

Capture-Aware Staggering for Concurrent Transmissions

Naveen Santhapuri
University of South Carolina
santhapu@cse.sc.edu

Srihari Nelakuditi
University of South Carolina
srihari@cse.sc.edu

Romit Roy Choudhury
Duke University
romit@ee.duke.edu

1. INTRODUCTION AND MOTIVATION

Spatial reuse in wireless networks is limited by the SINR threshold requirement. To satisfy this, neighbors of a receiver are expected to remain silent while a reception is in progress. Moreover, since 802.11 requires bidirectional exchange (i.e., nodes reverse their roles as transmitters and receivers), neighbors of *both* the transmitter and receiver must keep quiet for the entire duration of communication. This degrades spatial reuse, leading to low network throughput. Several approaches have been proposed to address this problem. In the spatial domain, power control, rate control, and carrier-sense adaptations have successfully identified possibilities of concurrency, while respecting the SINR requirement. In the temporal domain, optimizations such as piggybacked acknowledgments [1] have reduced *role reversals*, also enabling concurrency. Although useful, the improvements in all these schemes are bounded by the SINR requirement. Recent advancements in the physical layer are creating new possibilities. The SINR threshold is found to be a dynamic value, dependent on the relative order in which the signal and the interference arrive at the receiver. This implies that under certain conditions, it might be feasible to *capture* a data frame in the presence of concurrent interference. If harnessed carefully, this can be a powerful mechanism to improve the spatial reuse of wireless networks. This paper aims to design MAC protocols capable of exploiting capture capabilities in emerging wireless transceivers.

Capture has been understood as the ability to decode a sufficiently strong frame even if it arrives *after* the start of the interfering frame but *within* physical layer preamble time [2]. Recent advancements discussed in [3] demonstrate that it is possible to decode a frame even if it arrives at *any arbitrary point* after the start of the interfering frame, as long as a higher SINR threshold is satisfied. In Figure 1, observe that for different relative ordering of the interference and the frame of interest, the required SINR thresholds vary. Importantly, when a receiver is able to overhear the preamble of the interfering frame, it is better equipped to decode the frame of interest that arrives later. The high level intu-

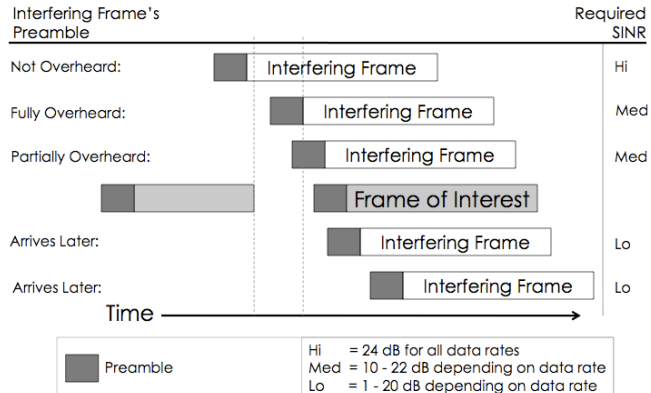


Figure 1: SINR variations with timing of packets

ition is that characterizing the behavior of the interfering frame (by receiving the preamble) is useful for separating out the signal from the interference. This paper incorporates this idea into a channel access scheme.

2. CAPTURE-AWARE STAGGERING (CAST)

The success of overlapping transmissions depends solely on satisfying the SINR requirement at the concurrent receivers. As we observed earlier, the SINR threshold is not fixed and varies with the timing of frame overlap. When only a particular order of frame overlap ensures concurrent reception, the MAC protocol needs to effect the suitable transmission order. So, capture-awareness needs to be incorporated into MAC protocols to enable concurrent communications where feasible.

To show how ordering helps, let us consider a simple example where all nodes are in the range of each other as in Fig. 2(a). Suppose, S1 and S2 have packets to transmit to R1 and R2 respectively. Now if S2 starts first, the two transmissions must occur sequentially. However, if S2 knows about the signal conditions at R1 and defers its transmission by one physical preamble time, two transmissions can finish in 1 packet time + 1 physical preamble time, as illustrated in Fig 2(b). This improvement in spatial reuse was possible because frames were *staggered* to take advantage of capture.




Figure 2: (a) 2 node pairs in range of each other; (b) Transmissions without and with staggering.

3. APPLICATION SCENARIOS OF CAST

We discuss our basic ideas and optimizations for a number of targeted architectures in the following.

Adhoc Multihop Wireless Networks: Assuming knowledge of two hop signal strengths, the staggering described in the previous section can be achieved in a distributed manner. When two neighboring transmitters have packets to transmit, the first one (primary transmitter) exchanges the RTS/CTS and waits for 1 physical preamble time before transmitting data. During this preamble time, the second (secondary) transmitter that overheard the RTS from the primary, will decide to initiate a concurrent communication based on the estimated SINR at its receiver, as well as at the primary receiver. Two options exist: (i) to start transmitting a concurrent DATA packet during the primary transmission’s wait time or, (ii) after primary starts transmitting DATA. In the example topology shown in Fig. 2(a), if S2 is the primary transmitter, S1 decides to start transmitting DATA¹ during the wait time. Both receptions can advance because low and medium SINR thresholds are satisfied at R1 and R2 respectively. Of course, several challenges arise. What happens when two secondary transmitters attempt to simultaneously initiate transmissions? How can the backoff mechanism be handled between such secondary transmitters? While our scheme addresses these questions systematically, we omit the discussion in the interest of space.

Gateway based Multihop Mesh Networks: A multihop wireless mesh network typically has a gateway that acts as an entry/exit point for all traffic. The gateway may be able to collect link quality, routing, and traffic information. A centralized gateway-assisted scheduling can fully exploit capture and staggering by eliminating (RTS/CTS) role reversals, and adopting piggybacked ACKs. Our idea is somewhat similar to the one proposed for enterprise WLANs in [4] but is designed to extract benefits from capture-awareness.

We now briefly describe our scheme. The gateway constructs a conflict graph and equivalence sets of links, taking into consideration the low and medium SINR

¹The primary transmitter’s RTS/CTS are not shown in Fig. 2(b). The secondary transmitter does not use RTS/CTS during a concurrent transmission

thresholds, as well as the order of active links for each packet (the downstream link must be scheduled after the upstream link for a given packet). All links in an equivalence set are activated together, ensuring that some are staggered appropriately. The order of staggering in a set depends on a priority assigned by the gateway (nearest to gateway get priority 1, next nearest priority 2, etc.). The links with priority i are scheduled at $(i - 1) \times$ (preamble time). This order will enable the more vulnerable receivers, to start receiving first and together achieve concurrency. The upward traffic problem can be addressed by assigning a proportional percentage of the time for upward traffic during which DCF is used to compete for the channel.

Enterprise Wireless LANs: The access points in enterprise WLANs are connected to a backend controller through a wired infrastructure. In such networks, nodes in the overlapping regions of multiple APs must undergo sequential communications to avoid interference. Our scheme addresses this by centrally scheduling the communications with appropriate staggering, and prescribing this schedule from the controller to individual APs. For example, consider a scenario with two APs, AP1 and AP2, each having two clients {1,2} and {3,4} respectively, with one packet to send to each of them. Assume that the clients 2 and 3 are in the overlapping region. Without CAST, it may take 4 slots ({AP1→1} {AP1→2} {AP2→3} {AP2→4}) to complete the transmissions. If the APs coordinate and stagger the transmissions, all of them can be completed in 2 slots ({AP1→1, AP2→3} {AP1→2, AP2→4}). This scheme can not only increase concurrency but also alleviate unfairness at clients of overlapping WLANs.

4. CONCLUSION

Neighbors of a receiver in wireless networks need to be silent to satisfy the SINR requirement, limiting the spatial reuse, which is further aggravated because of role reversals in 802.11. In this paper, we presented an approach to exploit capture by staggering of transmissions for enhancing spatial reuse. We are currently developing and evaluating the CAST approach further.

5. REFERENCES

- [1] N. Santhapuri et al, “Piggybacked-Ack-aided Concurrent Transmissions,” in *ICNP Poster Session*, Nov 2005.
- [2] A. Kochut et al, “Sniffing out the correct physical layer capture model in 802.11b,” in *ICNP*, 2004.
- [3] J. Lee et al, “An experimental study on the capture effect in 802.11a networks,” in *WinTech*, Sept. 2007.
- [4] N. Ahmed et al, “Extended abstract: Interference mitigation in enterprise w lans through speculative scheduling,” in *Proc. ACM Mobicom*, Sept. 2007.