

Interference Avoidance and Control *

Ramakrishna Gummadi*
*MIT CSAIL

Rabin Patra[†]

Hari Balakrishnan*
[†]UCB

Eric Brewer[†]

ABSTRACT

The throughput of a wireless network is often limited by interference caused by multiple concurrently active nodes. The conventional approach of using a “one-transmission-at-a-time” MAC protocol to combat interference leads to a significant loss of achievable throughput compared to schemes such as interference cancellation that keep all transmitters active simultaneously. Unfortunately, interference cancellation incurs significant computational complexity, and cannot be implemented with commodity hardware.

In this paper, we propose a practical approach for improving the throughput of interfering nodes using variable-width frequency allocation. We show that variable-width channels provide significant theoretical capacity improvements, comparable to interference cancellation for infrastructure networks. We design an algorithm that reduces interference by assigning orthogonal variable-width channels to transmitters. We evaluate a prototype implementation of this algorithm on an outdoor wireless network with ten long-distance links configured into point-to-point and point-to-multipoint topologies. We observe a throughput improvement of between 30% and 110% compared to the existing fixed-width channel allocation.

1 INTRODUCTION

Our goal is to build wireless networks with high aggregate throughput. So, extracting transmission concurrency is essential. But there is a trade-off: higher concurrency generally means higher interference. Previous work has been in one of two areas: MAC protocols that attempt to extract concurrency, and interference cancellation and its variants [8, 11].

For 802.11 networks, existing MAC protocols such as CSMA, Time-Based Fairness (TBF) [19] and CMAP [21] regulate concurrent transmissions carefully to ensure that collisions resulting from interference remain low. They allow only one transmitter to be active within a given channel at any time there is a risk of interference (i.e., whenever concurrent transmissions result in either packet being lost). However, serializing interfering transmissions imposes a fixed upper bound on the aggregate throughput, regardless of the number of interfering transmitters. Therefore, the average throughput per transmitter decreases with the number of interfering transmitters.

In contrast, interference cancellation (IC) deals with distinguishing between concurrently transmitted signals by de-

modulating and decoding all the interfering signals simultaneously. The theoretical concepts behind IC were developed in the 1960s [7, 18], especially in the context of spread-spectrum systems. Recently, researchers have investigated IC and related alternatives such as interference alignment and ZigZag decoding to mitigate the problems caused by interference [4, 8, 11]. Unfortunately, such receivers involve significant complexity because separating overlapping signals requires considerable signal processing. Moreover, the running time of such algorithms grows at least linearly with the number of concurrent transmissions a receiver overhears, none of which might ultimately be intended for the receiver.

In this paper, we ask the following question: is it possible approximate the optimal throughput provided by IC using simpler techniques, and, if so, under what conditions? We demonstrate a spectrum allocation algorithm that assigns variable-width channels to transmitters and keeps all transmitters active concurrently, thereby achieving a higher capacity than any fixed-width channel assignment scheme such as CSMA or TBF. Our result suggests that we should control interference while maintaining high concurrency.

The allure of using variable-width channels to control interference is that commodity wireless chipsets, such as Atheros and PRISM, support variable-width channels ranging from at least 5 MHz to 40 MHz [12, 13]. Recently, Moscibroda et al. [13] have studied variable-width channels to improve network throughput by allocating spectrum to APs (Access Points) based on their load. Similarly, Chandra et al. [5] have studied variable-width channels to improve a single link’s throughput and energy efficiency. Here, we study variable-width channels for their ability to improve throughput among multiple interfering transmitters with backlogged flows (i.e., flows which always have some data to send).

We show that, for infrastructure networks, using orthogonal variable-width channels on the uplink from the clients to the AP not only achieves the optimum sum-capacity of n concurrent transmitters predicted by Shannon’s theorem, but also improves the aggregate throughput over any fixed-width TDMA scheme such as CSMA or TBF by an additional $\Theta(\log_2(n))$ bits/s/Hz. The intuition is that maintaining the transmitters on non-overlapping channels theoretically eliminates interference, while narrowing their channel widths allows the total transmitted power to be the sum of all transmitters. Thus, the aggregate transmit and received powers are increased, without adding interference. We believe that this approach also exhibit good gains for mesh networks, though we do not discuss that setting in this paper.

*This work was supported by the National Science Foundation under awards CNS-0721702 and CNS-0520032, and by Foxconn.

We use this intuition to develop a spectrum allocation scheme called VWID (Variable WIDTH channels). When there are n concurrent transmissions on a given channel, VWID attempts to split the spectrum into n non-overlapping variable-width channels. The width of each channel is allocated to maximize each transmitter's throughput, subject to the constraint that no interfering transmitter receives lower throughput than it would have with fixed-width allocations. Thus, VWID provisions spectrum to reduce interference specifically, and complements MAC protocols such as CSMA that provide additional functionality such as ACKs and retransmissions to deal with both noise and interference.

We have implemented a VWID prototype and conducted a preliminary evaluation on an outdoor wireless testbed consisting of ten medium to long-distance links deployed in an urban area. We configured the testbed into point-to-point and point-to-multipoint topologies, thus representing networks typically encountered in rural point-to-point [15] and point-to-multipoint [16] settings. Even though our implementation is unoptimized, we find that VWID provides per-node throughput improvements ranging from 30%–110% by provisioning orthogonal variable-width channels to reduce interference.

2 VARIABLE-WIDTH CHANNELS IMPROVE THROUGHPUT

We analyze the throughput improvement produced by encouraging multiple concurrent transmissions using orthogonal variable-width channels compared to TDMA schemes such as CSMA and Time-Based Fairness (TBF) that use fixed-width channels. Using variable-width channels and enabling concurrent transmissions on these channels always increases the aggregate throughput compared to using fixed-width channels, because the total transmitted and received powers are increased, while interference is still kept in check. We can also address inter-node fairness by using sufficient channel widths that guarantee that every transmitter obtains at least the throughput it obtains under the original fixed-width allocation.

We consider a single cell with n clients and an AP. The AP has a single radio and antenna. Our primary result is that providing n concurrent transmissions between the AP and the n clients using orthogonal variable-width channels, whose width is proportional to received SINRs (signal to interference plus noise), can achieve higher aggregate throughput (by an additional $\theta(\log_2(n))$ bits/s/Hz) beyond the status quo.

Assume that the transmissions between the the clients and the AP are in the uplink, and that there is demand on all n links. Consider two backlogged transmitters 1 and 2 whose signals are received with powers P_1 and P_2 . The receiver noise power is N per Hz. If transmitter 1 alone is active, the capacity C_1 of 1, assuming a Gaussian channel, is given by the Shannon-Hartley theorem: $C_1 = \log_2(1 + \frac{P_1}{N})$

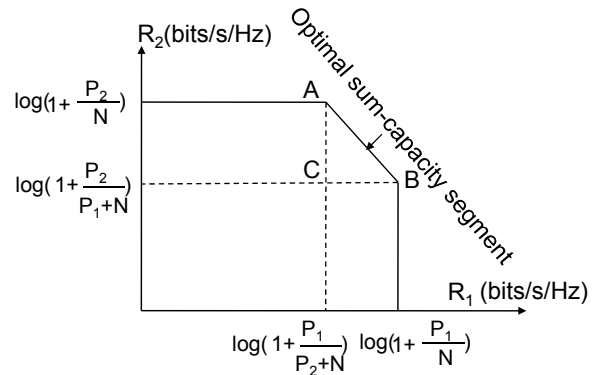


Figure 1: Achievable throughputs and the optimal capacity pentagon.

bits/s/Hz [7, 20]. Transmitter 1 can achieve any throughput rate R_1 that is less than C_1 [7, 20].

If both 1 and 2 are concurrently active, the theoretical capacity achievable by the two users simultaneously is called the sum-capacity. It consists of all throughput rate pairs (R_1, R_2) such that:

$$\begin{aligned} R_1 &< \log_2\left(1 + \frac{P_1}{N}\right) \text{ bits/s/Hz,} \\ R_2 &< \log_2\left(1 + \frac{P_2}{N}\right) \text{ bits/s/Hz,} \\ R_1 + R_2 &< \log_2\left(1 + \frac{P_1 + P_2}{N}\right) \text{ bits/s/Hz.} \end{aligned} \quad (1)$$

The achievable rates boundary, called the Cover-Wyner pentagon [1], is shown in Figure 1. The line segment A–B with slope -1 represents the optimal sum-capacity and is given by $R_1 + R_2 = \log_2(1 + \frac{P_1 + P_2}{N})$. The reason is that, no matter how the two users code their transmissions, independently or cooperatively, it is not possible for them to exceed the capacity limit that occurs when there is a single user with total received power $P_1 + P_2$. If both transmitters send on same frequencies, the rate pair at point A on the optimal sum-capacity segment in Figure 1 can be achieved by successive interference cancellation, in which the receiver first treats 2's signal as noise, recovers 1's signal, subtracts 1's signal from the total signal, and finally decodes 2's signal. Point B is vice-versa.

If we use variable width channels for 1 and 2 such that the total width is equal to the spectrum available to the receiver, we achieve non-interfering throughput rates for 1 and 2 that are given by:

$$\begin{aligned} R_1 &< \alpha \log_2\left(1 + \frac{P_1}{\alpha N}\right) \text{ bits/s/Hz,} \\ R_2 &< (1 - \alpha) \log_2\left(1 + \frac{P_2}{(1 - \alpha)N}\right) \text{ bits/s/Hz.} \end{aligned} \quad (2)$$

where α is the fraction of the spectrum allocated to 1 ($0 \leq \alpha \leq 1$). The noise term for R_1 in Equation 2 is reduced by a factor α because the signal is now confined to a narrower band, while noise still occupies the entire band with power N per Hz.

Theorem 2.1. *If transmitters 1 and 2 are continuously backlogged, then the aggregate throughput achieved with variable width channels is strictly higher than that with any TDMA scheme such as CSMA or TBF.*

Proof. The total rate R is given by:

$$\begin{aligned} R &= R_1 + R_2 \\ &= \alpha \log_2 \left(1 + \frac{P_1}{\alpha N} \right) + (1 - \alpha) \log_2 \left(1 + \frac{P_2}{(1 - \alpha)N} \right) \text{ bits/s/Hz.} \end{aligned}$$

Maximizing R by setting $\frac{d}{d\alpha} R = 0$ gives $\alpha = \frac{P_1}{P_1 + P_2}$, at which value $R = \log_2 \left(1 + \frac{P_1 + P_2}{N} \right)$ bits/s/Hz, which is optimal. Thus, we achieve the optimal throughput when transmitters are assigned channel widths proportional to their received power at the AP.

But no TDMA scheme, such as CSMA or TBF, is optimal (i.e., its throughput does not lie on the A–B segment) because TDMA only keeps one transmitter active at a time, thereby reducing the total transmitted and received powers. In particular, we calculate the CSMA and TBF throughputs below.

CSMA throughput. To a first order, CSMA allows equal number of channel accesses to nodes. The total achievable capacity under CSMA is as follows: Transmitter 1 takes $\frac{1}{R_1}$ time to send a bit, while 2 takes $\frac{1}{R_2}$ time to send its bit. Thus, CSMA sends 2 bits in time $\frac{1}{R_1} + \frac{1}{R_2}$. Thus, CSMA rate is $\frac{2}{\frac{1}{R_1} + \frac{1}{R_2}} = \frac{2R_1 R_2}{R_1 + R_2} \leq \sqrt{R_1 R_2}$ because the arithmetic mean of two positive numbers is not smaller than the geometric mean. Substituting $R_1 = \log_2 \left(1 + \frac{P_1}{N} \right)$, $R_2 = \log_2 \left(1 + \frac{P_2}{N} \right)$, we find that this rate is less than the optimal rate $R = \log_2 \left(1 + \frac{P_1 + P_2}{N} \right)$. Moreover, the relative performance of CSMA to optimal can be seen to be arbitrarily bad if, say, $P_1 \ll P_2$, because transmitter 1 ends up monopolizing the channel.

TBF Throughput. TBF allows equal channel access and fares better than CSMA, but its capacity is also lower than the optimal power-proportional variable-width allocation. This is because, in one second, transmitter 1 sends R_1 bits and transmitter 2 sends R_2 bits. So, the achieved rate $= \frac{R_1 + R_2}{2} < R = \log_2 \left(1 + \frac{P_1 + P_2}{N} \right)$, again using basic algebra. In the worst case (i.e., when $P_1 \ll P_2$), TDMA’s rate is half the optimal power-proportional allocation with two transmitters. \square

To obtain more insight into how concurrent transmissions improve throughput, consider an example with two transmitters t_1, t_2 whose SINRs at the receiver are 1 each. So, t_1 and t_2 achieve a throughput of $\log_2(1 + 1) = 1$ bit/s/Hz individually. When they transmit concurrently, a MAC such as CSMA shares the channel by time-division multiplexing it, so that each transmitter achieves a rate of 0.5 bit/s/Hz. On the other hand, dividing the channel into two and making t_1 and t_2 transmit concurrently allows each transmitter to achieve a throughput of $\frac{1}{2} \log_2(1 + 2) = 0.79$ bit/s/Hz, as shown




Figure 2: Throughputs of two transmitters when SINR=1.

in Figure 2. Each transmitter thus improves its throughput by about 30%, while the aggregate throughput increases by about 60%. Moreover, this throughput gain of variable-width channels relative to CSMA increases both with the imbalance between the received signal strengths of the concurrent transmitters, and with the number of concurrent transmitters. For example, if t_1 is 8 times (or 9 dB) stronger than t_2 (which can happen frequently with 802.11), the throughput improvement with variable-width channels increases to more than 2 \times .

Variable-width channels not only achieve higher throughput than any TDMA scheme with two concurrent transmitters, but we can also show that, n concurrent transmitters using variable-width channels can improve their aggregate throughput by an additional $\theta(\log_2(n))$ bits/s/Hz over TDMA. The reason is that the aggregate capacity of n transmitters using variable-width channels is $\log_2 \left(1 + \frac{nP}{N} \right)$ bits/s/Hz, assuming that the received powers P of all transmitters are equal. But TDMA schemes can only achieve a capacity of $\log_2 \left(1 + \frac{P}{N} \right)$ in this case. So, for large n , variable-width channels provide an additional $\theta(\log_2(n))$ bits/s/Hz increase in aggregate capacity.

3 VWID DESIGN AND IMPLEMENTATION

Using the insight that variable-width channels improve throughput, we develop an initial version of a variable-width channel assignment algorithm called VWID.

Our platform consists of high power Atheros 802.11a Ubiquiti XR5 radios (600 mW) that work in the 5 GHz spectrum. We use the Ubiquiti radio driver that allows variable channel widths (5, 10 and 20 MHz). While the normal 20 MHz-wide channel supports a maximum bit-rate of 54 Mbps according to the 802.11a standard, the half-width (i.e., 10 MHz) channel supports up to 27 Mbps, while the quarter-width (i.e., 5 MHz) channel supports up to 13.5 Mbps. In practice, we find that the achievable UDP throughput for outdoor links in our testbed is about 10 Mbps, due to interference [6] and multipath [2]. We found that both 10 MHz and 20 MHz channels attain this throughput, and that the 5 MHz channel obtains more than 8 Mbps.

Given a chunk of spectrum and a set of mutually interfering links, VWID assigns non-overlapping variable-width (i.e., 5, 10 or 20 MHz) channels to these links so as to control

Algorithm 3.1: ALLOCCHANNELS(InterferingLinks n)

Step 1 : Measure throughput $t[i]$ of each link i with n interferers
Step 2 : **for** $j \leftarrow 0$ **to** $n - 1$
 do $\begin{cases} \text{Step 3 : Measure } t'[i, c] \text{ of link } i, \text{ channel } c \text{ in } \binom{n}{j} \text{ choices} \\ \text{Step 4 : if } t'[i, c] < t[i] \text{ eliminate channel } c \end{cases}$
Step 5 : Return the channel c with highest $t'[i, c]$ for each link i

Figure 3: VWID channel selection algorithm.

interference. VWID only decreases the channel width for a link if doing so increases its throughput, thereby maintaining fairness. However, even with the fairness constraint, because VWID keeps every transmitter active while controlling interference, we find that per-node throughputs are higher for all links in many ($> 90\%$) scenarios, although these throughputs may be lower than those without the fairness constraint.

Our current implementation of VWID assigns variable-width channels to links within a single 20 MHz channel. The current best practice is to operate outdoor long-distance point-to-point or point-to-multipoint networks on a single channel because of spectrum scarcity, hardware limitations, and ease of management [3, 14, 15, 16]. So, while VWID can in principle handle any amount of spectrum, we have only instantiated and evaluated VWID for selecting 5, 10 or 20 MHz channel widths for a link.

In addition to selecting the channel width, VWID must also select the channel positions for 5 and 10 MHz channels within the 20 MHz channel. For example, assigning 5 MHz channels at 5.185 GHz and 5.195 GHz for two interfering links might be better in practice than assigning two channels at 5.185 GHz and 5.19 GHz, because the former provides more channel separation even if neither provides perfect orthogonality. So, for every link, VWID considers 4 choices for placing 5 MHz channels and 2 choices for placing 10 MHz channels, in addition to retaining the 20 MHz channel option. So, we have seven channel choices for each link.

VWID measures the throughput of each link for the seven channel-width choices under interference from other interfering links, and picks the channel and channel width that provides the highest throughput for that link. While we have shown that a power-proportional channel width allocation is optimal (§2), since commodity cards do not report SINR measurements accurately, we use throughput measurements. While the worst-case complexity of VWID is 7^n for n interfering links, we can prune the search space because we can reject channel widths and link combinations that violate the fairness constraint. For example, if a 5 MHz-wide channel tested under no interference is unable to provide more throughput than with a 20 MHz channel under interference, we can reject it immediately from all possible combinations with other links. Thus, in practice, VWID is efficient (for example, it only considers 16 combinations for four interfering links used in §4). The pseudocode for VWID is shown




Figure 4: A point-to-multipoint topology configured on the outdoor testbed.

in Figure 3. We defer the study of a more efficient channel-width assignment algorithm and its run-time dynamics such as measurement overhead, channel allocation effectiveness and stability for future work.

4 EVALUATION

We ran our experiments on our campus testbed, which consists of 6 wireless nodes and 10 links, 8 of which ranged from 1 km to 4 km, and 2 of which are co-located between different radios at P (Figure 4). Subsets of these links interfere with one another at either end-point, and each link interferes with at least one other link. The wireless nodes are based on 266 MHz x86 Geode single board computers running Linux kernel 2.6.19.2. The node at P has three wireless radios, the one at B has two radios and all the other nodes (S, B, E and Y) have one radio each. The nodes have directional antennas of 25 dBi gain. However, because of the relatively short distances involved, we were able to configure the links into various topologies such as point-to-point and point-to-multipoint by assigning the right transmit powers to the links. We selected a fixed bit-rate for each radio based on the maximum sustainable throughput (i.e., without getting disconnected after a while) across all its links.

We chose this outdoor setup because researchers have observed that interference imposes significant limits on achieved throughput, regardless of the supported bit-rates [6]. We modified the base driver to give us more control over MAC layer parameters such as disabling of ACKs, and changing the retransmission and Clear Channel Assessment (CCA) behavior. We experimented with various CCA settings that regulate the backoff behavior based on detected energy from concurrent transmissions. We disabled the CCA altogether and also varied the CCA energy-detection threshold between the card's minimum and maximum values. We measured unidirectional UDP and (bidirectional) TCP throughput under various CCA, ACK, and retransmissions configurations. We present results for UDP

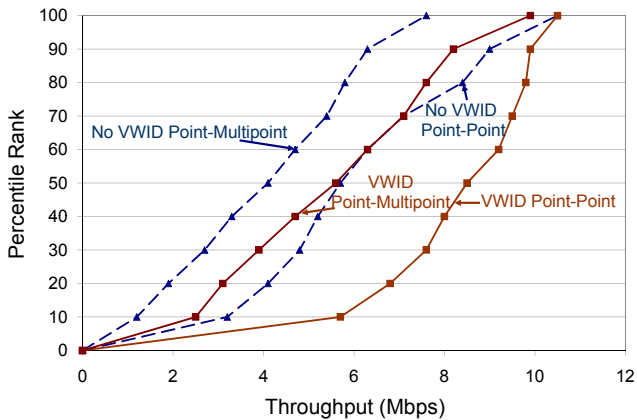


Figure 5: VWID increases CSMA throughput by 30%–110%.

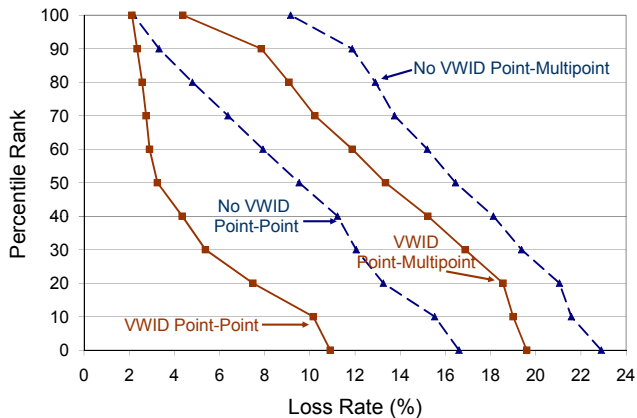


Figure 6: Loss rates correlate inversely with throughputs.

throughput with and without VWID under CSMA with the default ACKs, CCA and retransmission settings, because it achieved the highest throughput while maintaining fairness. While MAC configurations with CCA or ACKs switched off provide higher throughput for stronger links, they are either unstable or highly unfair to weaker links, so we do not consider them here.

Figure 5 shows the bi-directionally averaged percentile throughputs of one-way UDP flows across ten links as a CDF, and Figure 6 shows the corresponding loss rates as a CDF. Our main result is that, except for a high-throughput local link between two radios at node B, VWID improves throughput between 30%–110%. The highest improvements are for low-throughput links in low-throughput point-multipoint networks, because they suffer most from interference effects. So even modest interference relief is significant, which is a desirable outcome. While CCA mitigates some collisions, seven of the links have hidden terminals, leading to collision losses. The loss rates results exhibit good inverse correlation with the throughput plots and confirm that VWID controls interference and reduces losses, even if it means assigning narrower-width but orthogonal channels. Links in the point-multipoint topology show the biggest improvement, followed by point-point links. The reason is that

point-multipoint links are forced to share bandwidth if there are not enough radios at the APs (which is the case with APs at location P in Figure 4).

We also observed that, although TCP obtains lower throughput than UDP, its relative gains with VWID are higher. The reason is that the impact of reduced interference is more significant, as in the case with point-multipoint links in Figure 5. While we have used VWID with a CSMA MAC because it provided the highest throughput in our setup, our results are also applicable to new MAC protocols based on TDMA [15, 17] that have been proposed to avoid interference and provide concurrency with links tens of kilometers long. The motivation behind these these TDMA protocols is that nodes with multiple wireless radios operating on the same wireless channel are constrained from transmitting on one radio while simultaneously receiving on another radio. While this scheduling eliminates collisions, it forces the wireless nodes to synchronize their transmissions on all their outgoing links (and, similarly, receptions), thus making the scheduling of links and flow allocation more difficult [14]. We found that VWID creates more non-overlapping variable-width channels that relieves this scheduling pressure.

5 RELATED WORK

Current networks use either interference suppression on a packet-by-packet basis using MAC protocols [10, 14, 15, 16, 19, 21], or cope with interference using interference cancellation and related techniques such as interference subtraction, interference alignment and ZigZag decoding [4, 7, 8, 11, 18, 20, 22]. We have introduced the idea of interference control as a potential practical alternative, in which multiple transmitters operate concurrently, but take precautions to ensure they do not interfere with one another significantly. We use variable-width channels to achieve interference control, and ensure their orthogonality to avoid interference further.

While commodity hardware has supported variable-width channels out of necessity of narrow-width operation outside the unlicensed bands, this potential seems to have been recognized only recently. Moscibroda et al. [13] have used them for adjusting an AP’s channel width based on load, while Chandra et al. [5] have examined their properties in detail for the single-link case. We complement them by examining interference control using variable-width channels. As newer standards such as 802.11-2007 mandate narrow-width channels even in unlicensed bands, we can expect more commodity hardware to offer variable-width channel support.

6 CONCLUSIONS AND FUTURE WORK

We have shown that maintaining high concurrency by keeping multiple transmitters active concurrently, while controlling interference, increases the total system power without increasing interference, and hence increases aggregate throughput. We have examined the theoretical and practical potential of controlling interference using variable-width

channels. We have developed and implemented a preliminary channel allocation algorithm called VWID based on insights from a theoretical analysis of infrastructure networks. We evaluated VWID on a small campus testbed of outdoor links configured into point-to-point and point-to-multipoint topologies, and observed up to 2x throughput improvements with narrower-width but orthogonal channels that reduce interference, and, consequently, packet losses.

Our analysis and evaluation of variable-width channels are by no means complete, and point to several pieces of future work. First, while we believe our analysis extends to mesh networks, we would like to characterize the capacity region of mesh networks with variable-width channels. Further, our current VWID algorithm is simplistic, in that it uses a brute-force algorithm that has exponential worst-case complexity. We are currently enlarging the campus testbed and deploying VWID to carry real Internet traffic to residential users. We also plan to deploy VWID in wireless networks used in developing regions that we have access to, and learn more about the strengths and weaknesses of interference control. Also, in an accompanying paper [9], we study throughput improvement strategies for bursty traffic using spread-spectrum codes, and we plan to extend them to variable-width frequencies.

Acknowledgments. We thank Mythili Vutukuru and the anonymous reviewers for their valuable comments.

REFERENCES

- [1] R. Ahlswede. Multi-way communication channels. In *ISIT'71*.
- [2] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *MobiCom'05*.
- [3] S. Biswas and R. Morris. Opportunistic Routing in Multi-Hop Wireless Networks. In *HotNets'03*.
- [4] V. Cadambe, S. A. Jafar, and S. Shamai. Interference alignment on the deterministic channel and application to fully connected AWGN interference networks. Nov 2007. URL <http://arxiv.org/abs/0711.2547v1>.
- [5] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl. A case for adapting channel width in wireless networks. In *SIGCOMM'08*.
- [6] D. Gokhale, S. Sen, K. Chebrolu, and B. Raman. On the feasibility of the link abstraction in (rural) mesh networks. In *INFOCOM'08*.
- [7] A. Goldsmith. *Wireless Communications*. Cambridge University Press, 2005.
- [8] S. Gollakota and D. Katabi. ZigZag decoding: Combating hidden terminals in wireless networks. In *SIGCOMM'08*.
- [9] R. Gummadi and H. Balakrishnan. Wireless networks should spread spectrum based on demands. In *HotNets'08*.
- [10] R. Gummadi, R. Patra, S. Nedeveschi, S. Surana, and E. Brewer. A radio multiplexing architecture for high throughput point to multipoint wireless networks. In *WiNS-DR'08*.
- [11] D. Halperin, J. Ammer, T. Anderson, and D. Wetherall. Interference Cancellation: Better receivers for a new Wireless MAC. In *HotNets'07*.
- [12] HFA3863 Baseband Processor Data Sheet, <http://pdos.csail.mit.edu/decouto/papers/802-11-docs/hfa3863.ps>.
- [13] T. Moscibroda, R. Chandra, Y. Wu, S. Sengupta, and P. Bahl. Load-aware spectrum distribution in wireless LANs. In *ICNP'08*.
- [14] S. Nedeveschi, R. Patra, S. Surana, S. Ratnasamy, L. Subramanian and E. Brewer. An adaptive, high performance MAC for long-distance multihop wireless networks. In *MobiCom'08*.
- [15] R. Patra, S. Nedeveschi, S. Surana, A. Sheth, L. Subramanian and E. Brewer. WiLDNet: Design and implementation of high performance WiFi based long distance networks. In *NSDI, 2007*.
- [16] R. Patra, S. Surana, S. Nedeveschi, and E. Brewer. Optimal scheduling and power control for TDMA based point to multipoint wireless networks. In *NSDR, 2008*.
- [17] B. Raman and K. Chebrolu. Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks. In *MOBICOM'05*.
- [18] P. Stavroulakis. Interference suppression techniques: A twenty-year survey. *International Journal of Satellite Communications and Networking, Vol. 21, No. 1, 2003*.
- [19] G. Tan and J. Gutttag. Time-based fairness improves performance in multi-rate WLANs. In *USENIX'04*.
- [20] D. Tse and P. Viswanath. *Fundamentals of wireless communication*. Cambridge University Press, 2005.
- [21] M. Vutukuru, K. Jamieson, and H. Balakrishnan. Harnessing Exposed Terminals in Wireless Networks. In *NSDI'08*.
- [22] L.-L. Xie and P. R. Kumar. A network information theory for wireless communication: Scaling laws and optimal operation. In *IEEE Transactions on Information Theory, Vol 50, No. 5, 2004*.