvement of up to 42%, and 73% with preliminary indoor walking mobility experiments. Outdoor vehicular measurements show an improvement of up to 27%, and 70% with the fairness index and coupling coefficient respectively.**

Index Terms*—WiMAX testbed, 802.16e, VNTS, Group fairness.

I. INTRODUCTION

One of the general trends in wireless networking research is the growing use of experimental testbeds for realistic protocol evaluation. Open networking testbeds such as ORBIT [11], Dieselnet [13] and Kansei [5] have been widely used in the past 5 years for evaluation of new wireless/mobile architectures and protocols based on available radio technologies. With the emergence of so-called "4G" networks, there is a need to support open experimentation with wide-area cellular radios such as mobile WiMAX or LTE. A recent initiative by GENI [2] is aimed at making an open WiMAX base station available to GENI and ORBIT outdoor testbed users. As a part of this initiative, we address the design challenges of integrating a WiMAX basestation as a part of a virtualized wireless testbed.

Virtualization of the WiMAX Basestation (BS) provides a convenient approach to provide separate environments for different slices. A slice refers to the subset of BS resources and WiMAX clients allocated to a user/experimenter. The problem of ensuring radio fairness and policies across slices is specially hard since the channel for mobile wireless devices changes continuously, thereby consuming varying amount of resources at the BS transmitter.

This study describes challenges in ensuring air time fairness when access to the 802.16e scheduler is unavailable. Evaluation is provided for a carrier grade BS deployed at the Rutgers university. The WiMAX clients are accessible to experimenters through virtual machines (VMs) running on the ORBIT network [11]. The entire experimenter space from the VMs to its corresponding clients refers to a single slice. Implementation and evaluation of our proposed $VNTS$ architecture is provided under both walking and vehicular mobility.

Rest of the paper is organized as follows. Section II provides a brief survey of related work. Section III provides pilot measurements that help justify the use of our $VNTS$ mechanism for administering fairness policies across slices. Section IV discusses details and design considerations for implementation of the VNTS architecture. Section V provides a thorough evaluation of the system under varying network conditions. Finally, Section VI gives concluding remarks.

II. RELATED WORK

WiMAX being a relatively new technology, most of the recent work in this area has been focussed on theoretical modeling. There have been significant efforts with the development of models to modify the scheduler for quality of service [12], [8], [14], [9] guarantees and or for ensuring fairness [4], [6]. However, in our problem we take a practical approach that QoS is provided as per pre-set classes by the built in proprietary scheduler. By treating the scheduler as a blackbox device, we provide an architecture that provides air time fairness across slices.

A study in slice control [10] has addressed issues of space versus time multiplexing of 802.11 links for interference minimization. However, since we operate in $802.16e$, links with a single BS are interference free. A token-passing based air time fairness mechanism was implemented in [15]. Though this approach is suitable for implementation on 802.11 devices, it does not demonstrate performance with varying traffic loads, and frame sizes. To the best of our knowledge we are not aware of any study which implements an air time fairness mechanism for user groups in $802.16e$ devices at this time.

III. MOTIVATION

We begin with a brief description of the experiment apparatus followed by a pilot experiment to motivate the study.

A. Hardware Setup

The deployed BS is capable of up to 8 modulation and coding schemes (MCS) and implements two mechanisms for downlink (DL) rate adaptation. DL-link adaptation is active

$Parameter$	$Value$
$Channel\ \ Rate$	Adaptive
Frequency	2.59GHz
DL/UL Ratio	35 : 12
Bandwidth	10Mhz
Client Chipset	Beceem

Fig. 1. Basestation (BS) settings for all experiments. Explicit change in parameters are as mentioned in the experiments.

and changes rate every 200frames. Downlink-Uplink symbol ratio is variable between $(35 : 12) - (26 : 21)$. A Profile-A BS allows the creation of pre-provisioned Best effort service flows which are used in all experiments. Other BS settings unless mentioned otherwise are as shown in Figure 1.

The system architecture integrated in the ORBIT framework is as shown in the Figure 2. The Application Service Network gateway $(ASN - GW)$ is used for connecting the basestation to the outside world. The internal control interfaces of the basestation (network and RF) as well as the ASN-GW are networked as a part of the *instrument* network. The external interface of the ASN-GW is a part of the *outdoor* network. Client connectivity is provided by PCMCIA cards. Due to limited hardware we have to limit one client per slice. However, this can be easily extended to multiple clients per slice using our setup. Experimenters can include clients as a part of their slices through virtual machine instances [3] running on the *VM Host* machine. Access through the VM host machine also allows administrators to control slice access to pre-provisioned service flows, thereby determining slice QoS. It is important to note that QoS only provides traffic prioritization, while our mechanism aims at providing fairness.

B. Baseline Experiment - Femtocell Mobility

Our baseline experiment emulates mobility in a Femtocell deployment, and is repeated in all indoor measurements. As shown in Figure 2, the virtual machines VM1 and VM2 are sending traffic to a stationary and and mobile client respectively. The link from each VM to the corresponding client constitutes of a slice. The stationary client is located such that it has a CINR greater than 30 which allows the basestation to send traffic comfortably at $64QAM\frac{5}{6}$. An experimenter walks with the mobile client as per the coverage map shown in the Figure 3. As per the RSSI trace for the walk, the link degrades in a corner of the corridor and improves as the experimenter returns to the starting position. Each VM is configured to saturate the link to its client with UDP traffic. Other parameters as shown in Figure 1. The observed downlink throughput for both the clients is as shown in Figure 4(a). We observe that as the mobile client reaches areas where the RSSI drops below certain threshold, the rate adaptation scheme at the basestation selects a more robust modulation and coding scheme(MCS). However, in the process the link with the mobile client ends up consuming a lot more radio resource at the basestation, which affects performance of the stationary client. Thus we observe that while the BS scheduler is capable of providing QoS, it does not ensure radio resource fairness across links.

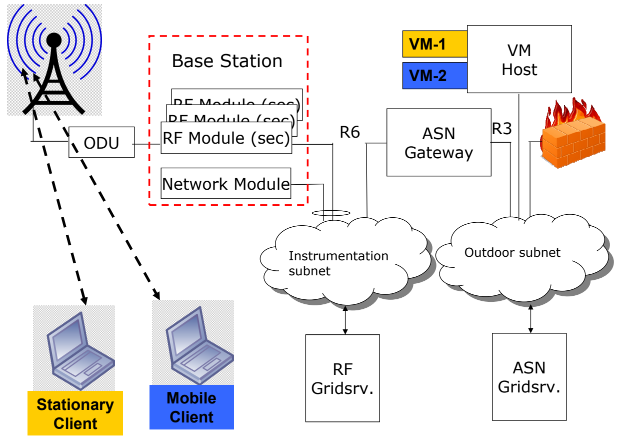

Fig. 2. Integration of the WiMAX setup into the ORBIT testbed.

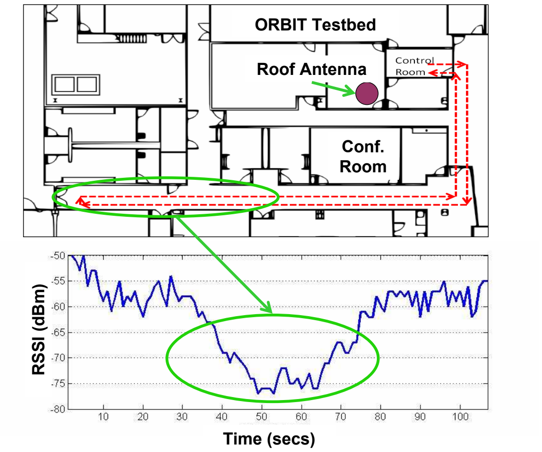

Fig. 3. Walking path used by the mobile client for indoor experiments *(Topo-1)*. Typical, variations observed in RSSI are as shown.

For simplicity, we will consider the initial problem of assigning 50% of BS resources to each of the two clients. A conservative solution to this problem is through static shaping. Assuming that we have some way of knowing about the movement of the mobile client, we could calculate throughput based on 50% channel time at the slowest MCS that would be used by the BS to reach the mobile client. In this case, the basestation uses $\frac{1}{2}QPSK$ when link quality deteriorates the most. Hence we use this information to statically shape the throughput of the slice with the mobile client. The goal of this shaping mechanism is to limit the offered load for the slice in such a way, that the slice can use only the allocated share of basestation resources. The results of an experiment that demonstrates performance with such a static shaping policy is as shown in Figure 4(b). The measurements show that this time the throughput of the slice with the stationary node becomes independent as a virtue of the shaping for the mobile client. Pro-active traffic shaping results in lower throughput for the mobile client, which helps prevent the

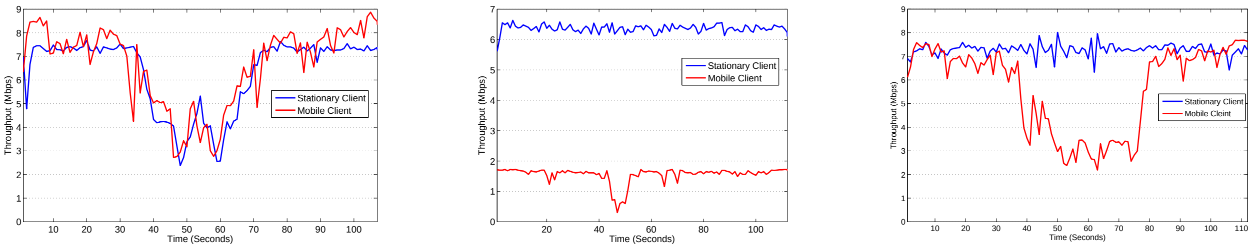

(a) No Shaping
(b) Static Shaping
(c) Adaptive Shaping (VNTS)

Fig. 4. Performance improvement using VNTS for walking mobility shown in Figure 3 (*Topo-1*).

saturation of the basestation and eliminates coupling between slice performances. However, we also observe that while fairness increases, aggregate downlink throughput decreases significantly even when channel to the mobile client is good. Static shaping will also require the knowledge of mobility and link saturation in advance.

Hence, no shaping gives better overall channel utilization with no slice fairness, and static shaping gives us better fairness with significantly lesser utilization. To alleviate this problem we propose and implement the virtual network traffic shaping (VNTS) technique, which adaptively controls slice throughput. Results from the experiment repeated with *VNTS* are as shown in the Figure 4(c). Even as the channel for the mobile client deteriorates, the VNTS mechanism is able to appropriately limit the basestation utilization for the mobile client (slice) thereby providing fairness to the stationary client. The next section will provide details on the design and working of the VNTS architecture.

IV. VNTS ARCHITECTURE

The VNTS architecture is as shown in the Figure 5 and consist primarily of the VNTS engine and the VNTS controller. The VNTS engine is responsible for performing the traffic shaping from the virtual machines, and the VNTS controller component is responsible for controlling the VNTS engine by determining slice throughput and current operating conditions.

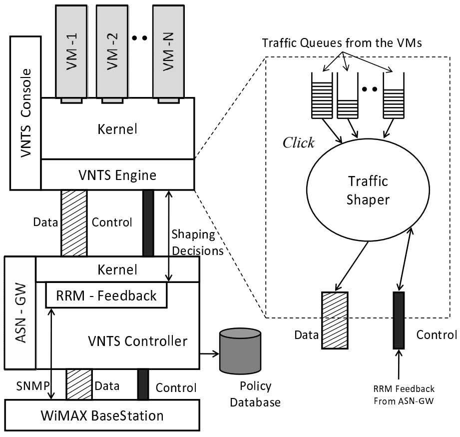

Fig. 5. Overview of version 1 architecture for virtualization with the Wimax Basestation

```
Input: MCS Rate $A_i$, Slice Weight $W_k$,
Number of clients per slice $N_k$
Output: Shaping rate per client $\gamma_i$
while $True$ do
    foreach Slice $k$ do
        $N_k = getNumClientsForSlice(k)$
        foreach Client $i \in$ Slice $k$ do
            $SNMPGet(A_i)$
            $\gamma_i = A_i \times \frac{W_i}{N_k}$
            Set($\gamma_i$)
        end
    end
    $sleep(UpdateInterval)$
end
```

Algorithm 1: Algorithm for adaptive virtual network traffic shaping (VNTS) controller.

A. *VNTS Engine*

The virtual network shaping engine is implemented in the Click Modular router [1] run in user mode as shown in Figure 5. This mechanism is responsible for-

1) Separating VM traffic based on a slice identifiers. We are currently using MAC identifiers of virtual machine interfaces as the slice identifiers.
2) Arping, routing and shaping as per policies to and from the virtual machines.
3) Providing element handlers that allow dynamic control of virtual machine traffic by using the VNTS controller.

B. *VNTS Controller*

Resource blocks at the BS are defined as a set of time - frequency tiles in the OFDMA radio that are allocated to individual links by the BS scheduler. Isolation between slices is compromised when the WiMAX basestation runs out of resource blocks and the slice with a poor link, eats into the resources of the slice with a better link. Since allocation of these resource blocks is only accessible to the BS scheduler, we aim to alleviate this problem by input to the scheduler.

The controller is based on our observation that the number of resource blocks used per slice at the WiMAX BS are directly proportional to the offered load per slice in terms of

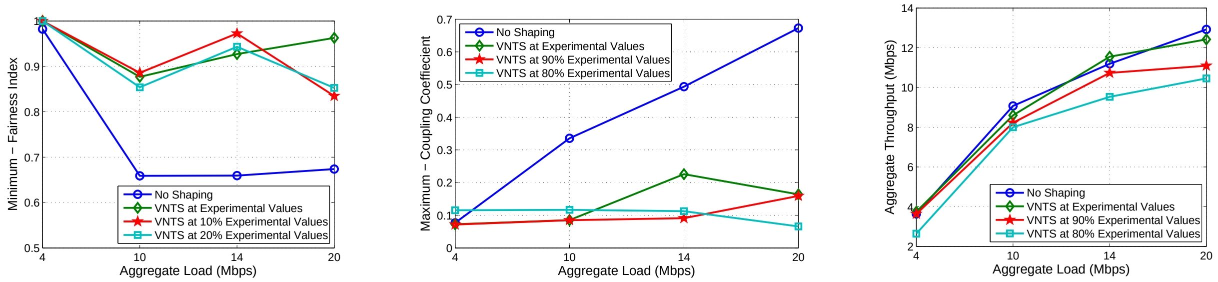

(a) Fairness Index (max) (b) Coupling Coefficient (min) (c) Channel utilization

Fig. 6. Performance for indoor walking experiments *(Topo-1)* with various shaping policies.

channel time required per second. Hence the end goal of our VNTS controller is to detect saturation, and limit the offered load (in Mbps) per slice in such a way that channel time required per second for every slice scales in proportion to the weight assigned to that slice.

The complete algorithm is as shown in Algorithm 1. The rate at which a client is receiving downlink data is determined by sending $SNMP$-GET queries to the basestation. The current rate adaptation algorithms on the WiMAX BS change MCS every 200 frames. Hence, there is a good chance that the MCS received from the SNMP query varies from the MCS used for the batch of frames for which shaping is done. Averaging the MCS bit rate obtained for every 200 frames would provide a solution to this problem, but this would demand a large amount of control overhead and make it infeasible with current latencies for SNMP call completion. Instead, we shape the traffic every 1sec with a conservative estimate of the MCS based on the BS feedback. The control algorithm scales the saturation throughput for slices in proportion to their weights. The $Set()$ function is used to remotely set throughput limit for each virtual machine. $UpdateInterval$ is a parameter that determines how often the control loop is run. Both of these parameters are controllable. By default, this loop is executed every second, to achieve repeated control. The sleep duration for the loop can eventually be made adaptive based on observed link conditions.

V. EVALUATION

In this section we present results from experimental evaluation of the VNTS architecture. We begin with a brief description of the metrics used for slice fairness followed by results from indoor and outdoor experiments.

A. Metrics

In our evaluations, we modify and use the Jain fairness index [7] for determining weighted fairness across flows for varying levels of offered loads. Let the throughput observed at a client i be given by T_i, bit rate achieved with the current modulation and coding scheme for the slice i be given by A_i. Number of clients (flows) belonging to a slice k be given by n_k. Total number of such concurrent slices is given by N. Fraction of channel time used by all links in slice k can be calculated as ϕ_k-

$$\forall\ Clients_i \in Slice_k,\quad \phi_k = \sum_{i=1}^{n_k} \frac{T_i}{A_i} \tag{1}$$

$$I = \frac{\left(\sum_{k=1}^{N} \phi_k\right)^2}{N \times \sum_{k=1}^{N} \phi_k^2} \tag{2}$$

The fairness index (I) determines the global variation in channel utilization across slices. We further modify the index to evaluate fairness under saturation with different slice weights while also accounting for performance deterioration due to bad channel quality. To measure worst - case performance, our experiments will determine minimum value of this index for different scenarios. Due to the availability of only two clients at this time, each slice constitutes of a single client.

Another metric used in our measurement is defined as the coupling coefficient. The coupling coefficient is used to measure the performance impact of the mobile slice on throughput obtained by a stationary slice. The coefficient C_k for every client k is measured as-

$$C_k = \frac{(T_k - T_k^{fix})}{T_k} \tag{3}$$

T_k^{fix} is the average throughput measured at client k, when all clients are stationary and in similar channel conditions. T_k denotes the average throughput of the stationary client k over one second, with the other client being mobile. The smaller the coupling coefficient, lesser the coupling between client k and other clients. In our measurements we focus on improving the worst case performance. Hence we will plot the maximum value of the coupling index seen over the duration of the experiments in all measurements.

B. Policy Conservativeness

Limiting slice throughput to more conservative values will enable better performance isolation between slices. However, this will also lead to a lesser net utilization of the available resources at the BS. To determine optimum rates for shaping, we perform more experiments with walking mobility using the same layout from Figure 3.

We measure performances by shaping at 80% (most conservative), 90% and 100% of the prescribed channel rates for

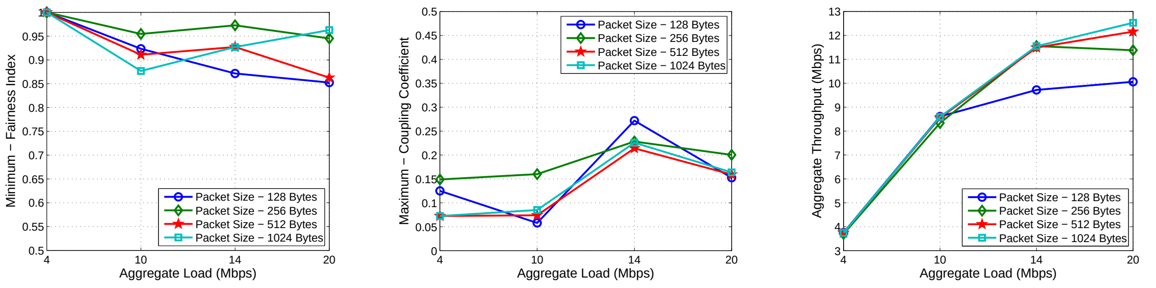

(a) Fairness Index (I) (b) Coupling Coefficient (C_k) (c) Channel utilization

Fig. 7. Performance with varying frame size. Frame sizes are varied for both flows.

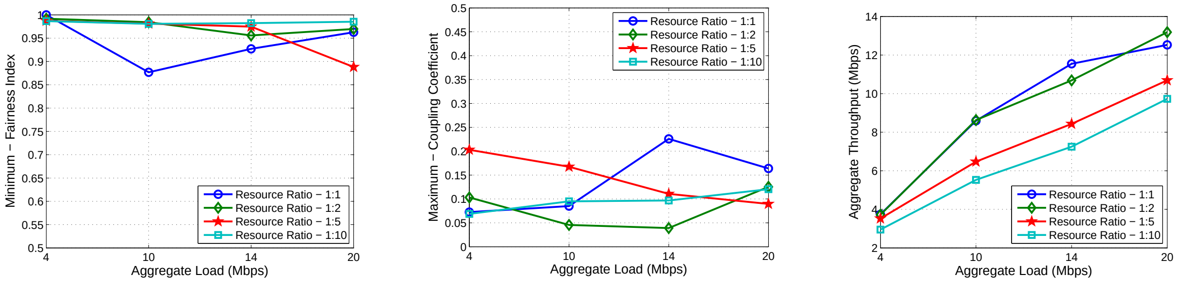

(a) Fairness Index (I) (b) Coupling Coefficient (C_k) (c) Channel utilization

Fig. 8. Performance with ratio of flow weights across slices. Flow weight for the mobile client is fixed while that for the stationary client is increased.

each of the modulation and coding schemes. We repeat the same baseline experiment as in the previous section, by having a stationary client and a mobile client moving as shown in Figure 3. Figure 6(a) plots the minimum value of the fairness index from equation (2) for each slice based on the radio resources promised to that slice as a function of aggregate load on the system. Without shaping, the fairness index hits lows of up to 0.67. However, with adaptive shaping at 80%, 90% and 100% of saturation values, we see the index ranging from $0.82 - 1$ while being independent of the system load. We also see an improvement of up to 42% in certain cases. Performance of the coupling coefficient for stationary client is as shown in Figure 6(b). We observe that without the VNTS mechanism coupling between slices reaches up to 0.66 under saturation, which can be limited to around 0.22 by the use of the VNTS mechanism. It is also seen that irrespective of the conservativeness of the shaping scheme, similar reduction in coupling is observed. Since we do not see a significant difference in fairness or coupling by varying conservativeness, we use channel utilization as a metric for deciding the best shaping rates.

Figure 6(c) shows the total channel utilization of the slices as a function of aggregate offered load. These results are as per intuition and show that the mean channel utilization across all system loads is the best when we use the prescribed channel rates without making the policy conservative. Since the throughput drop due to the use of a conservative policy is much more than the benefits in fairness, we chose prescribed channel rates as the default shaping values.

C. Varying Frame Sizes

As seen in previous experiments, slice fairness improves significantly after using our VNTS mechanism. We will now evaluate the performance of our system in the presence of varying frame sizes for slices. Experiment setup is the same as described earlier. Frame sizes used for the UDP traffic are varied as $128, 256, 512$ and 1024Bytes.

Figure 7(a) shows the minimum value of the fairness index with the use of different frame sizes as a function of aggregate offered load across clients. We observe that the fairness index varies between $0.85 - 1$. This shows that the fairness obtained across slices is independent of the frame sizes used by the slice. In all experiments, worst case fairness index or coupling performances are similar for loads $> 10Mbps$ because, as seen in Figure 4(a), the aggregate link capacity with mobility drops below $10M$ causing the BS to run out of resources despite low offered loads.

Results in Figure 7(b) show that despite varying the frame size for slices, we see that the coupling coefficient for the stationary client reaches a maximum value of 0.27 which is significantly lesser than that achieved without the VNTS mechanism. Finally, results in Figure 7(c) show that the VNTS mechanism works fairly across the clients and the aggregate throughput is similar to that achieved without shaping.

D. Varying Flow Weights

To study weighted fairness performance, we vary the weights of downlink resources given to the mobile and stationary slice in the following proportion- $1 : 1$, $1 : 2$ $1 : 5$, and

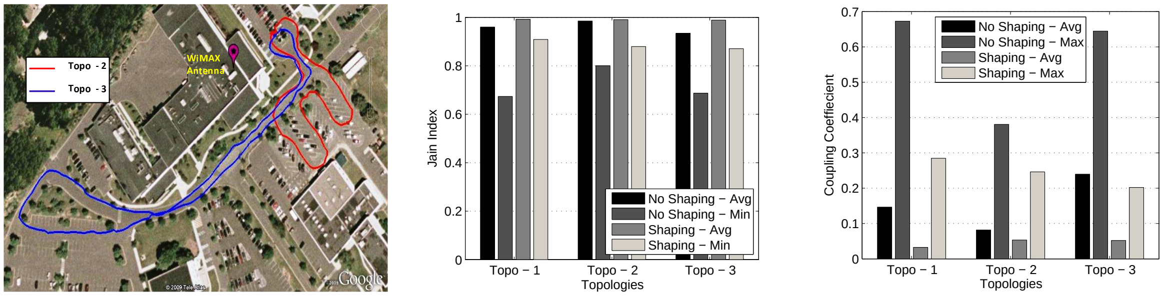

(a) Experimental topology (b) Fairness Index (I) (c) Coupling Coefficient (C_k)

Fig. 9. Relative performance of indoor and Vehicular experiments.

$1:10$. The results are plotted for different values of aggregate load on the system and are as shown in the Figure 8(a), 8(b) and 8(c). Each client loads the system equally during these tests. Across all loads we observe that the minimum value of fairness index in Figure 8(a) is maintained under acceptable limits. Typically as the weight for the stationary client is increased, fairness improves since we expect lesser impact of unfairness from the mobile slice. The Figure 8(b) shows that maximum coupling for the stationary client is also kept within limits (< 0.25) under all experiment conditions. Channel utilization is only limited by the offered load, and weights assigned to each slice.

E. Vehicular Measurements

To further validate the working of our VNTS system, we perform experiments with the mobile client placed in a vehicle. Experiment setup is same as that in the previous sections, with the exception that performance is determined only under saturation conditions, where each slice has a downlink offered load of 10Mbps. Outdoor results are measured for two topologies (*Topo-2,3*) as shown in Figure 9(a). Average velocity of the vehicle is maintained greater than $10Mph$ for both *Topo-2,3*. These results are compared with those obtained from the indoor experiment *Topo-1* as shown in Figure 3.

Figure 9(b) shows the modified fairness index for the three experiment setups. We observe that in all the cases, we obtain better fairness using our VNTS mechanism. Average values of the fairness index show improvements across all the topologies, and the worst case performance (minimum fairness points) is significantly improved with the use of VNTS. Figure 9(c) plots the performance of the coupling factor for the stationary client across all topologies. In the average case, we observe that the coupling reduces for all topologies with the use of our VNTS mechanism. The VNTS mechanism also improves the worst case performance for all topologies with reduction in coupling of up to 70% in some cases. Thus, using our VNTS mechanism, we are able to significantly improve fairness across slices even under vehicular mobility.

VI. CONCLUSIONS AND FUTURE WORK

This study motivates the need for an airtime fairness mechanism to be built with a virtualized WiMAX basestation. Design and prototyping results are presented to demonstrate the feasibility of such an approach outside the $802.16e$ scheduler thereby making the implementation platform independent. Preliminary evaluation of the VNTS architecture shows that it is possible to assign a certain percentage of the BS resources to experimenters or slice users while assuring significant isolation even under bad channel conditions. Future work involves more extensive evaluation, and using more client feedback to improve performance.

REFERENCES

[1] Click modular router. *http://read.cs.ucla.edu/click/*.

[2] GENI. *http://www.geni.net/*.

[3] Linux kernel virtual machines. *http://www.linux-kvm.org/*.

[4] S.-T. Cheng, B.-F. Chen, and C.-L. Chou. Fairness-based scheduling algorithm for tdd mode ieee 802.16 broadband wireless access systems. In *In proceedings of IEEE APSCC*, pages 931–936, Dec. 2008.

[5] E. Ertin, A. Arora, R. Ramnath, M. Nesterenko, V. Naik, S. Bapat, V. Kulathumani, M. Sridharan, H. Zhang, and H. Cao. Kansei: a testbed for sensing at scale. In *In proceedings of IPSN*, pages 399–406, 2006.

[6] F. Hou, J. She, P.-H. Ho, and X. Shen. Performance analysis of weighted proportional fairness scheduling in ieee 802.16 networks. In *In proceedings of IEEE ICC*, pages 3452–3456, May 2008.

[7] R. K. Jain, D.-M. W. Chiu, and W. R. Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. Technical report, DEC, September 1984.

[8] J. Lakkakorpi, A. Sayenko, and J. Moilanen. Comparison of different scheduling algorithms for wimax base station: Deficit round-robin vs. proportional fair vs. weighted deficit round-robin. In *In proceedings of IEEE WCNC*, pages 1991–1996, April 2008.

[9] Y.-N. Lin, C.-W. Wu, Y.-D. Lin, and Y.-C. Lai. A latency and modulation aware bandwidth allocation algorithm for wimax base stations. In *In proceedings of IEEE WCNC*, pages 1408–1413, April 2008.

[10] R. Mahindra, G. Bhanage, G. Hadjichristofi, I. Seskar, D. Raychaudhuri, and Y. Zhang. Space versus time separation for wireless virtualization on an indoor grid. In *proceedings of NGI*, pages 215–222, April 2008.

[11] D. Raychaudhuri, M. Ott, and I. Secker. Orbit radio grid tested for evaluation of next-generation wireless network protocols. In *In proceedings of IEEE TRIDENTCOM*, pages 308–309, 2005.

[12] A. Sayenko, O. Alanen, J. Karhula, and T. Hämäläinen. Ensuring the qos requirements in 802.16 scheduling. In *In proceedings of ACM MSWiM*, pages 108–117, New York, NY, USA, 2006.

[13] H. Soroush, N. Banerjee, A. Balasubramanian, M. D. Corner, B. N. Levine, and B. Lynn. DOME: A Diverse Outdoor Mobile Testbed. In *In proceedings of HotPlanet*, June 2009.

[14] J. Sun, Y. Yao, and H. Zhu. Quality of service scheduling for 802.16 broadband wireless access systems. In *In proceedings of IEEE VTC*, volume 3, pages 1221–1225, May 2006.

[15] G. Tan and J. Guttag. Time-based fairness improves performance in multi-rate WLANs. In *Proceedings of the 2004 USENIX Technical Conference*. USENIX Association, June 2004.