Successive Interference Cancellation: A Back-of-the-Envelope Perspective

Souvik Sen
Duke University

Naveen Santhapuri
Duke University

Romit Roy Choudhury
Duke University

Srihari Nelakuditi
Univ. of South Carolina

ABSTRACT

Successive interference cancellation (SIC) is a physical layer capability that allows a receiver to decode packets that arrive simultaneously. While the technique is well known in communications literature, emerging software radios are making practical experimentation feasible. This motivates us to study the extent of throughput gains possible with SIC from a MAC layer perspective. Contrary to our initial expectation, we find that the gains from SIC are not easily available in many realistic situations. Moreover, we observe that the scope for SIC gets squeezed by the advances in bitrate adaptation, casting doubt on the future of SIC based protocols.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless Communication

General Terms

Experimentation, Performance

1. INTRODUCTION

Successive interference cancellation (SIC) is a well-known physical layer technique [1]. Briefly, SIC is the ability of a receiver to receive two or more signals concurrently (that otherwise cause a collision in today's systems). SIC is possible because the receiver may be able to decode the stronger signal, subtract it from the combined signal, and extract the weaker one from the residue. A natural question is: given SIC capable radios, what are the implications on MAC protocol design? What are the scope and limitations?

Inspired by these questions, we systematically study the ideal gains available from SIC. We focus on two simple topologies: (1) two transmitters sending to a common receiver, and (2) two transmitters sending to distinct receivers. We find that the characteristics of SIC in these simple topologies reflect on the behavior of larger, real world networks, such as enterprise or residential wireless LANs. We interpret the implications from a protocol designer's perspective. Our key observations may be summarized as follows:

(1) In the case of distinct receivers ($T_1 \rightarrow R_1$ and $T_2 \rightarrow R_2$), the gains from SIC are marginal.

(2) In the case of common receivers ($T_1 \rightarrow R_1 \leftarrow T_2$), SIC may offer modest MAC layer throughput gains if transmitters are carefully coordinated with techniques such as transmitter pairing and power reduction. However, somewhat counter-intuitively, the throughput gain is maximized when the system is forced to operate below the physical (PHY) layer capacity.

(3) We find that these behaviors hold even under various real-world network architectures (e.g., enterprise WLANs, where multiple APs are connected via a wired backbone).

Our observations may appear to be at odds with the high throughput improvements with SIC reported in [2]. The root of this discrepancy is in that, unlike [2], our study assumes that packets are transmitted at the ideal bitrate. One could argue that a practical rate adaptation scheme is unlikely to operate at the ideal bitrate, and there will always be a slack for SIC to exploit. Although true, this slack is fast disappearing with the recent advances in rate adaptation [3, 4]. Moreover, we believe that there is value in understanding the stand-alone benefits from SIC, when other factors are operating at the optimal point. This paper is targeted to improve this understanding.

2. CAPACITY GAINS WITH SIC

We begin with a brief PHY-centric overview of SIC, and build up the MAC layer interpretations subsequently.

2.1 SIC

Let us define *collision* as the simultaneous arrival of two or more packet transmissions at a receiver. Traditionally, only the strongest signal can be decoded, treating the other signal as interference. However, SIC facilitates recovery of even the weaker signal. For this, the bits of the stronger signal are decoded as before. The original (stronger) signal is then reconstructed from these bits, and subtracted (i.e., *cancelled*) from the combined signal. The bits of the weaker packet are then decoded from this *residue*. This can be an iterative process to recover multiple packets and hence it is termed *successive interference cancellation*. This

Hotnets '10, October 20–21, 2010, Monterey, CA, USA.

paper focuses on the simpler case of two packets only, i.e., interference cancellation is performed only once.

2.2 The Role of Bitrates

The decodability of two packets with SIC depends on their relative signal strengths and transmission bitrates. Let S_1^1 and S_1^2 be the received signal strengths at a common receiver R_1 from two transmitters T_1 and T_2 (see Fig. 1). Suppose B is the bandwidth and N_0 is the noise of the channel. When both T_1 and T_2 transmit concurrently, R_1 must decode the stronger signal first, say S_1^1, treating the weaker signal, say S_1^2, as interference. To be able to decode the stronger signal S_1^1, Shannon's theorem says that the highest feasible rate $\hat{r}_1^1$ for T_1's transmission to R_1 is

$$\hat{r}_1^1 = B\log_2(1 + \frac{S_1^1}{S_1^2 + N_0}) \tag{1}$$

Only if T_1 transmits at a rate $\hat{r}_1^1$ or below, it can be decoded successfully by R_1. After that, R_1 can attempt to decode T_2's signal. Assuming *perfect cancellation* of T_1's signal, the best feasible bitrate $\hat{r}_1^2$ for T_2 is

$$\hat{r}_1^2 = B\log_2(1 + \frac{S_1^2}{N_0}) \tag{2}$$

Interestingly, to facilitate SIC, *the stronger transmitter T_1's rate $\hat{r}_1^1$ may have to be lower than the weaker transmitter T_2's rate $\hat{r}_1^2$*. As we will see soon, this has important ramifications in SIC-aware MAC protocol design.

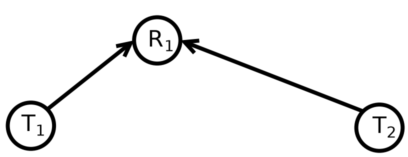

Figure 1: Two transmitters sharing a common receiver.

2.3 Capacity with SIC

Let us now contrast the capacity of a wireless channel with and without SIC, for a common receiver scenario as in Fig. 1. Without SIC, only one of T_1 or T_2 can transmit at a time, so the capacity of the channel, $C_{-\text{SIC}}$, is

$$C_{-\text{SIC}} = \max(B\log_2(1 + \frac{S_1^1}{N_0}), B\log_2(1 + \frac{S_1^2}{N_0})) \tag{3}$$

With SIC, it is possible to simultaneously receive two transmissions. The highest bitrates at which T_1 and T_2 can successfully transmit concurrently are $\hat{r}_1^1$ and $\hat{r}_1^2$ as given by (1) and (2). The corresponding channel capacity with SIC, $C_{+\text{SIC}}$, as derived in earlier works [5], is

$$\begin{aligned} C_{+\text{SIC}} &= B\log_2(1 + \frac{S_1^1}{S_1^2 + N_0}) + B\log_2(1 + \frac{S_1^2}{N_0}) \\ &= B\log_2(1 + \frac{S_1^1 + S_1^2}{N_0}) \end{aligned} \tag{4}$$

Now consider the relative capacity gain with SIC, $\frac{C_{+\text{SIC}}}{C_{-\text{SIC}}}$. The gains are plotted in Fig. 2 as shades of color (lighter the

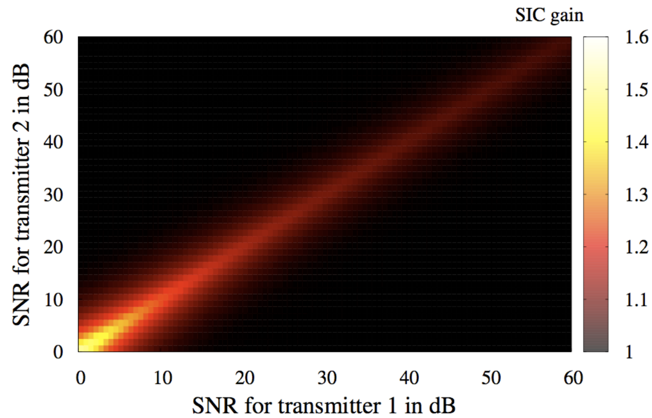

Figure 2: SIC capacity gains are not high in general but are larger when RSSs are smaller and similar.

shade, higher the gain). S_1^1 is on the x-axis and S_1^2 is on the y-axis. The key observation is that the channel capacity with SIC is always better than the individual capacities of any single transmitter, and *the relative gain is more when the received signal strengths (RSSs) are similar.*

3. SIC: MAC LAYER PERSPECTIVE

Similar RSSs increase the relative capacity gain with SIC; this implies that in trying to realize these gains, the transmission rates of the two packets have to be dissimilar. This is evident from Equation (1) and (2), where rate $\hat{r}_1^1$ depends on the ratio of the two RSS values, but $\hat{r}_1^2$ only depends on the ratio with noise. Therefore, when the two RSS are similar, $\hat{r}_1^1$ will be low, and $\hat{r}_1^2$ will be much higher in comparison. Converting this to the transmission time of packets, we note that one packet will incur a long air-time, while the other packet (transmitted in parallel) will finish much quicker. Ironically, the transmitter that experiences longer air time (i.e., T_1) actually has a stronger signal to the receiver; but its rate must still be low, because, to achieve SIC, it has to cope with the interference from the other transmitter (T_2). This rate disparity wastes channel capacity, creating a "hole" as shown in Fig. 3. Filling the "hole" by increasing the packet size or transmitting a train of packets is impractical – protocol limits on packet sizes prevents the former, while PHY layer synchronization issues make the latter difficult (see Section 5.4). Hence, the MAC protocol throughput suffers in practical settings even though the PHY layer is configured to attain SIC capacity.

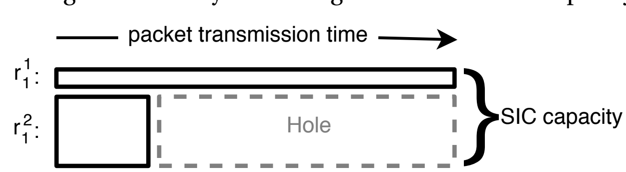

Figure 3: "Hole" created due to rate disparity with SIC.

One way of minimizing the hole could be power control; T_1 and T_2 could coordinate their transmit powers so that their bitrates become equal. While this will indeed improve MAC throughput, recall that the PHY layer capacity gain (with SIC) is maximized when the rates are disparate. Put differently, the strategy that maximizes MAC throughput forces the system to operate below PHY layer capacity, exposing an interesting tension between MAC and PHY layer perspectives. The optimal strategy for one is suboptimal for the other.

The rest of this section expands on these observations, characterizing the nature of the gains when coordination among nodes is an option. However, not all scenarios are amenable to lightweight coordination (e.g., two links in neighboring homes). We will study this scenario too, and quantify the gains without link coordination. Equipped with an understanding of these building-block scenarios, we will visit generic network architectures in the next section.

3.1 Two Transmitters to the Same Receiver

Consider the scenario in Fig. 1. Assume that transmitters T_1 and T_2 are transmitting one packet of length L bits each, to the common receiver R_1. Without SIC, they have to transmit sequentially and the total time needed (discounting MAC related overheads such as backoff) is

$$Z_{-\text{SIC}} = \frac{L}{B\log_2(1+\frac{S_1^1}{N_0})} + \frac{L}{B\log_2(1+\frac{S_1^2}{N_0})} \tag{5}$$

With SIC, both the packets are transmitted concurrently and therefore the completion time is dictated by the lower bitrate transmission. Assuming $S_1^1 > S_1^2$, the total time needed to transmit both packets with SIC is

$$Z_{+\text{SIC}} = \max\left(\frac{L}{B\log_2(1+\frac{S_1^1}{S_1^2+N_0})}, \frac{L}{B\log_2(1+\frac{S_1^2}{N_0})}\right) \tag{6}$$

The SIC gain, i.e., the ratio $\frac{Z_{-\text{SIC}}}{Z_{+\text{SIC}}}$ ($Z_{+\text{SIC}}$ is in the denominator reflecting the gain in time), is plotted in Fig. 4. Observe that, as the difference between RSSs increases (i.e., as we move from the middle towards the axes), the gains begin to increase up to a point and then start decreasing. The reason is as follows. Equation (6) is maximized when denominators of two terms are minimized. Since the two terms have inverse relationship to each other, the maximum of the two would be minimized when they are equal, i.e., $\frac{S_1^1}{S_1^2+N_0} = \frac{S_1^2}{N_0}$. Hence, the SIC gain peaks when S_1^1 is roughly the square (twice in terms of SNR in dB) of S_1^2. This is exactly when the bitrates become equal, eliminating the "hole" in Fig. 3. Thus, from a protocol designer's perspective, the two transmitters should be coordinated such that the stronger transmitter's SNR at the receiver is close to twice that of the weaker transmitter's.

3.2 Two Transmitters to Different Receivers

Now consider T_1 and T_2 transmitting concurrently to different receivers R_1 and R_2, respectively. Let S_j^i denote the RSS from transmitter T_i to receiver R_j. We study the feasibility of SIC in each of the 4 possible cases shown in Fig. 5.

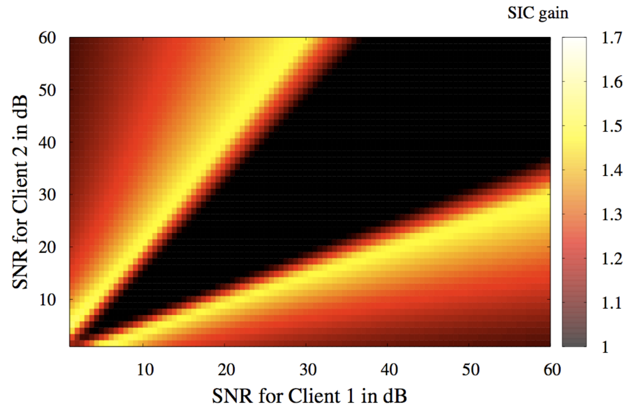

Figure 4: Two transmitters to the same receiver: SIC gains most when RSSs are such that the resulting bitrates are the same for both transmissions.

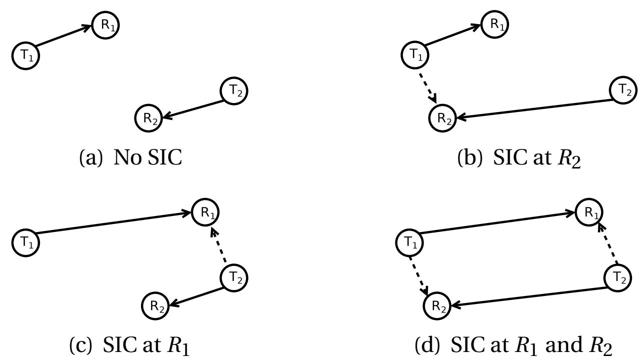

Figure 5: Two transmitters to different receivers: signal of interest (solid) and strong interference (dashed)

(Fig. 5(a)) $S_1^1 > S_1^2$ and $S_2^2 > S_2^1$: Signal of interest is stronger than the interference. So SIC is not needed here.

(Fig. 5(b)) $S_1^1 > S_1^2$ and $S_2^2 < S_2^1$: The RSS of T_1 is stronger at R_1, so SIC not needed at R_1. But at R_2, RSS of T_2 is weaker than that of T_1, so SIC can aid R_2. However, for SIC at R_2, it needs to decode T_1's transmission. The optimal rate for T_1's transmission to R_1 is $B\log_2(1+\frac{S_1^1}{S_1^2+N_0})$ whereas the permissible rate for T_1's signal at R_2 is $B\log_2(1+\frac{S_2^1}{S_2^2+N_0})$. Therefore, SIC is feasible at R_2 only if $\frac{S_2^1}{S_2^2+N_0} > \frac{S_1^1}{S_1^2+N_0}$. Neglecting noise, this means that the ratio of S_2^1 and S_2^2 should be greater than that of S_1^1 and S_1^2. Translating this RSS relationship to relative distances, the necessary conditions for SIC are: (1) T_1 has to be *closer* to R_2 than its own receiver R_1; (2) R_2 has to be *closer* to T_1 than its own transmitter T_2. Even if these conditions hold (i.e., SIC is feasible), gains may not be obvious – serial transmissions on the two links may finish sooner than concurrent transmissions. We will evaluate the gains at the end of this section.

(Fig. 5(c)) $S_1^1 < S_1^2$ and $S_2^2 > S_2^1$: Similar to the above case with the roles of strong and weak pairs reversed.

(Fig. 5(d)) $S_1^1 < S_1^2$ and $S_2^2 < S_2^1$: SIC is needed at both receivers. So the conditions similar to those above have to be satisfied at both R_1 and R_2. The difference is that, in this case, the optimal rates for the pairs $T_1 \rightarrow R_1$ and $T_2 \rightarrow R_2$ are $B\log_2(1+\frac{S_1^1}{N_0})$ and $B\log_2(1+\frac{S_2^2}{N_0})$ respectively, i.e., as if there is no interference, owing to SIC at each receiver. Therefore, SIC is feasible at R_2 only if $\frac{S_2^1}{S_2^2+N_0} > \frac{S_1^1}{N_0}$ (and a similar condition at R_1). When both conditions are satisfied, $Z_{+\mathrm{SIC}}$ would be

$$Z_{+\mathrm{SIC}} = \max\left(\frac{L}{B\log_2(1+\frac{S_1^1}{N_0})}, \frac{L}{B\log_2(1+\frac{S_2^2}{N_0})}\right) \qquad (7)$$

and $Z_{-\mathrm{SIC}}$ is the sum of the two terms in Equation (7). Although $Z_{+\mathrm{SIC}} < Z_{-\mathrm{SIC}}$ (raising hopes for SIC), topologies like Fig. 5(d) are unfortunately not common in real life.

To evaluate the performance of SIC in these scenarios, with four RSS variables, we use the Monte Carlo method. We fix the positions of the transmitters separated by a certain *range*. The receivers are then placed randomly within the range of their transmitters. We compute RSS based on the the transmitter-receiver distance, using path loss exponent α=4. Using these RSS values, the gain with SIC is computed as $\frac{Z_{-\mathrm{SIC}}}{Z_{+\mathrm{SIC}}}$. The simulation is repeated over 10,000 times. Fig. 6 shows results for different ranges (gains from lower pathloss exponents and other ranges, not reported here, are even lower). These results confirm that topological conditions for SIC are stringent, resulting in limited gains in most cases.

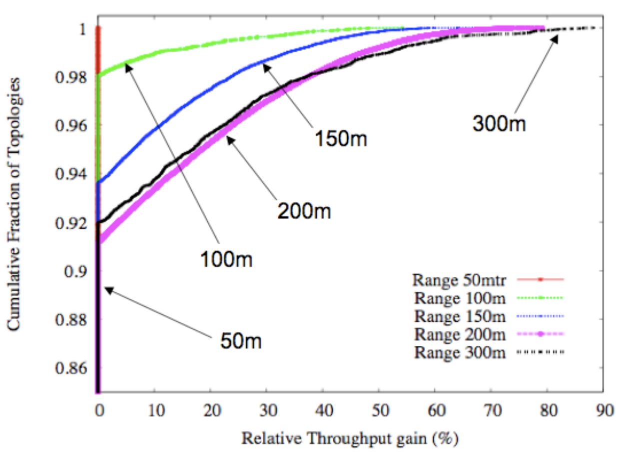

Figure 6: Two transmissions to different receivers: no gain from SIC in 90% of the cases.

4. SIC IN DIFFERENT ARCHITECTURES

We have used toy topologies to characterize the benefits of SIC; this section argues that the observations are generalizable to different wireless architectures.

4.1 Enterprise Wireless LANs

APs in enterprise wireless LAN (EWLAN) environments, such as corporate campuses, are connected through a wired backbone (see Figure 7(a)). This enables coordination of download and upload traffic.

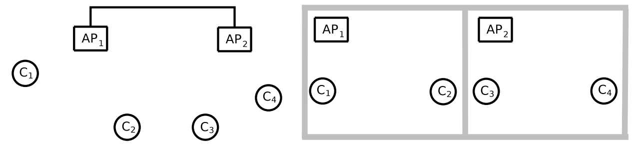

Figure 7: Different architectures: (a) Enterprise WLAN; (b) Residential WLAN (adjacent apartments)

Upload Traffic: Two Clients to One AP

First, consider the simple scenario of two clients, say C_1 and C_2, each having a packet to send to AP_1. This is no different from the two transmitter and one receiver scenario in Section 3.1. Hence, SIC can improve upload performance, particularly when the RSSs at AP_1 from C_1 and C_2 are such that their bitrates are equal under SIC.

Download Traffic: Two APs to One Client

Now consider the download traffic case where a client C_2 is within the range of AP_1 and AP_2. Since these APs are connected through a wired backbone, packets can be delivered to C_2 via either of the APs. SIC achieves higher throughput for C_2 by having both APs send packets simultaneously to C_2. With SIC, this scenario is no different from the above upload scenario. Without SIC, the two packets have to be transmitted sequentially. However, *we have the option of minimizing the total time by transmitting both the packets from the stronger AP.* So, the time for transmitting two packets without SIC is

$$Z_{-\mathrm{SIC}} = \frac{2L}{\max(B\log_2(1+\frac{S_2^1}{N_0}), B\log_2(1+\frac{S_2^2}{N_0}))} \qquad (8)$$

The gain from SIC then is (8)/(6), plotted in Fig. 8. Modest gains are available when the RSSs are such that one is roughly square (twice in terms of dB) of the other. But overall, the gains with SIC are quite limited in this download scenario.

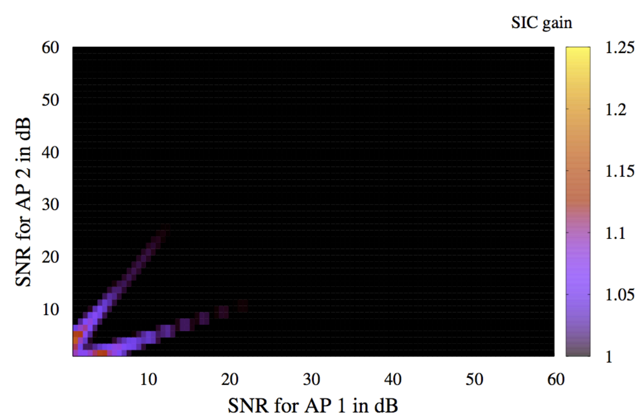

Figure 8: Download traffic from two APs to one client in an EWLAN: very little benefit from SIC.

Upload Traffic: Two Clients to Two APs

We already observed that SIC does not gain much in cases where two transmitters are sending packets to different receivers. We argue that the gains are even less in EWLANs. SIC is needed only when C_1 or C_2 sends a packet to AP_2 simultaneously when C_3 or C_4 is sending to AP_1. Given that a client has the choice of sending the packet to any of the APs, transmission to the closest AP is a better option. Then, each AP's signal of interest will be stronger than interference (as in Fig. 5(a)), and hence, SIC is not necessary.

Download Traffic: Two APs to Two Clients

Another scenario of two transmissions to different receivers is when two APs send packets to two different clients. Again, SIC figures only when AP_2 is delivering a packet to C_1 or C_2 in concurrence with AP_1's transmission to C_3 or C_4. Given that packets can be delivered to clients through any of the APs, there is no gain in choosing a farther AP, and thus SIC.

4.2 Residential Wireless LANs

Residential wireless LANs (RWLANs) in adjacent apartments could be like in Fig. 7(b). Unlike in EWLANs, in RWLANs, a client may not have the option of passing a packet through the neighbor's AP. Packets meant for C_2 can only be delivered through AP_1 even if C_2 is closer to AP_2. Strangely, this restriction provides some opportunities for SIC. In this example, if C_2 performs SIC, $AP_1 \rightarrow C_2$ can be concurrent with $AP_2 \rightarrow C_4$. On the other hand, $AP_1 \rightarrow C_2$ can not be concurrent with $AP_2 \rightarrow C_3$. This is because, the optimal rate for $AP_2 \rightarrow C_3$ is higher than that supported by RSS of AP_2 at C_2 and hence C_2 may not successfully decode $AP_2 \rightarrow C_3$ packet. In other words, SIC gains only when the client's own AP is farther than the neighbor's AP, and the client is closer to neighbor's AP than its own client. In essence, RWLANs offer some opportunities for SIC in apartment complexes crowded with APs, but the necessary conditions may not hold in most cases. Of course, upload traffic gains from SIC as in Section 3.1.

5. TECHNIQUES TO FACILITATE SIC

The summary thus far is that SIC is not quite helpful in the distinct receiver scenario, but beneficial in the case of a common receiver, such as in upload traffic from clients to the AP. Moreover, SIC offers the best possible gain when the RSS of the concurrent signals at the receiver are such that they yield the same bitrate for both transmitters. Building on these observations, this section explores opportunities to enable SIC to extract the gains, where available.

5.1 Client Pairing

When multiple clients have packets to the AP, we can reduce the upload time by allowing two clients to transmit concurrently. But not all client pairs produce the same gain with SIC. An ideal pair would have a RSS difference appropriate for achieving the same bitrate for both the clients. Therefore, among all the possible pairings, we could choose those that minimize the overall upload time.

Consider a case where 4 clients have a packet each to their common AP. Without SIC, they transmit sequentially as in Fig. 9(a). Each client's transmission rate depends on its proximity to the AP. Fig. 9(a) shows that C_1, C_2, C_3, and C_4 transmit their packets in 1, 2, 4, 8 time units respectively, for a total of 15 units. Note that, these values are not precise, meant for illustration only. With SIC, there are three possible pairings $(C_1|C_2, C_3|C_4)$, $(C_1|C_3, C_2|C_4)$, and $(C_1|C_4, C_2|C_3)$ as in Fig. 9(b), Fig. 9(c), and Fig. 9(d). The corresponding transmission times are 11.5, 12, and 13 respectively. Clearly, appropriate pairing of clients would reduce the overall time needed with SIC to transmit the packets to the AP.

5.2 Power Reduction

SIC achieves higher gain if a client is allowed to adjust its power depending on its partner client. Interestingly, SIC gain can be increased by reducing the power of the weaker client, when the RSSs at the AP of both clients are close. In such a case, the stronger client becomes the bottleneck as it experiences a lower rate than the weaker client. By reducing the power of the weaker client and thus widening the difference in RSSs at the AP, we can increase the rate for the stronger client while decreasing rate for weaker client, and thereby achieve the best completion time for both the packets. For example, by lowering power for C_2 and equalizing the rates of C_1 and C_2, the overall time can be reduced from 11.5 to 11 units, as shown in Fig. 9(e). Thus, dynamic reduction of weaker client's power to equalize the transmission times of both clients can improve the upload performance under SIC.

(a) Serial transmissions (time = 15)

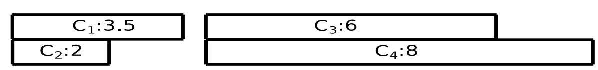

(b) Pairing of $C_1|C_2$ and $C_3|C_4$ (time = 11.5)

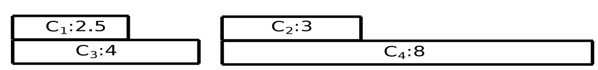

(c) Pairing of $C_1|C_3$ and $C_2|C_4$ (time = 12)

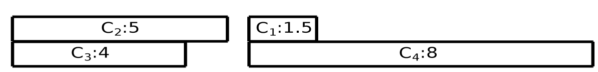

(d) Pairing of $C_1|C_4$ and $C_2|C_3$ (time = 13)

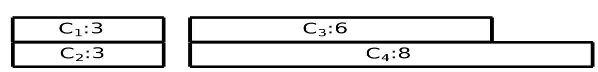

(e) Power reduction by C_2 helps C_1 (time = 11)

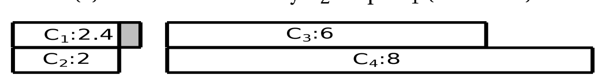

(f) Multirate packetization by C_1 (time = 10.4)

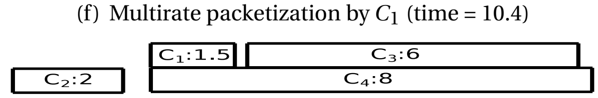

(g) Packing of $(C_1, C_3)|C_4$ (time = 10)

Figure 9: Illustration of client pairing, power control, and multirate packetization with SIC.

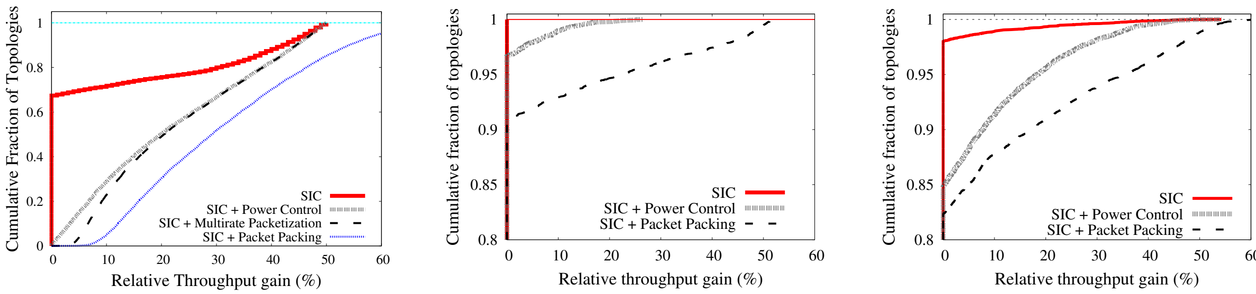

(a) One receiver (b) Two receivers in EWLAN (c) Two receivers in RWLAN

Figure 10: Performance of SIC with power control and packet packing.

5.3 Multirate Packetization

When power control is not an option, multirate packetization [6], where different parts of the packet are transmitted at different rates (Fig. 9(f)), can be effective. The packet from C_1 with the slower rate is sent at a faster rate (optimal for the given SNR) after the completion of the packet from C_2. The total time of completion is about 10.4 as opposed to 11.5 (Fig. 9(b)) without multirate packets. So multirate packetization can complement SIC in scenarios where the difference in RSS values of the clients is not high.

5.4 Packet Packing

Fig. 9(g) shows an alternative approach to power control, i.e., send multiple packets (C_3 and C_1) serially at higher rates before the packet at the lower rate (C_4) finishes. This is difficult today as practical SIC receivers will require some parts of C_3 and C_1 to be in the clear for reliable transmitter-receiver synchronization. Packet C_3 may achieve this by starting before C_4, however, C_1 cannot do the same. Future advancements in SIC may allow such forms of packet packing, providing some gains in favorable settings.

5.5 Comparison

To evaluate the above techniques, we performed Monte Carlo simulations. Fig. 10 shows that in two transmitters to one receiver scenario, the gains with SIC alone are modest (20% of the cases gain over 20%), while the above techniques may offer some benefits (over 20% in 40% of the topologies). In the two-receiver cases, SIC alone has almost no gain and little gain even with these optimizations.

6. DISCUSSION AND CONCLUSION

The gist of this paper is that SIC may not be a promising tool to improve wireless network throughput, except in some restricted scenarios. While this could be unexpected at first glance, it makes sense upon a closer look. Specifically, the very first step in canceling an interference is to decode its bits. Decoding, however, is not only dependent on the RSS of the interfering signal, but also on the bitrate that the interferer is using to communicate to its own receiver. Even if the interference is strong, it may not be decodable if the interferer is also transmitting at a high bitrate. To make it worse, the SNR of the signal of interest should also be sufficiently low to allow for decoding of the interfering signal. Together, these conditions are quite restrictive, especially when any given transmitter is operating at near optimal bitrates.

Our gloomy assessment of the effectiveness of SIC may seem to go against the upbeat reports of throughput gains from recent works based on interference cancellation. Several recently proposed schemes such as ANC [7], ZigZag [8], CSMA/CN [9] and Full-Duplex [10] have successfully applied interference cancellation to demonstrate performance gain. The common thread among all these approaches is that bits of the interfering frame are known in advance. So they *need not be concerned with decoding* but only with modeling and subtracting the interference. Our study in no way contradicts these works but in a sense reinforces them, i.e., interference cancellation should be used where the interference is known through some out-of-band mechanism. When the interference is unknown, and links share a common node, modest gains may be feasible through various forms of coordination. When the interference is unknown and links have no common node, the gains from SIC are probably not worthwhile.

7. REFERENCES

[1] S. Verdu, "Multiuser Detection," in *Cambridge University Press*, 1998.

[2] D. Halperin, T. Anderson, and D. Wetherall, "Taking the Sting out of Carrier Sense: Interference Cancellation for Wireless LANs," in *Mobicom*, 2008.

[3] M. Vutukuru et. al., "Cross-Layer Wireless Bit Rate Adaptation," in *SIGCOMM*, 2009.

[4] S. Sen et. al., "AccuRate: Constellation based Rate Estimation in Wireless Networks," in *NSDI*, 2010.

[5] D. Tse and P. Vishwanathan, "Multiuser Capacity and Opportunistic Communication," in *Fundamentals of Wireless Communication*, 2005.

[6] A. Misra et. al., "Supporting Continuous Mobility through Multi-rate Wireless Packetization," in *HotMobile*, 2008.

[7] S. Katti et. al., "Embracing Wireless Interference:Analog Network Coding," in *SIGCOMM*, 2007.

[8] S. Gollakota et. al., "Zig-Zag Decoding: Combating Hidden Terminals in Wireless Networks," in *SIGCOMM*, 2008.

[9] S. Sen, R. Roy Choudhury, and S. Nelakuditi, "CSMA/CN: Carrier Sense Multiple Access with Collision Notification," in *ACM MobiCom*, 2010.

[10] J. Choi et. al., "Achieving Single Channel, Full Duplex Wireless Communication," in *MobiCom*, 2010.