

A Case for Adapting Channel Width in Wireless Networks

Ramya Raghavendra, Paramvir Bahl, Ranveer Chandra, Ratul Mahajan, Thomas Moscibroda
ramya@cs.ucsb.edu, {bahl, ranveer, ratul, moscitho,}@microsoft.com

April 12, 2008

One of the core design principles of IEEE 802.11 networks is the use of a simple, fixed channelization structure. The entire available spectrum is divided into smaller channels¹ of equal channel-width (i.e., or bandwidth, in electrical engineering terms), and every node operates on a specific channel at any time. Each channel is allocated the same amount of spectrum, albeit on different frequencies. For example, the 2.4GHz ISM band has 11 channels whose center frequencies are separated by 5 MHz width, of which three are non-overlapping channels each occupying 20MHz of the frequency spectrum. The reason for defining channels of fixed, a-priori defined channel-width is primarily simplicity. Network planning and coordination as well as the PHY and MAC layer protocols become simpler if the width of the channel remains fixed.

This fixed channelization structure is deeply ingrained in the wireless networking community and, by now, is typically taken for given. Consequently, the focus of research has been almost exclusively on improving wireless network performance within the scope of fixed channelization. The problem with this approach is that it ignores the possibility that the fixed, equal-width channelization structure itself may be a fundamental cause of many problems encountered in today's IEEE 802.11-based networks.

In this work, we take a fresh look at the channelization in IEEE 802.11 networks how that it is 1) practically doable and 2) actually beneficial to adaptively change the channel-width of wireless communication channels. Our experiments show that even with current, off-the-shelf network cards such as the Atheros card, it is possible to dynamically change the channel-width to, say, 5MHz, 10MHz, 20MHz, or 40MHz in software, as shown in Figure 1(a). This indicates that the hardware limitations of the IEEE 802.11-based fixed channelization are by no means insurmountable, and gives us a completely unexplored knob (among the popular existing knobs that researchers have investigated are transmission power, coding schemes and modulations, or frequency assignment), that can be used to optimize wireless systems.

We performed extensive measurements both in a controlled setting using signal attenuator, and in real-world settings both in an indoor network and outdoors. Intriguingly, our measurement results show that this novel knob may be a particularly powerful one. Our measurements indicate that adaptively varying the channel-width of channels has positive impact on many of the other major desirable metrics in wireless networking: *range/connectivity*, and *power consumption, throughput, and system capacity*.


Specifically, through our measurements, we make the following conclusions:

- At small communication distances, throughput increases with channel width. The increase is not proportional to the channel width due to MAC layer overheads.
- Decreasing the channel increases communication range. We get a 3dB improvement by halving the channel width due to better SNR. Narrower channel widths also have better resilience to delay spread. This is illustrated in Figure 1(b). When the sender and receiver were separated by a large distance, higher throughput was obtained on the narrow channels.
- Narrower channel widths consume lesser battery power when sending and receiving packets, as well as in the idle states. A 5 MHz channel width consumes 40% lesser power when idle, and 20% lesser power when sending packets than 40 MHz channel width.

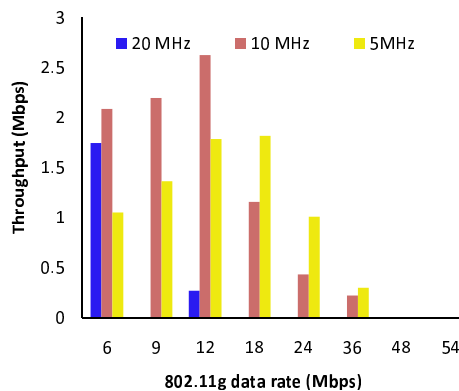
A wide variety of applications can leverage these properties of narrow channels to improve wireless performance. For instance, we can now have higher range at still save power; two quantities which were always in conflict earlier. In this section, we discuss a few applications. The reader can probably think of many others. In general, we believe that dynamically adapting channel width can benefit most aspects of wireless networks.

1. Load balancing for hotspots or enterprise WLANs. Large scale WLANs can have a skewed distribution of clients near individual APs, and load balancing becomes crucial to provide capacity and ensure fairness. To address this problem, researchers have proposed a multiplicity of solutions that can broadly be classified into three categories, depending on which configurable network parameter is used for the optimization. The most popular techniques are based on i) client-AP association ii) transmission power control (cell-breathing) and iii) intelligent channel assignment (i.e., adapting the center frequency); or combinations thereof. These solutions, while useful, provide limited benefits.

¹A channel in a wireless network is the frequency spectrum block over which nodes spread their transmissions. A channel is defined by its center frequency and channel-width.



(a) Screenshot of the spectrum analyzer showing 20MHz, 10MHz and 5MHz signals.



(b) UDP throughput of a flow when the sender and receiver are separated by 162 meters outdoors. The x-axis represents the G data rates for different physical layer coding schemes. Signals propagate better over lower bandwidths

For instance, one way to reduce load on an AP with many close clients is to reduce its power so that it serves fewer clients; but this can potentially cause clients close to this AP to associate to distant APs, which hurts performance. Similarly, while adapting center frequency helps minimize interference, it does not reduce load on APs with many close clients. In a sense, these solutions try to alleviate the symptoms, rather than solving the cause of the problem.

In contrast, adapting channel width can provide a more direct and conceptually cleaner solution. APs could be dynamically allocated channels of different widths (centered on different frequencies) where the width of an AP's channel is determined based on traffic demand of clients near it and the number of interfering APs in its vicinity. We have proposed a spectrum allocation scheme for this problem in [1]. Preliminary results show that this scheme significantly improves total network throughput and prevents starvation for clients that connect to popular APs.

2. Better support for QoS intensive applications. VoIP is becoming increasingly desirable for WLANs, but its usage is marred by interference from bursty data traffic and poor battery lifetimes. The capabilities provided by narrow channels—both in terms of increased range and reduction in battery power consumption—make it suitable for VoIP applications. One concrete solution is to set aside a “voice channel” on which data traffic is not sent. The width of this channel can be based on the workload. It will typically be narrow given the low bandwidth requirements of voice traffic. Thus, we can obtain high VoIP performance without setting aside a bigger than needed chunk of the spectrum and at the same time, improve the battery lifetime of power constrained WiFi devices.

3. Improved connectivity at edges. One of the benefits of narrow channels is an increase in range. This capability can be harnessed to provide enhanced wireless coverage. APs on the periphery of office buildings can be assigned narrower channels to increase their coverage region, especially in corners that are otherwise hard to reach [2] without a highly dense deployment of APs.

4. Improved capacity in mesh (multi-hop) networks. Early experiences with city-wide wireless mesh networks suggest that their total capacity is an important limitation, e.g. [3]. Adapting channel widths can help alleviate the capacity problem. When different channel-widths are allocated to “links” based on the traffic they carry, many links may operate on a narrower channel, leaving more spectrum for the heavily-loaded links. (This assumes that nodes have multiple radios so that they can operate on multiple channels.) Also, as multiple flows on narrow channels provide higher overall throughput than a single wide channel the total capacity of the mesh backhaul can be improved.

References

- [1] P. Bahl, R. Chandra, T. Moscibroda, Y. Wu, and Y. Yuan. Load-Aware Channel-Width Assignments in Wireless LANs. Technical Report MSR-TR-2007-79, Microsoft Research, June 2007.
- [2] R. Chandra, J. Padhye, A. Wolman, and B. Zill. A Location-based Management System for Enterprise Wireless LANs. In *Proc. of NSDI*, 2007.
- [3] Questions for Tropos: Does Google's mountain view network fold under pressure? <http://www.muniwireless.com/article/articleview/5395>.