Increasing Spatial Diversity at the Wireless Client

Szymon Jakubczak
Massachusetts Institute of Technology
szym@mit.edu

ABSTRACT

Diversity methods have been shown recently to benefit wireless users that experience "dead spots". When a single wireless transmitter-receiver pair cannot establish a reliable link, receptions from multiple independent receivers can be combined in a variety of ways. However, the practical benefits of such methods depend on the availability of receiver diversity and have so far been limited to the upload direction that carries a small fraction of typical consumer traffic.

In this paper we suggest that wireless users create their own receiver diversity by using external antennae, external wireless interfaces or employing their smartphone devices. Our measurements in the wild indicate that such diversity yields significant potential gains even if the ability to exploit it is limited by the current WLAN architecture. We propose a proof-of-concept implementation and evaluate it experimentally.

1. INTRODUCTION

Despite proliferation of the wireless LAN technology over the last decade, wireless network connections are still regarded as generally lossy, variable and hardly predictable. The light at the end of the tunnel is *wireless diversity*: the fact that the reliability of the wireless channel varies in time and space and that transient outage events caused by multi-path fading, shadowing and interference are largely independent between two different locations.

Spatial diversity methods have been shown in recent years to benefit wireless users that experience "dead spots". When the channel between the client device and the access point (AP) experiences poor performance, a spatially-independent realization of the channel might experience better quality. Thus, rather than moving the receiver into a "live spot", receptions from multiple independent receivers can be combined in a variety of ways, with the granularity ranging from bursts of packets to soft values, in order to improve the probability of successful reception.

What makes spatial diversity attractive is that, unlike frequency or time diversity, it does not consume additional resources on the wireless medium. However, the fundamental challenge in deploying these diversity-based solutions is that spatial diversity must exist in the first place, and most consumer devices have form factors of size insufficient to generate the neccessary spatial independence to maximize the diversity potential.

Until now exploiting spatial diversity by consumer devices has been limited by their small form-factor to MIMO and integrated antenna designs which address multi-path fading, but fall short of achieveing the full diversity potential of spatial independence. As a result the majority of the proposed spatial-diversity methods focus on the uplink or mesh scenario where multiple independent APs are likely to be available. The download direction which carries the majority of traffic in common consumer scenarios (e.g., web browsing) gets the short end of the stick and often benefits from temporal diversity only.

In this paper we propose a simple and readily available way to overcome this problem and bring the spatial diversity gains to the downlink. We advocate that the user deploys an additional receiver herself, thus creating receiver diversity. Such receiver could take the form of a passive external antenna, an external wireless interface wired to the main device or even a cooperating wireless device such as a smartphone as shown in Fig.1.

To support our proposal, we use packet-level traces collected in multiple local wifi hotspots to estimate and quantify the potential opportunistic gain from spatial diversity created by one additional receiver. We show that by simply employing their *smartphone* as an additional receiver for the laptop computer, users could substantially improve their download performance in common usage scenarios.

We discuss the system design requirements to enable such gains in the case when the diversity is pro-

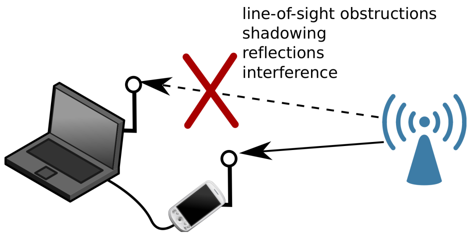

Figure 1: By employing their wifi-enabled smartphone as an external receiver, wireless users can better tap into the benefits of larger scale spatial diversity.

vided by an external device and present a proof-of-concept implementation under the current limited access to the wireless network stack.

2. FADING AND SPATIAL DIVERSITY

We begin by reviewing the mechanism of fading and spatial diversity, as well as the methods proposed to exploit diversity. Although, in this work we focus on static networks, we note that static wireless devices can be surrounded by a mobile environment, e.g. in the presence of humans.

2.1 Causes of Fading

Even in absence of external interference, a typical wireless channel is not stationary, but rather experiences time-, frequency- and space-dependent variation in the signal attenuation (and hence effective SNR), i.e., fading. The two major causes of fading are multi-path interference and shadowing.

Multi-path interference occurs when the radio signal reaches the receiver antenna by two or more paths of different optical length, for example one direct or line-of-sight (LOS) and the other reflected off the wall. Each signal "copy" will experience a different attenuation, delay and phase shift depending on the path length, number and type of reflections. Depending on their relative phase, this can lead to constructive or destructive interference of the two copies of the signal, amplifying or attenuating the signal power at the receiver. The phase shift dependends on the signal frequency, hence multi-path fading is frequency-selective. A small, on the order of the carrier wavelength, change in the path length is sufficient to turn constructive interference into a *deep fade*, or outage.

Because multi-path fading is fast, it is typically dealt with by exploiting frequency, time and small spatial diversity. Since different carriers will experience different fading, by using a multi-carrier system, such as OFDM, we can exploit the good channel realizations by interleaving the coded bits or QAM I/Q dimensions over the length of the OFDM symbol in the frequency domain [5]. Similarly, by interleaving the data in time or simple retransmissions, we can deal with short transient deep fades. Finally, by using multi-antenna arrays in MIMO systems, we can exploit the small-scale spatial diversity necessary to overcome multi-path-induced fades [5].

On the other hand, shadow fading is caused by transient or location-dependent obstruction of the line-of-sight path, for instance by a passing person. Shadow fading is typically slow, leading to burst losses on the order of tens to hundreds of packets or longer term outages which cannot be handled effectively using retransmissions or interleaving due to the incurred delay. Shadow fading cannot be overcome with compact multi-antenna arrays due to the much higher spatial correlation in LOS attenuation [5,13,15]. Therefore spatial diversity in shadow fading is fundamentally limited by the antenna packing problem on small form factor devices, whether client or access point.

Many spatial diversity schemes depend on availability of multiple physically distributed access points, in order to overcome shadow fading, often experienced as identifiable large dead spots when the obstacle is static [11,12,18]. The benefits of such systems is, however, limited by two factors:

- While multiple access points might be available in large wifi access hotspots such as university campuses, corporate offices or airports, smaller hotspots are often limited to a single access point.
- Multiple access points create spatial receiver diversity only in the uplink direction, which carries a small fraction of the traffic in most common Internet access scenarios. As a result, diversity schemes applied to the downlink exploit temporal diversity only [11,18].

We note that although transmitter diversity methods exist (e.g., space-time block codes, or single frequency networks [5]) they either require strict synchronization because the signal combining occurs on one receiver antenna (and thus must be aligned to within a symbol) or they reduce to simple route switching where only one transmitter is active at a time [12].

2.2 Diversity Combining

The key insight of receiver diversity schemes is that using multiple independent receivers offers significant gains over picking the best single receiver. In this section we discuss methods used to combine signals arriving at two receiver antennas. In Section 4 we will address link-layer implications of diversity combining.

Combining multiple receptions on an erasure channel is straightforward: simply use the bits/symbols/packets present in one received copy to fill in erasures in the other copy. For an erasure channel or in the outage model, if we have L independent channels and each of them has a probability of erasure/outage (e.g., packet loss) of p then the probability that all of them are under outage at the same time is reduced to p^L.

We note that this sort of model of spatial receiver diversity is the foundation for packet-level opportunistic routing schemes [2,4,8,9] where the receivers are distributed and need to forward the received packets further. In a distributed setting, the challenge is the coordination required to avoid forwarding multiple redundant copies, but when the two receivers are connected by a fast link of abundant capacity, e.g., Ethernet or USB, implementing such a scheme does not pose a substantial difficulty.

Multiple recent spatial diversity schemes seek performance improvent by reducing the granularity of such combining. If each packet is split into K chunks and chunks are lost independently with probability q, then with L independent receivers, we find that the probability that the packet-level combining fails of $\left(1-(1-q)^K\right)^L$ is further reduced at the chunk-level to $1-(1-q^L)^K$.

For instance, by fragmenting a packet into smaller chunks, each with an error-detection code (e.g., CRC), one could combine two errorneous packet receptions by choosing correct chunk from either copy at the price of the additional bits required for CRC [6]. A similar technique but relying on a single per-packet CRC to determine which combination of chunks is correct is proposed in [11].

Below the bits, some schemes operate on *soft values* provided by the physical layer [18]. A simple scheme would always choose the more confident value. A more efficient scheme, known as maximal-ratio combining, adds the two copies by weighing them in respect to their SNR [3,18]. Soft combining can be performed at the even lower level of the baseband signal [5].

Although all of the described schemes exploit the broadcast nature of the wireless transmission, one can also exploit diversity assuming unicast communications. For example, a smart MAC would let transmitters whose receivers experience better transient channel realizations take priority over transmitters who would need to use a low, inefficient bit-rate, while their receivers are experiencing a deep

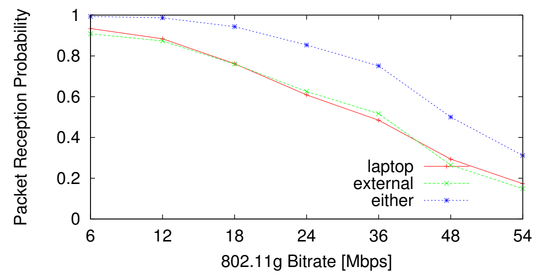

Figure 2: Average probability of successful reception for the observed bit-rates

	Throughput [kBps]			
	laptop	external	both	gain
SAME	1040	994	1104	1.09×
SEP	1074	1020	1460	1.36×

Table 1: Expected throughput estimated using the ETT metric. $gain = \frac{both}{(external+laptop)/2}$

fade [14,16]. For a single user, this reduces to fast route updates as in [12].

3. DIVERSITY POTENTIAL

In order to support our position, we collect packet-level traces using off-the-shelf 802.11 interfaces, as we focus on practical benefits available today.

Method: We run three wireless devices in monitor mode: one sniffer is placed near the access point and the other two, laptop and an external interface, are placed in a usage configuration. To quantify the spatial diversity gain, we compare two cases: SEP, when the two user receivers are about 1 meter apart, and SAME, when the two user receivers are placed adjacent to each other.

We collect this trace in a crowded hotspot location with only one AP. The human crowd is a responsible for both wireless traffic and shadow fading. We collect 10 traces, each 10 minutes long for both configurations. The packets of interest are large, longer than 1000 bytes data packets sent from the AP (FROM_DS). For each captured packet we are interested in the 802.11 bit-rate and a unique fingerprint that we can use to correlate the traces at all three receivers in a method similar to [10].

Devices: The wireless devices used for this experiment were: laptop with a Intel 4965agn mini-PCI card with the iwlagn driver, a Zonet USB interface using Ralink chipset and the rt2870sta driver and an Atheros AR2454 WiSoc embedded device (first-generation Meraki node) as the sniffer. We note that the Intel device has two diversity antennas embedded around the laptop screen.

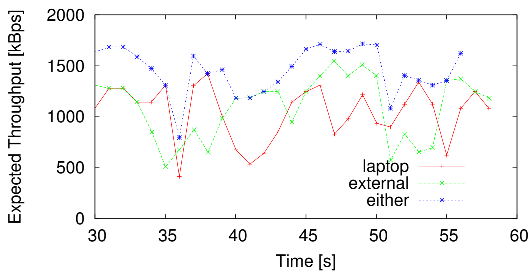

Figure 3: Expected throughput in one second window, estimated using the ETT metric

Results: Figure 2 shows the average packet delivery rate across all traces for each of the observed bit rates. It shows that both receivers together successfully captured substantially more packets than each separate, which hints at the possible opportunistic gain.

Given delivery probability, we can compute the Expected Transmission Time (ETT) as the shortest time to successfully deliver a single packet, taking into account the bit-rate, expected number of retransmissions and backoff [1].Figure 3 shows the expected throughput computed using the ETT metric for a fragment of one of the traces. Deep fades due to shadowing are clearly observable. Although one receiver might be within a passing person's shadow, the other receiver, thanks to the sufficient separation sees uncorrelated fading.

We quantify the potential overall gain in Table 1. The table reports the average expected throughput across all traces. Although adding an additional receiver adjacent to the laptop (SAME) does marginally improve the throughput, the gains from spatial diversity in the separated case (SEP) are much more substantial. We have observed in our experiments that the separate receiver could improve the throughput by 36%.

4. PRACTICAL DESIGN

Although the gains observed in the previous section are significant, we note that a practical implementation of such per-packet diversity scheme would require modifications to the link layer. In this sections we focus on client-side only design that would exploit existing features in the APs to benefit from multi-interface receiver diversity.

Acknowledging Opprtunistic Receptions: The 802.11 MAC requires that a successfully decoded unicast DATA packet is immediately followed by a synchronous ACK. This behavior is crucial for the MAC: automatic retry request (ARQ), auto-bit-rate adaptation and the binary exponential backoff all depend on it. To maximize gain from opportunistic reception in the standard 802.11 MAC, one would need to provide an ACK and suppress retransmission in case any of the receivers captured the packet. The technical difficulty comes from real-time requirements[1] which is why ACKs and other time-sensitive parts of the 802.11 protocol are normally implemented in hardware. Solutions proposed in the past include request-for-acknowledgment, a form of asynchronous ACKing [11,18].

A practical system would live above the link layer to be compatible with existing 802.11 devices. For instance, Divert [12] can provide fine-grained route switching on the downlink while being backward compatible with unmodified clients. However, it requires modifications at the AP side. Our goal is to provide the receiver diversity in the downlink direction without modifying the APs.

Diversity via Airtime Fairness: We make two observations to simplify this design. First, our experimental results indicate that shadow fading is slow and thus per-packet switching might not be strictly necessary in order to tap into the benefits from spatial diversity. Second, modern APs offer sophisticated client queueing features which enable us to exploit diversity with only minimum modification to the client.

After a decade of 802.11 proliferation, modern APs are now designed to be aware of client diversity. Specifically, bit-rate adaptation is now commonly performed individually for each associated client. Furthermore, to reduce jitter and improve overall medium utilization, APs frequently offer per-client queueing. This prevents head-of-line blocking by packets addressed to a receiver experiencing an outage. The queueing discipline often implements time-based fairness that takes into account the fact that transmissions to different clients might require substantially different amount of the wireless air-time [14,16]. For example, such technology is available today under names such as Airtime Fairness [17].

Thus, consider the case when a user connects two wireless clients to the same modern AP. Not only will the bit-rate adaptation for each client operate independently from the other, but also, the traffic to each client will compete for the wireless air-time in a way that promotes the traffic to the client with better instantaneous channel conditions. This way, the user will be able to receive some traffic on the

[1]ACK timeout is tuned to the round trip times of the wireless medium (propagation delay) and the hardware interface (slot time) and could be as low as 10us.

currently better interface while the other interface is within a shadow fade.

Although such design relies on changes at the client-side only and above the MAC layer, we note that its diversity gain is reduced in comparison to the fully opportunistic scheme described in Section 3. Looking back at Fig. 3, the potential gain of omniscient per-packet unicast routing is limited to the *envelope* of the curves *laptop* and *external*. On the other hand. the potential gain of air-time fairness is limited to the *arithmetic average* of the two curves, as the two clients would effectively time-share the medium in equal proportions. Applying this calculations to our traces we obtain the overall gain of 20% and 14% respectively. Finally, we note that to prevent potential exploitation, the queues of the two clients belonging to one user should each be assigned a weight of 0.5 in the fairness scheduler.

Diversity-Aware Applications: Even with two wireless interfaces, the user would see no benefit unless the application actually utilizes the diversity. A completely transparent solution could reside in the system kernel and bind each new socket to either interface in a multi-home routing configuration[2]. However, the very coarse granularity of such solution would likely make the benefit negligible.

To maximally utilize the multi-homing potential, the application should be aware of the diversity and balance its load across the two interfaces. The key to good balance is to keep the committment low, i.e., to not request more than the bandwidth-delay product necessary to utilize the bottleneck. Otherwise, a straggling request scheduled via the weaker receiver might unnecessarily delay the total transfer time [7]. We consider a hybrid solution for HTTP: an unmodified web browser using a diversity proxy which load-balances small HTTP requests across the two wireless devices.

Proof-of-Concept Experiment: To test our concept in practice, we configured an Android smartphone (Motorola A855 Droid) to forward IP packets via its USB device. We then configured a Linux laptop for multi-home routing using its main wireless interface and the RNDIS USB interface connected to the phone. We also manuually installed an ARP entry to ensure the AP knows how that the IP of the USB interface is reachable via the phone. In our prototype, two TCP connections are opened but bound to two different IP addresses, although routed via the same gateway (AP). We obtained the following result:

	Overall Throughput [kBps]
one TCP	920
two TCP at main interface	942
two TCP at two interfaces	1120

Although the gain is modest, it is available virtually for *free*, and requires only simple software configuration of the user's smartphone.

5. CONCLUSION

Wireless users today depend on spatial diversity provided by multiple APs, in order to overcome dead spots caused by shadow fading. We show in this paper that this leaves a substantial unrealized gain that cannot be achieved within the small form-factor of consumer devices. Thus, we enable users to increase their own downlink spatial diversity by employing a smartphone as an external receiver, and demonstrate a proof-of-concept software-only implementation finding up to 20% gain in throughput.

6. REFERENCES

[1] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *MOBICOM*, 2005.

[2] S. Biswas and R. Morris. Opportunistic routing in multi-hop wireless networks. In *SIGCOMM*, 2005.

[3] D. Brennan. Linear diversity combining techniques,. In *Proc. IRE*, volume 47, page 10751102, June 1959.

[4] S. Chachulski, M. Jennings, S. Katti, and D. Katabi. Trading structure for randomness in wireless opportunistic routing. In *Proc. of ACM SIGCOMM 2007, Kyoto, Japan*.

[5] S. Diggavi, N. Al-Dhahir, A. Stamoulis, and A. Calderbank. Great expectations: The value of spatial diversity in wireless networks. *Proc. of the IEEE*, 92(2):219–270, 2004.

[6] R. Ganti, P. Jayachandran, H. Luo, and T. Abdelzaher. Datalink streaming in wireless sensor networks. In *Sensys*. ACM, 2006.

[7] S. Jakubczak, D. G. Andersen, M. Kaminsky, K. Papagiannaki, and S. Seshan. Link-alike: using wireless to share network resources in a neighborhood. *SIGMOBILE Mob. Comput. Commun. Rev.*, 2008.

[8] J. Laneman and G. Wornell. Exploiting distributed spatial diversity in wireless networks. In *Allerton*. Citeseer, 2000.

[9] P. Larsson. Selection diversity forwarding in a multihop packet radio network with fading channel and capture. *SIGMOBILE Mob. Comput. Commun. Rev.*, 2001.

[10] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Analyzing the mac-level behavior of wireless networks in the wild. *SIGCOMM Comput. Commun. Rev.*, 2006.

[11] A. K. Miu, H. Balakrishnan, and C. E. Koksal. Improving loss resilience with multi-radio diversity in wireless networks. In *MOBICOM*, 2005.

[12] A. K. Miu, G. Tan, H. Balakrishnan, and J. Apostolopoulos. Divert: Fine-grained Path Selection for Wireless LANs. In *Mobisys*, Boston, MA, June 2004.

[13] N. Patwari and P. Agrawal. Nesh: A joint shadowing model for links in a multi-hop network. 2008.

[14] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly. Opportunistic media access for multirate ad hoc networks. In *MOBICOM*. ACM, 2002.

[15] T. B. Sorensen. Correlation model for shadow fading in a small urban macro cell. In *Proc. Personal, Indoor and Mobile Radio Communications (PIMRC)*, 1998.

[16] G. Tan and J. Guttag. Time-based fairness improves performance in multi-rate wlans. In *USENIX*, 2004.

[17] J. Wexler. Wi-fi vendors duke it out over airtime fairness. `http://www.networkworld.com/newsletters/wireless/2009/022309wireless2.html`, Feb 2009.

[18] G. Woo, P. Kheradpour, and D. Katabi. Beyond the bits: Cooperative packet recovery using phy information. In *Proc. of ACM MobiCom 2007, Montreal, Canada*.

[2]Advanced routing features of Linux 2.6 allow source routing via different devices even with the same gateway