Space Versus Time Separation For Wireless Virtualization On An Indoor Grid

R. Mahindra*†, G. D. Bhanage*, G. Hadjichristofi*‡, I. Seskar*, D. Raychaudhuri*, Y.Y. Zhang*
*WINLAB, Rutgers University, RT 1 South, Technology Center, NJ 08902, USA
†NEC Labs America, USA
‡University Of Cyprus, Nicosia, Cyprus
Email: {rajesh, gautamb, ghadjich, seskar, ray, yyzhang}@winlab.rutgers.edu

***Abstract*—The decreasing cost of wireless hardware and ever increasing number of wireless testbeds has led to a shift in the protocol evaluation paradigm from simulations towards emulation. In addition, with a large number of users demanding experimental resources and lack of space and time for deploying more hardware, fair resource sharing among independent co-existing experiments is important. We study the proposed approaches to wireless virtualization with a focus on schemes conserving wireless channels rather than nodes. Our detailed comparison reveals that while experiments sharing a channel by space separation achieve better efficiency than those relying on time separation of a channel, the isolation between experiments in both cases is comparable. We propose and implement a policy manager to alleviate the isolation problem and suggest scenarios in which either of the schemes would provide a suitable virtualization solution.**

I. Introduction

The GENI Project [1], supported by NSF, aims to provide a flexible, programmable, shared experimental infrastructure for investigation of future Internet protocols and software. GENI will consist of a global-scale wired network with programmable and virtualizable network elements (routers, switches, servers) along with several wireless access network deployments intended to support experimentation with mobile computing devices, embedded sensors, radio routers, etc. In [3], the authors discuss the importance of wireless virtualization in the integration of wired-wireless testbeds. This project is aimed at finding solutions to the virtualization of wireless network resources to provide capabilities for simultaneous support of multiple concurrent experiments ("slices") on the same set of radio devices.

The ORBIT [4] is a 400 node wireless testbed sponsored by NSF for indoor wireless experimentation. It is a multi-user radio grid that allows "sequential" short term access to the radio resources for repeatable experiments. Time scheduling is done so that users have exclusive access to the entire grid during their slot. Excessive usage leading to lack of available time slots in the light of the GENI initiative has further motivated efforts for ORBIT virtualization. Thus, virtualization in the ORBIT context refers to the ability of splitting the wireless testbed resources among multiple concurrent experiments with each experimenter controlling a "slice" of the radio grid.

Suitableness of a wireless virtualization scheme is decided by:

- Resource Constraints: Different virtualization schemes can help conserve different resources (number of nodes, available orthogonal channels, ability of the experiment control mechanism to handle parallel experiments).
- Efficiency: Sharing of resources by virtualization introduces additional overheads. For example, in case of a UML [5] based approach to virtualization excessive resource utilization may be seen in the form of context switching. The virtualization scheme should be efficient such that there is minimal management overhead, since it eventually decides the maximum number of parallel experiments.
- Inter-experiment interference: Virtualization of any resource almost always results in some form of compromised performance for co-existing experiments. While mapping virtualization to scientific experiments it is necessary to quantify any performance degradation associated with experiments.
- Experiment Repeatability: An important aspect with performing indoor controlled experiments is to ensure the repeatability of results. Improper resource sharing may result in unpredictable performance across multiple experiment runs.

A wide range of wireless virtualization schemes have been proposed [2]. We select and compare two approaches to wireless virtualization: space separation and time separation; based on their suitableness for deployment on the ORBIT testbed. Empirical evaluation of sample scenarios are used for comparison and deduction of overheads with wireless virtualization. Despite having minimal overhead in terms of CPU utilization in both approaches, we show the usefulness for a policy management mechanism that dynamically allocates channel resources for experiments.

The contributions of this paper are to:

1) Compare strengths and drawbacks of space and time based virtualization among other schemes and determine their suitableness for deployment on a type of wireless testbed.
2) Provide empirical measurements from a systematic setup and use them to determine efficiency of the virtualization schemes.
3) Propose metrics to compare interference between experiments and provide an implementation of a Click based policy manager. We show that this policy manager makes it possible to select a better efficiency virtualization scheme irrespective of the level of inter-experiment interference.

This research is based on devices that use the same MAC and physical layers. In addition, all devices in the virtualization schemes use MAC and PHY layers that are compatible with the 802.11 standard. Our study does not aim to provide a comprehensive virtualization solution across heterogeneous wireless devices and drivers, but can serve as a reference to show the trends in performance that may be observed with the use of the two aforementioned virtualization approaches. This study lays out the criteria, which could be used for deciding virtualization schemes on testbeds. We believe that our study sets the foundation for some of the key design issues and deployment strategies for wireless virtualization on large-scale network testbeds like GENI.

The rest of the paper is organized as follows. Section II describes some of the important approaches of wireless virtualization. Section III and IV compare the performance difference between SDMA and VAP-based virtualization schemes. Section V discusses inter-experiment effects that may be seen in channel multiplexed wireless virtualized schemes. Finally, in Section VI, we propose a policy manager for ensuring fairness between virtualized experiments.

II. Virtualization Schemes

Wireless virtualization approaches may be conveniently classified along the space, time, and frequency axes as:

- Frequency separation channel sharing
- Space separation channel sharing
- Time separation channel sharing

Before providing an overview of these approaches, we briefly describe our experimental testbed.

A. Virtualization Platform

ORBIT is a two-tier laboratory emulator/field trial network testbed designed to achieve reproducibility of experimentation, while also supporting evaluation of protocols and applications in real-world settings. The laboratory-based wireless network emulator uses a novel approach involving a large two-dimensional grid of 400 802.11 radio nodes, which can be dynamically interconnected into specified topologies with reproducible wireless channel models. The majority of the ORBIT nodes are fitted with Atheros 5212 based cards while the remaining have Intel cards. We used Atheros cards for our experiments.

B. Outline Of Virtualization Approaches

The outline of design principles of GENI in [2] presented the following virtualization techniques that are intended to share a set of wireless resources amongst multiple users.

Frequency Separation Channel Multiplexing: Frequency separation implies partitioning of the experiments in the frequency domain with different experiments assigned orthogonal channels to prevent interference. Multiple experiments are executed on the same physical nodes, with each experiment being executed in an instance of an OS or virtual OS. Thus, the resources of a physical node are split into multiple virtual nodes. This virtualization would introduce a finite channel switching delay when switching from one virtual node to another. In most facilities, there is a provision for multiple wireless interfaces. Hence, individual experiments could be mapped to different wireless interfaces on the same physical node eliminating switching delays.

Space Separation Channel Multiplexing: The space separation approach splits the testbed resources to provide sufficient spatial separation between the transmitting nodes and avoid interference across individual experiments. During this allocation, a subset of the physical resources is assigned to a specific experiment. This separation provides virtualization across multiple nodes eliminating the need for experimenters to share experiment nodes. Space separation will be broadly referred to as space division multiple access (SDMA) scheme.

Time Separation Channel Multiplexing: Time separation or Time division multiple access (TDMA) virtualization partitions the network in time domain across multiple experiments. Multiple experiments run on the same physical nodes, with each experiment sharing the wireless resources in time. Time sharing of the channel is discussed in further detail in subsection *E*.

C. Most Suited Approaches

Selection of a virtualization scheme primarily depends on the resource being conserved. Wireless virtualization can be targeted at either the conservation of nodes (hardware) or channels (frequency). Frequency multiplexing of the wireless channel, allows for node conservation where the same node could be shared using a UML like mechanism on multiple channels to emulate different experiments. Keeping in mind Moore's Law, the concern for deployment of a virtualized testbed would be more on channel conservation especially with the availability of relatively cheaper commodity wireless devices. For instance, with access to 800 wireless interfaces on the ORBIT grid the focus was more on channel conservation rather than node conservation. Frequency multiplexing may not scale well with provision of only three orthogonal channels in 802.11b mode of experimentation. Since time and space separation allow for channel conservation we will compare and contrast these approaches for selecting an apt virtualization setting.

D. Space Separation on ORBIT

The ORBIT wireless testbed is located in a 20 meter x 20 meter space and hence the nodes are in close physical proximity of one another. Under these conditions, partitioning the resources in space to avoid interference would not be practically possible. This limitation holds true for most of the emulator testbeds. Artificial stretching of the distance between the nodes is achieved by controlling transmission power of the nodes and using "noise injection" to emulate barriers between the nodes of different experiments. However, observation from previous studies [8] reveals that it is considerably difficult to create and limit the effect of noise locally with the current noise injection subsystem on ORBIT. Therefore, our experiments explore the possibility for virtualizing the

ORBIT grid using SDMA by controlling power in addition to providing spatial separation. In this artificially stretched SDMA scheme, the individual experiments are multiplexed on the same channel.

E. Time Separation On ORBIT

Time sharing on the ORBIT wireless grid can be achieved using two approaches:

- Explicit TDMA implementation
- Virtual access points (VAP)

TDMA: TDMA has been implemented and tested on the ORBIT grid in [6]. This approach runs multiple UML instances on the same node, which use the same wireless device and the same wireless channel. They ensure through tight synchronization, that at any time all the nodes are running the same experiment slice across the network of nodes.

Efficiency of implementation and overall performance seen with a TDMA scheme will greatly depend on:

- Experiment Synchronization: In TDMA, there is a stringent need for high degree of time synchronization between all the experiment nodes. Moreover, wireless experiments can potentially involve a high number of experimental nodes. Though tools like the network timing protocol (NTP) [13] can provide distributed time synchronization, accuracies achieved may not be sufficient.
- Design Dilemma: The choice of time slot allotted to the different experiments is another design issue for the TDMA approach. A small value may not be possible due to practical limitations of wireless hardware like switching time and a large value would adversely affect the performance in delay sensitive experiments. Since, in this approach, several concurrent experiments share one or more physical nodes, there is also a need to provide isolation on every node between the experiments.

The TDMA approach requires design and deployment of a complicated infrastructure on current testbeds like ORBIT, which does not seem plausible. To offset these disadvantages we consider the use of virtual access points as a mechanism for channel multiplexing.

Virtual Access Points (VAPs) A VAP is defined as a logical abstraction that runs on a physical access point while emulating the behavior of a conventional access point to all the stations in the network [7]. Using a VAP allows for two or more AP mechanisms to share the same channel thereby helping channel and energy conservation. In contrast to the TDMA approach for channel multiplexing, VAPs are more suitable for running short and long-term experiments with less stringent constraints on the current testbed resources.

The concept of VAPs is incorporated in the 802.11 driver, which operates just above the MAC layer and below the IP layer. The driver provides the multiple AP abstraction to the higher layers though it is operating on a single lower layer. Hence all the protocols operating on the machine are agnostic to the presence of the abstraction. As we will show in the coming sections this setup can be extremely

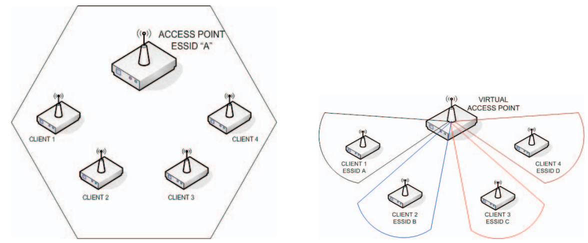

(a) A physical access point and four clients

(b) Four virtual access points and their individual clients

Fig. 1. Experimental setup for performance evaluation with physical and virtual access points.

useful for minimizing down link interference with multiple infrastructure mode setups. Compared to the TDMA approach, the VAP does not require tight synchronization among the different experiment nodes. However, this scheme requires traffic shaping and is limited to fixed star topology wireless networks.

Since channel conservation is of prime importance, we choose to evaluate the space and time separation approaches for virtualization on the ORBIT grid. As the VAP approach provides a more plausible solution to time multiplexing over conventional TDMA approaches intended for long term experimentation, we will use it for further quantitative evaluation with space separation.

III. Throughput Comparison

Throughput, latency and jitter are usually the three main parameters, which determine a users utilization and experience on a network device. Throughput for individual experiments in a virtualized environment is expected to be lesser than those under single user conditions. However, performance under these conditions is largely contingent on how fairly the resources are shared.

A virtualized channel is shared among multiple users running simultaneous experiments and the end performance can largely be a function of individual experiment parameters rather than just a fair share between users.

Prior to investigating and comparing VAP and SDMA based virtualization schemes, we discuss briefly the implementation and operation of a VAP, which is a relatively new concept. This study should be insightful in determining whether a VAP provides significant advantages over a conventional physical access point setup.

A. Virtual Access Point Overhead

A VAP creates an abstraction of multiple physical access points running from the same hardware for the clients associating with it. Creation of these logical entities requires state maintenance and independent management signaling for each of the networks managed by each VAP.

Before we evaluate the benefits of using VAPs, we consider it important to determine the overheads of maintaining the state of multiple networks at a single hardware device. The experimental setup for comparison is as shown in Figure 1(a) and Figure 1(b). Figure 1(a) shows a setup with one AP and

$Parameter$	$Value$
$Channel\ \ Rate$	36Mb/sec
$Aggregate\ \ Offered\ \ Load$	50Mb/sec
$Experiment\ \ Duration$	5 Minutes
$Averaging\ \ Duration$	Per Second
$Operation\ \ Mode$	802.11a
$Traffic\ \ type$	Uplink
$Chipset$	ATHEROS
$Driver$	Madwifi(0.9.3.1)

Fig. 2. Experimental Parameters Used With ORBIT Nodes

all four clients within the same network. Figure 1(b) has the same nodes. However, each of the clients now belongs to a different logical network created by the VAPs. Care was taken to ensure that there is no capture within the network by choosing client nodes such that they had comparable RSSI at the access point. Results are evaluated for both uplink and downlink performance with a saturated channel and equal offered load per client. Other experiment parameters were maintained as shown in Figure 2.

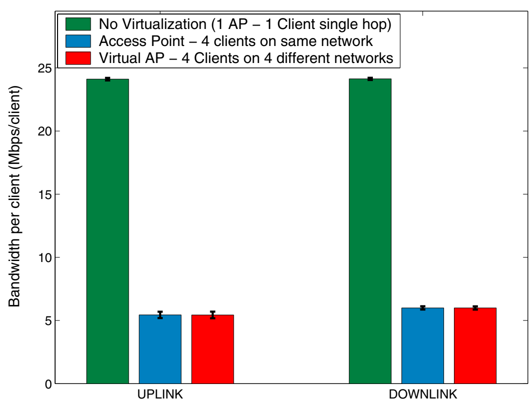

Fig. 3. Impact of virtualizing using channel multiplexing approaches.

Figure 3 plots the observed per client throughput ($\frac{Mbits/sec}{client}$) for uplink and downlink traffic. Performance of a single client with a single access point is taken as a reference for comparison. Key observations that can be made from the results are:

- As with any time sharing approach, the entire bandwidth (which is seen in the scenario with 1 client) is now shared across 4 clients.
- Uplink traffic sees a slight deterioration in performance with both the AP and the VAP as compared to the reference flow with 1 client.
- There is no added deterioration with uplink traffic using VAPs for having clients on multiple networks, as compared to an AP with all clients in one network. Hence, we can conclude that the deterioration is seen in both cases, which leads to a net channel throughput decrease of 9.75%. This decrease for the virtualized scenario as compared to no vitualization is due to the increased channel contention overhead.
- Downlink overheads for both AP and VAP with 4 clients are neglibible as compared to that with a single client.
- Error bars for both cases show little variance in throughput.

Hence we can conclude that using a VAP adds no conspicuous overhead to the throughput performance of an AP. We confirmed this behavior by investigating the source code for the MADWifi [9] driver where the VAPs are created. The driver does minimal additional processing to differentiate between the packets received for the different virtual interfaces. The above study suggests that experiments evaluating aggregate throughput with test setups running a single AP or multiple VAP should generate comparable results with the channel utilization being determined by the number of clients. Based on this conclusion, we can now compare the performance of virtualization with VAP and that with space separation based on:

1) Offered load
2) Packet sizes

B. Variation With Offered Load

Performance comparison of the VAP versus space separation (SDMA) uses the experiment setup as shown in Figure 4(a) and Figure 4(b). We compare the performance of both virtualization schemes by mapping four co-existing experiments. Each individual experiment consist of an AP-client single hop wireless.

Figure 5 shows the results for the aggregate throughput for virtualized experiments with varying offered load. We observe that below saturation both SDMA and VAP have the same performance. However, as the offered load is pushed into the saturation limits of the channel, there is a clear difference in the throughput.

The difference in performance observed in Figure 5 is due to the physical layer capture [11]. Capture is the phenomenon by which a receiver correctly decodes one of the many simultaneously colliding packets due to relatively higher received signal strength. Physical layer capture was detected either by sniffing packets from the channel with multiple sniffers (since the sniffers themselves are susceptible to capture) or by comparing the number of MAC retries with a case without capture. Figure 6 shows the the aggregate number of MAC retries with the VAP and the SDMA case. It is clearly seen that the number of MAC retries with SDMA were significantly lesser than with VAP since the receivers are able to decode colliding packets due to capture.

C. Variation With Packet Sizes

Packet sizes in a saturated channel impact both the MAC and physical layer overhead, as well as the aggregate channel access time. The goal of varying the packet sizes with experiments is to test if they have similar effect on performance with both the VAP and SDMA approach.

The setup of these experiments is the same as shown in Figures 4(a) and 4(b). To determine the effect of node

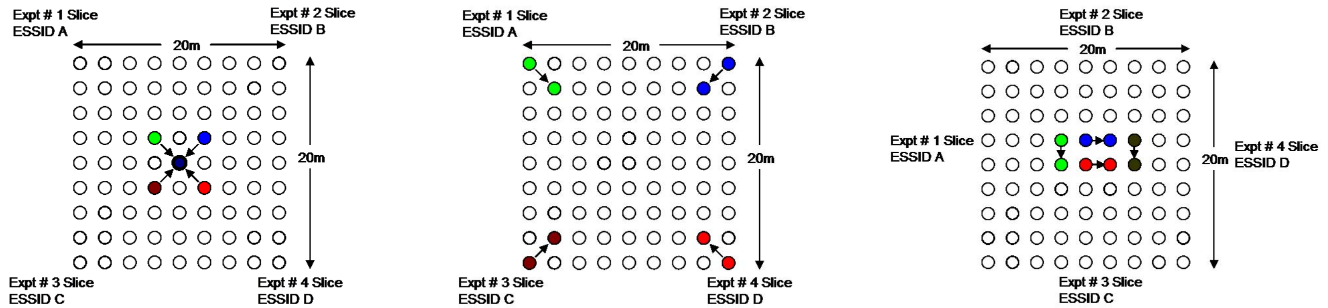

(a) Topology for VAP-based virtualization.

(b) SDMA scenario with maximum spatial separation between experiment nodes

(c) SDMA scenario with experiment nodes placed close together

Fig. 4. Experimental setup for performance evaluation of VAP and SDMA schemes on ORBIT.

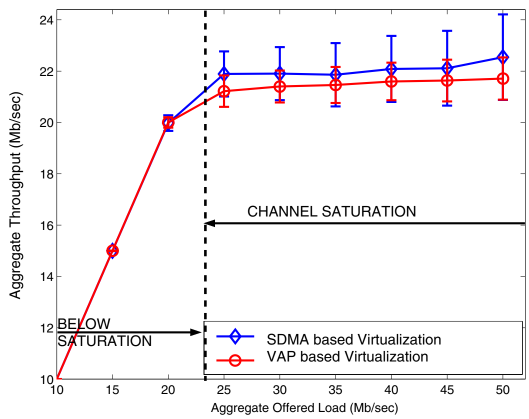

Fig. 5. A comparison of available bandwidth for SDMA and VAP based virtualization schemes supporting four concurrent experiments.

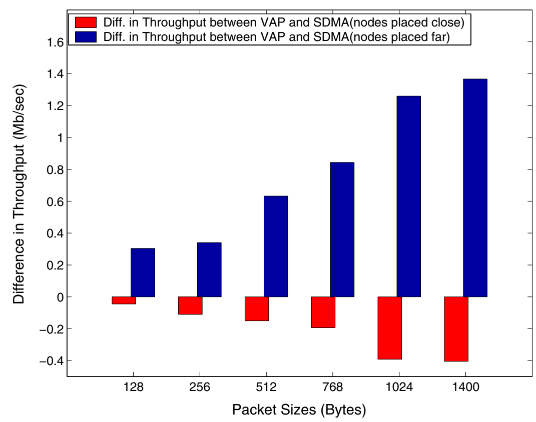

Fig. 7. A comparison of available bandwidth for SDMA and VAP showing the effect of space and transmission power control.

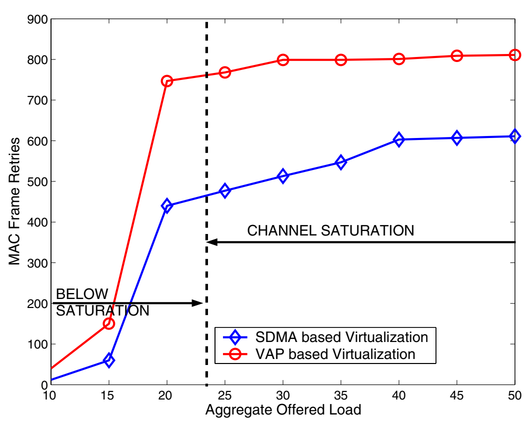

Fig. 6. A comparison of number of MAC frame retries for SDMA and VAP based virtualization schemes supporting four concurrent experiments.

positioning on the capture effect with SDMA, we measure SDMA performance with two setups as shown in Figures 4(b) and 4(c). In Figure 4(b) the nodes of the experiments are setup far from one another. In Figure 4(c) the experiments are setup next to each other. For each experiment packet sizes were varied and the aggregate throughput was measured. In Figure 7, we plot the difference in throughput of each of the SDMA setups from the VAP experiment and show the performance gains.

The SDMA setup with nodes placed far away had the advantage of decreased interference and improved performance with higher capture. The positive increase in difference in throughput shows that the benefits of capture increase with packet sizes. The SDMA setting without spatial separation shows a degraded performance as compared to the VAP setting. The MAC-ACKS in the downlink see lesser interference and collisions in the VAP due to time scheduled downlink transmission and hence the setting has a better performance as compared to the SDMA without spatial correlation. As the packet size increases this difference is even more pronounced since the effect of a collision is more pronounced for larger packet sizes.

IV. Delay-Jitter Comparisons

Experimenters often use delay as a metric measured for performance of an experimental setup. Jitter, defined as the variation of delay is also an important metric in the performance of real time traffic, such as voice or video. We will compare the effect of time and space separation for virtualization on both observed delay and jitter per experiment.

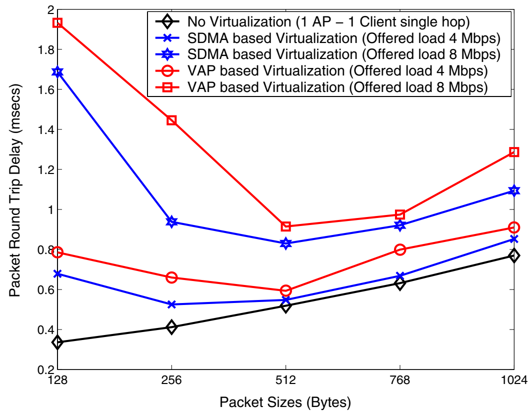

Fig. 8. Round trip delay variations with packet size for VAP and SDMA based virtualization schemes as compared to the non-virtualized scenario.

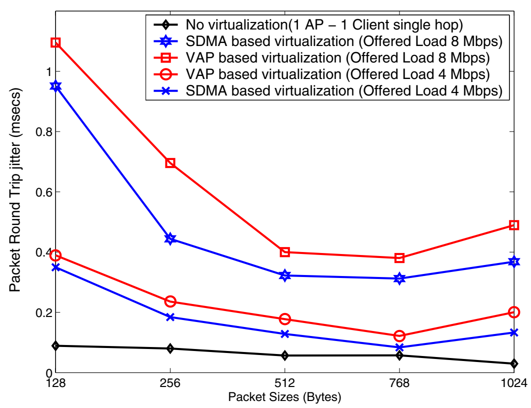

Fig. 9. Round trip jitter variations with packet size for VAP and SDMA based virtualization schemes as compared to the non-virtualized scenario.

Experiment setup for delay and jitter measurements is the same as that shown in Figures 4(a) and 4(b). Figure 8 shows the round trip delay measurements for the following cases: 1)No Virtualization, 2)SDMA and 3)VAP with different offered loads. We use two different offered loads to test the deterioration in delays with varying offered loads. With no virtualization, experiments show a linear increase in delay with packet sizes due to an increase in transmission times. This deduction is based on the assumption that the individual experiments have a one hop wireless topology with single flows. Hence there are no CSMA contentions. However, in the case of virtualization, experimenters have a V-shaped curve for delay results. The nodes of every experiment face CSMA contentions with nodes from other experiments. Delay values decrease with packet size for smaller packets as the CSMA contentions decreases with lesser number of packets. However for large packet sizes, the transmission and queueing times are more prominent than CSMA contentions and the delay increases with packet size. The per-packet delays for SDMA experiments are lower as a result of capture effect. Capture ensures that the MAC frames are received despite collision, which lowers the net MAC retries for getting a packet across and consequently the queueing delays.

Figure 9 shows the round trip jitter as a function of different packet sizes and offered load. The trend for jitter follows the same pattern as that for delay i.e., high for small packets, decreases for bigger sizes and slightly increases for the biggest packets sizes. However, unlike delay, the jitter decreases with packet size for no virtualization scenario. Since we measure RTT jitter, there is contention even with one hop, single flow topologies. Hence, as the packet size increases, for a constant offered load the number of contending packets decrease resulting in decreased jitter.

V. Inter-Experiment Interference Illustrations

Repeatability of experiments is strongly related to the isolation in the experimentation environment. Often it is seen that abuse of resource by one device sharing a resource leads to a deterioration in performance for other experiments sharing the same platform. We will elaborate the consequences of these inter-experimental effects with time and space separation for virtualization and suggest approaches to mitigate them. In this section, we use the same experiment set-up as used in the throughput, delay and jitter characterization of VAP and SDMA-based virtualization schemes. The experiment setup is shown in Figures 4(a) and 4(b).

A. Metrics

For lack of accurate delay characterization tools, we consider inter-experimental effect primarily in terms of throughput. To quantify the inter-experiment effects we define a coupling factor between virtualized experiments as:

$$\sigma_{(nv\_num,\ \ v\_num)} = \frac{(T_{non-virtualized} - T_{virtualized})}{T_{non-virtualized}} \quad (1)$$

$\sigma_{(nv\_num, v\_num)}$ indicates the coupling between non-virtualized experiment nv_num and virtualized experiment v_num. $T_{non-virtualized}$ and $T_{virtualized}$ represent the aggregate throughput of the experiments in the non-virtualized and virtualized cases respectively. A σ of 0 indicates an ideal experiment setup where there is no interference between experiments while a σ of 1 indicates complete interference of one experiment with the others. The Coupling Factor gives a direct indication of the level of interference expected between the virtualized experiments sharing a common wireless medium. Another approach is to use correlation among the throughput of virtualized experiments.

B. Coupling Factors

1) Throughput Coupling Factor: In this subsection we study the transient behavior of the experiments using VAP and SDMA-based virtualization schemes. In the scenario with four concurrent experiments for both VAP and SDMA, we observe the impact of the fourth experiment on the first three experiments for different traffic scenarios of the fourth experiment. We plot the coupling factor for the first three experiments with varying offered loads for both VAP and SDMA-based approaches. The packet size used by all four experiments was set to 1024 bytes. The plot of the throughput coupling factor is shown in Figure 10:

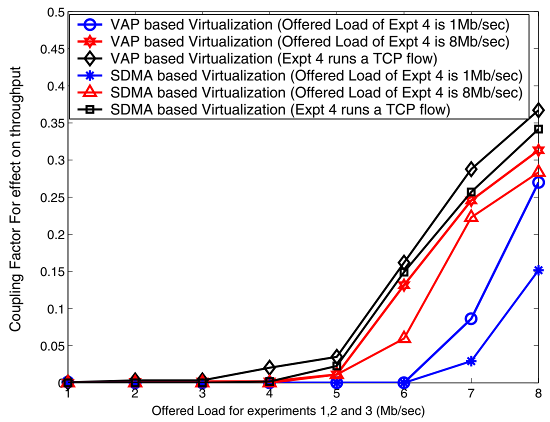

Fig. 10. Coupling Factor for effect on throughput of experiments due to other experiments.

- In the initial runs we kept the offered load of the fourth experiment at 1 Mbps and found that the coupling factors for both virtualization schemes is negligible for low offered loads but start to become prominent after the offered loads for the three experiments crosses 6 Mbps. Once the channel was driven into saturation, it effects the performance of the experiments. The effect is less for SDMA, since the performance of SDMA is better than the VAP.
- In the case where the offered load of the fourth experiment was about 8 Mbps the channel saturated at lower values of offered loads of the first three experiments and therefore the coupling factor was higher.
- In the case where the fourth experiment uses TCP, the coupling on other experiments observed was relatively higher than that with UDP. TCP flow pumped traffic at the maximum possible rate and its effect is more significant on the other experiments than that observed with a UDP flow. This increase can also be accounted by the overhead of the TCP-ACK traffic that increases the amount of contention among the different experiment flows.

2) Jitter Coupling: Similar to throughput results, the experimental measurements of packet jitter was affected by traffic from other experiments. We investigated jitter coupling in VAP and SDMA-based virtualization approaches by streaming a video from a client to an AP as a part of one experiment and running UDP flows as part of the other three experiment. Figure 11 shows the plot for jitter coupling factor values for videos of different bit-rates for VAP and SDMA-based virtualization scenarios. The jitter values were calculated for a real-time experiment that streams videos of different bit-rates from a client to an AP. With no virtualization, it was observed that the jitter of the video does not depend upon its bit-rate. However, in the virtualized case as the bit-rate increases the jitter value increases. Moreover, the jitter values of the video increased as the channel approaches saturation due to increase in the offered load of the other 3 UDP experiments. Similar to throughput results, the jitter coupling was more for the VAP setting as compared to that with the SDMA virtualization.

C. Summary

A comparative evaluation of the coupling factors for various offered loads and packet sizes have been shown. The coupling factor could be thoroughly evaluated by calculating it as a matrix. Where each experiment coupling with all other is measured with varying experimentation parameters such as packet sizes, offered loads and channel rates. On an average it is seen that the setup with SDMA behaves better than that with VAP in terms of relative coupling. However, the absolute values of coupling in both the cases are significantly high, thereby making the setup unsuitable for scientific experimentation. Thus to assuage inter-experiment effects we propose and implement a policy manager.

VI. TRAFFIC SHAPING/POLICY MANAGEMENT FOR VIRTUALIZATION

Results in Section V show that though the throughput coupling factor is lower for SDMA as compared to a VAP based approach, it is a non-zero entity and needs to be limited. Enforcing resource management across multiple experiments requires a systematic control framework for bandwidth assurance across multiple experiments.

A. Policy Manager

We will describe the implementation and testing of our policy manager with a VAP based setup. However, the same mechanism can be replicated and used without change for SDMA. A policy manager primarily performs the following functions:

- Admission control: To make a decision to allow or deny an experiment to share a VAP (or a slice for SDMA) with other experiments depending on its bandwidth requirements.
- Assigning And Enforcing Bandwidth: To allot the different experiments a maximum bandwidth value based on

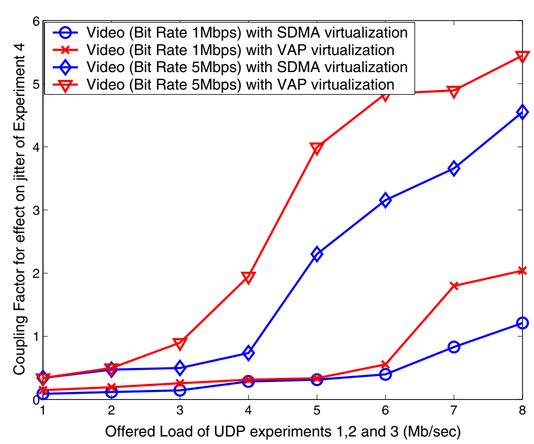

Fig. 11. Effect on jitter measurements in channel multiplexing virtualization approaches.

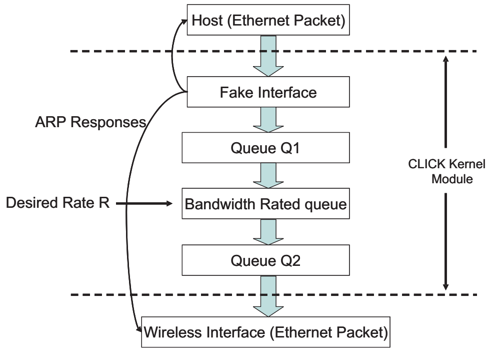

Fig. 12. Click Modular Router Elements for Bandwidth Shaping.

the number of experiments on a single VAP (slice) and their bandwidth requirements.

The Policy manager could be integrated with the experiment scheduling and resource tracking mechanisms to ensure that each of the experiments get a fair share of the resources. The experiments would be rate limited to their maximum assigned bandwidth, even before the experiment execution is started. Our implementation is based on a kernel module created with the CLICK modular router. The configuration setup for the CLICK[12] module is as shown in Figure 12.

Figure 13 shows the throughput results. For demonstration purposes, we allow each of the experiments to have unbounded bandwidth for the first 75 secs. We see that increased offered loads for experiments 1 and 2 results in performance degradation for experiment 3. Enforcement of the policy manager results in limiting experiment 1 and 2 to $6Mbits/sec$ and $4Mbits/sec$ respectively. Thus by artificially reducing the bandwidth available to each of the experiments we reduce the inter-experiment coupling factors to 0.

VII. Scheme Selection

Previous sections show the relative efficiency and inter-experiment coupling with the use of Space and time separation. Apart from these quantitative aspects other considerations for selection of a scheme are:

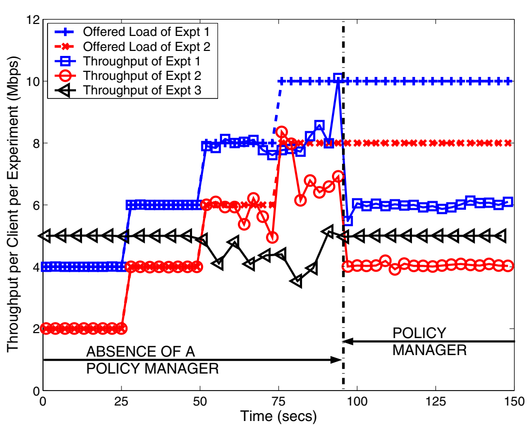

Fig. 13. Application of a policy manager in enforcing channel throughput.

- Topology: VAPs are limited to infrastructure mode setups, while SDMA can work with ad hoc as well.
- Space Separation: Achieving isolation and efficiency with SDMA requires considerable spatial separation between slices (of the order of $10dB$) or artificial stretching [8] of the testbed by use of noise.
- Scalability: Number of experiment slices with SDMA is limited due to the number of nodes, space constraints of the testbed and or the granularity of the noise generation mechanism.

Thus the quantitative approaches along with a qualitative comparison would possibly yield the best virtualization for a given testbed.

VIII. Conclusions And Discussion

Our study shows two approaches to channel conservation for a wireless testbed. Evaluation of the space and time separation scheme reveal benefits and weaknesses for both. Space separation provides relatively higher efficiency, lesser coupling between experiments. We layout selection criterions for each of these schemes based on the requirement of the testbed and finally propose and implement a policy manager for controlling inter-experiment interference. Finally, incorporating arbitrary topologies in a slice or across VAPs allocated to the experiment may be challenging or impossible for some experiments.

References

[1] Global Environment For Network Innovations (GENI). http://www.geni.net/
[2] Sanjoy Paul and Srini Seshan. Virtualization and Slicing of Wireless Networks. Technical Report GENI Design Document 06-17, GENI Wireless Working Group, September 2006.
[3] R. Mahindra et al., Integration of Heterogeneous Testbeds, Proceedings ACM TridentCom, Austria, March 2008.
[4] D. Raychaudhuri et al., Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols in WCNC LA, California, March 2005
[5] UML Jeff Dike A user-mode port of the Linux kernel, 5th Annual Linux Showcase and Conference, Oakland, California, November 2001
[6] G. Smith et al., Wireless Virtualization on Commodity 802.11 Hardware, WiNTECH07, Montreal, Canada, September 2007.
[7] B. Aboba. Virtual Access Points. IEEE document, IEEE 802.11-03/154r1. http://www.drizzle.com/~aboba/IEEE/11-03-154r1-I-Virtual-Access-Points.doc.
[8] S. Kaul, M. Gruteser, and I. Seskar, Creating Wireless Multi-hop Topologies on Space-Constrained Indoor Testbeds Through Noise Injection, Proceedings of IEEE Tridentcom, Barcelona, Spain, March 2006.
[9] MadWiFI Driver, www.madwifi.org.
[10] Virtual AP on madwifi http://madwifi.org/wiki/UserDocs/MultipleInterfaces
[11] S. Ganu, K. Ramachandran, M. Gruteser and I. Seskar, Methods for Restoring MAC Layer Fairness in IEEE 802.11 Networks with Physical Layer Capture, Proceedings REALMAN, Florence, Italy, May 2006.
[12] Click Modular Router. http://read.cs.ucla.edu/click/
[13] NTP: The Network Time Protocol in http://www.ntp.org/
[14] Manpage of tcpdump in http://www.tcpdump.org/
[15] IPERF TCP/UDP Traffic Generation Tool, http://dast.nlanr.net/Projects/Iperf/