This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication.

Successive Interference Cancellation: Carving out MAC Layer Opportunities

Souvik Sen, *Student Member, IEEE,* Naveen Santhapuri, *Member, IEEE,* Romit Roy Choudhury, *Member, IEEE,* and Srihari Nelakuditi, *Member, IEEE*

Abstract—Successive interference cancellation (SIC) is a PHY capability that allows a receiver to decode packets that arrive simultaneously. While the technique is well known in communications literature, emerging software radio platforms are making practical experimentation feasible. This motivates us to study the extent of throughput gains possible with SIC from a MAC layer perspective and scenarios where such gains are worth pursuing. We find that contrary to our initial expectation, the gains are not high when the bits of interfering signals are not known a priori to the receiver. Moreover, we observe that the scope for SIC gets squeezed by the advances in bitrate adaptation. In particular, our analysis shows that interfering one-to-one transmissions benefit less from SIC than scenarios with many-to-one transmissions (such as when clients upload data to a common access point). In view of this, we develop an SIC-aware scheduling algorithm that employs client pairing and power reduction to extract the most gains from SIC. We believe that our findings will be useful guidelines for moving forward with SIC-aware protocol research.

Index Terms—Wireless Communication, Interference Cancellation, Collision, Capacity

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 collision in current wireless networks based on IEEE 802.11 standard). 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. Emerging software radio platforms, such as GNU radios, are making practical implementations of SIC feasible [2], [3]. A natural question then is: given SIC capable radios, what are the implications on MAC protocol design? Can SIC be exploited at the MAC layer to improve throughput? What is the scope and what are the limitations?

Inspired by these questions, this paper is an attempt to interpret the PHY layer SIC capabilities from the MAC layer. We limit our focus to the special case of SIC, where only one signal is cancelled from another. We consider simple topological configurations that form the building blocks of larger networks, and systematically study the ideal gains available from SIC. Even among the simple building blocks, we recognize that certain topological patterns are amenable to SIC gains, while others are not. In an attempt to tap into some of these gains, we find that link layer coordination, such as SIC-aware link pairing and power control, are necessary. We verify these observations through theoretical formulations. Guided by these outcomes, we carve out the scenarios in which SIC-aware protocols are worth pursuing. We develop an algorithm for such scenarios, and evaluate its performance. Our key contributions in this paper may be summarized as follows:

(1) We show that SIC-aware MAC protocols offer significant throughput gains in restricted scenarios, mainly for upload in WLANs. Other topologies are not as amenable to SIC, particularly when each transmitter chooses its bitrate independently.

(2) The relative gain from SIC is maximized, when two transmitter's power levels are such that, with SIC, the feasible bitrate is equal for both the transmissions. This can be facilitated by suitable client pairing and power reduction.

(3) We develop a SIC-aware scheduling algorithm for WLANs. We show that such scheduling is equivalent to minimum weight perfect matching that is known to have efficient solutions.

Our observations may appear to be at odds with the high throughput improvements with SIC reported in [4][1]. The root of this discrepancy is in the bitrates used for packet transmissions. Our study assumes that each packet is transmitted at the best feasible rate supported by the channel to its receiver. Our

- *Souvik Sen, Naveen Santhapuri, and Romit Roy Choudhury are with the Departments of Computer Science, and Electrical and Computer Engineering, Duke University, Durham, NC 27708.*
- *Srihari Nelakuditi is with the Department of Computer Science and Engineering, University of South Carolina, Columbia, SC 29208.*

1. Several schemes such as ANC [5], ZigZag [6], CSMA/CN [7] and Full-Duplex [8] have successfully applied interference cancellation to demonstrate performance gain. However, they are applicable when interfering bits are known in advance. This paper focuses on SIC when none of the frames are known a priori to the receiver.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication.

intention is to capture the gains solely due to SIC, thereby isolating it from the gains realizable through ideal bitrate control. One could certainly argue that a practical bitrate adaptation scheme is unlikely to operate at the ideal bitrate at all times and there will always be a slack that SIC can harness. Although true, this slack is fast disappearing with more fine-grain bitrates (4 in 802.11b vs 8 in 802.11g vs 32 in 802.11n) and the recent advances in bitrate adaptation [9], [10], [11]. Moreover, we believe 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.

The subsequent two sections begin with a PHY-centric overview of SIC capacity, and gradually migrate to a MAC layer interpretation, namely throughput. Section 4 examines the efficacy of SIC in different architectures. We show that techniques like client pairing and power reduction help SIC in Section 5. We then develop an SIC-aware scheduling algorithm for upload traffic in WLANs and present trace based evaluation results in Section 6. We discuss the related work in Section 8, and conclude in Section 9.

2 Capacity Gains with SIC

Successive interference cancellation (SIC), as mentioned earlier, is the ability to decode information associated with the individual signals from their combined superimposed signal. It is generally expected that the capacity of a wireless channel can be improved significantly with the aid of SIC. We are interested in characterizing the extent of the gains possible with SIC, and how the gains relate to the relative strengths of the received signals. In this section, we first introduce SIC and then discuss its potential.

2.1 SIC

Let us define *collision* as the simultaneous arrival of two or more packet transmissions at a receiver. With traditional signal extraction, only the packet with the strongest signal can be decoded, treating all the other signals as interference. If the signal of interest is not the strongest signal, it cannot be recovered. On the other hand, SIC facilitates recovery of even a weaker signal of interest. With SIC, the bits of the strongest signal are decoded as before. The original signal corresponding to those bits is then reconstructed and subtracted (*cancelled*) from the combined signal, and the next strongest signal decoded from this *residue*. This can be an iterative process to recover multiple packets and hence it is termed *successive interference cancellation*. This paper focuses on the simpler case of two packets only, i.e., interference cancellation is performed only once. Nevertheless, for convenience, we still refer to it as SIC.

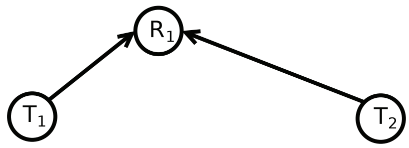

Fig. 1. Topology to illustrate the gains of SIC.

TABLE 1
Notation

B	bandwidth of the channel
N_0	noise in the channel
RSS	received signal strength
T_i	transmitter i
R_j	receiver j
S_j^i	RSS of transmitter T_i at receiver R_j
$\hat{r}_j^i$	optimal bitrate for transmission from T_i to R_j
L	packet length in bits
$C_{-\text{SIC}}$	channel capacity without SIC
$C_{+\text{SIC}}$	channel capacity with SIC
$Z_{-\text{SIC}}$	time to transmit two packets without SIC
$Z_{+\text{SIC}}$	time to transmit two packets with SIC

2.2 Bitrates and SIC Feasibility

The first step in SIC at a receiver of two transmissions is to decode the packet corresponding to the stronger signal correctly. Only then, it can cancel the stronger signal and obtain the packet corresponding to the weaker one. The decodability of the packets depends on their relative signal strengths and transmission bitrates. Let S_1^1 and S_1^2 be the received signal strengths (RSSs) at a receiver R_1 from two transmitters T_1 and T_2 as in 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 first attempts to decode the stronger signal, say S_1^1, treating the weaker signal, say S_1^2, as interference. R_1 can decode S_1^1 only if T_1 transmits *below* a certain bitrate. According to [12], the highest feasible bitrate $\hat{r}_1^1$ for T_1 is

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

If T_1 transmits at a rate higher than $\hat{r}_1^1$, it can not be decoded successfully by R_1, and consequently it can not decode T_2's signal either. Otherwise, R_1 can decode T_1's signal first and then 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

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

SIC is theoretically feasible if T_1 and T_2 transmit at a bitrate no higher than $\hat{r}_1^1$ and $\hat{r}_1^2$. 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.

2.3 Capacity with SIC

Now let us contrast the capacity of a wireless channel without and with SIC given the above scenario of two transmitters and one receiver as in Fig. 1. Without SIC, either T_1 or T_2, but not both, can transmit at a time. So the capacity of the channel without SIC is

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

With SIC, it is possible for a receiver 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 capacity $C_{+\text{SIC}}$ with SIC, as derived in earlier works [13], [14], is

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

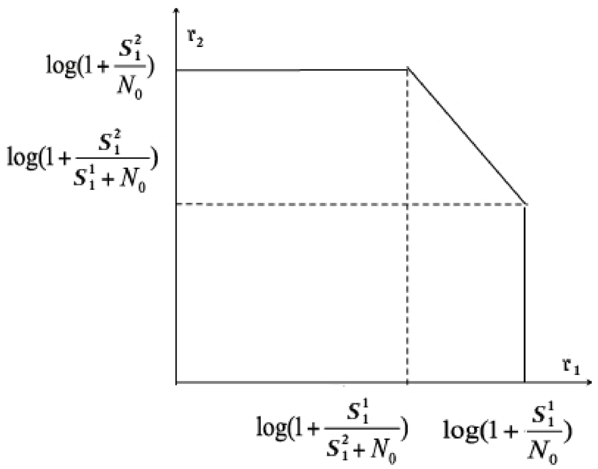

Fig. 2. Aggregate capacity of two transmitters with SIC is higher than the individual capacities.

Note that the capacity with SIC is greater than the individual capacities of either T_1 or T_2 to R_1. Fig. 2 (which is reproduced here from [12]) illustrates the gain with SIC. Intuitively, SIC with two transmitters offers a capacity equivalent to that of a single transmitter whose RSS at the receiver is $(S_1^1 + S_1^2)$. In other words, higher incident power at the receiver from the simultaneous transmitters can yield higher capacity if appropriate techniques are used to decode both the transmissions.

Now consider the relative capacity gain with SIC, $\frac{C_{+\text{SIC}}}{C_{-\text{SIC}}}$. The gains are plotted in Fig. 3 as shades of color with S_1^1 on x-axis and S_1^2 on y-axis. The lighter the shade, the higher the gain with SIC. 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.*

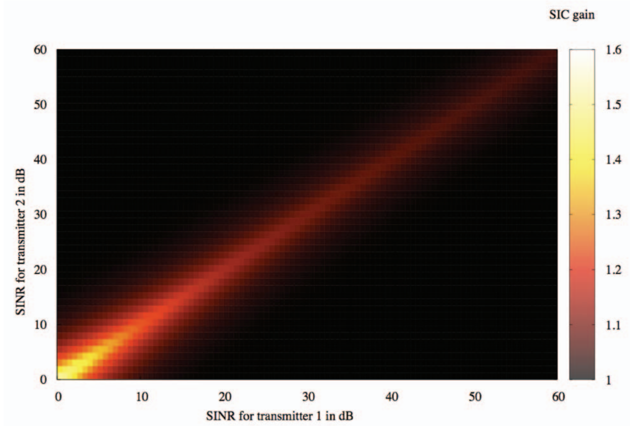

Fig. 3. SIC capacity gains are not high in general but are larger when RSSs are smaller and similar.

3 SIC: MAC Layer Perspective

The key observation from the previous section is that the channel capacity with SIC with two transmitters is always better than the individual capacities of any single transmitter, and the relative gain is maximized when the RSSs are similar. These gains can be realized when each transmitter continuously transmits packets, and one of them transmits at a much lower bitrate than the other. 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 values 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, 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). To achieve channel capacity, the gap in the air-times of packets can be filled by having T_2 transmit a large packet or a train of packets. It may not always be practical to fill the gap by increasing the packet size or transmitting a train of packets – protocol limits on packet sizes prevents the former, while PHY layer synchronization issues make the latter difficult (see Section 5.4). Moreover, it is likely that at an instant of time, each transmitter has a finite number of packets to be sent to its receiver, and it needs to get a fair share of the channel to transmit its packets without inordinate amount of delay. So from a MAC layer protocol perspective, the objective could be to complete the transmission of a pending packet in the shortest time possible.

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.

To obtain an intuition on MAC layer performance of SIC, we study the relative gains with SIC when *each transmitter has one packet* to transmit to its receiver. Since SIC in this paper is concerned with recovering at most two signals, we need to consider only two scenarios: i) two transmissions to the *same* receiver, i.e., both signals are of interest to the receiver; and ii) two transmissions to *different* receivers, i.e., only one signal is of interest to a receiver. We compare the minimum time needed for the two transmissions with and without SIC to determine the scenario where SIC is most beneficial.

3.1 Two Transmitters to the Same Receiver

Consider the scenario similar to Fig. 1 in the previous section, where two transmitters T_1 and T_2 are transmitting one packet of length L bits each to the same receiver R_1 with RSS values S_1^1 and S_1^2. Without SIC, they have to transmit sequentially and the total time (discounting MAC related overheads such as backoff) needed for transmitting the two packets is

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

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

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

The relative gain from SIC, 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. Once again, the axes indicate the RSS values S_1^1 and S_1^2, and the corresponding gain is higher if the shade is lighter. An interesting observation is 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 to decrease again. This behavior can be reasoned as follows. The equation (6) is maximized when denominators of two terms are minimized. Since the two terms over which maximum is taken 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. To summarize, SIC yields the best gain when a receiver can successfully decode two packets (of the same size) transmitted at the same bitrate. 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. Put another way, when the transmission times of two packets *overlap completely*, one of the packets, thanks to SIC, gets a free full ride.

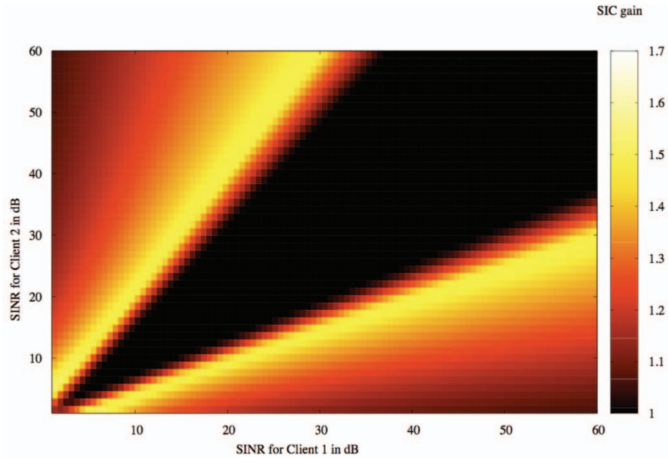

Fig. 4. Two transmitters to the same receiver: SIC gains most when RSSs are such that the resulting bitrates are the same for both transmissions.

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.

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

(Fig. 5b) $S_1^1 > S_1^2$ 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

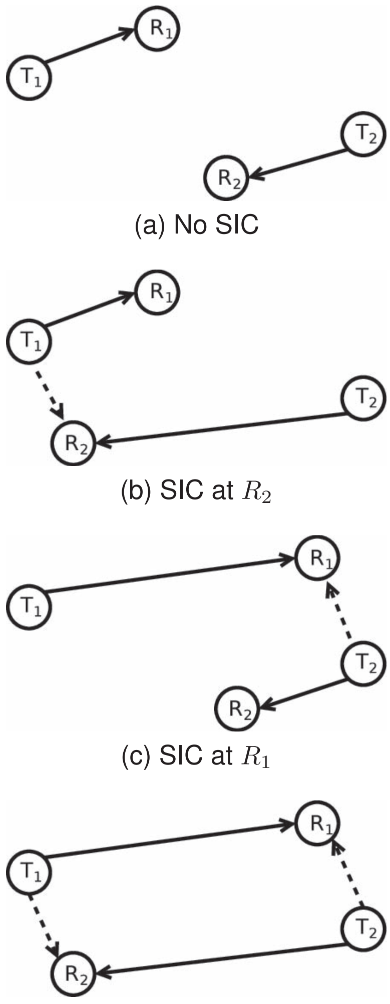

(a) No SIC

(b) SIC at R_2

(c) SIC at R_1

(d) SIC at R_1 and R_2

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

concurrent transmissions. The time to transmit two packets in this case with and without SIC are,

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

$$Z_{-\text{SIC}} = \frac{L}{B\log_2(1+\frac{S_1^1}{N_0})} + \frac{L}{B\log_2(1+\frac{S_2^2}{N_0})} \tag{8}$$

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

(Fig. 5d) $S_1^1 < S_1^2$ 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, $Z_{+\text{SIC}}$ would be

$$Z_{+\text{SIC}} = \max\left(\frac{L}{B\log_2(1+\frac{S_1^1}{N_0})}, \frac{L}{B\log_2(1+\frac{S_2^2}{N_0})}\right) \tag{9}$$

and $Z_{-\text{SIC}}$ is the sum of the two terms in Equation (9). Although $Z_{+\text{SIC}} < Z_{-\text{SIC}}$ (raising hopes for SIC), the requisite topological conditions like Fig. 5d are unfortunately not common in real life.

For example, in Fig. 5c, imagine that the signal to noise ratio (SNR) of T_1's transmission at R_1 is 40dB, and similarly SNR of $T_2 \rightarrow R_1$ and $T_2 \rightarrow R2$ are 50dB and 30dB respectively. Let r_{10}, r_{30}, and r_{40} be the highest bitrates feasible respectively with 10dB, 30dB, and 40dB. Now suppose T_1 transmits at rate r_{40} and T_2 transmits at rate r_{10}. Then, both transmissions can be decoded at R_1 since SINR of the stronger signal from T_2 is 10dB (50dB−40dB) and SNR of the weaker signal from T_1 is 40dB. On the other hand, if T_2 transmits (say to R_2) at rate r_{30}, then R_1 can not decode both the transmissions — it can not decode the stronger signal from T_2 since SINR of 10dB can not support bitrate of r_{30}, and it can not decode the weaker signal from T_1 without cancelling out the stronger signal. These examples highlight the role of transmission bitrates in determining the applicability of SIC in practice.

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 $\frac{Z_{-\text{SIC}}}{Z_{+\text{SIC}}}$. 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 with different receivers.

4 SIC IN DIFFERENT ARCHITECTURES

It may seem that we only considered toy topologies to make sweeping conclusions. We argue that the scenarios we discussed form building blocks in many wireless network settings and therefore are broadly applicable. Here we list different wireless architectures and discuss how SIC benefits them.

4.1 Enterprise Wireless LANs

The APs under enterprise wireless LAN (EWLAN) environments such as corporate campuses are connected through a wired backbone as shown in 7a. They have an additional wireless interface through which they communicate to wireless clients. These

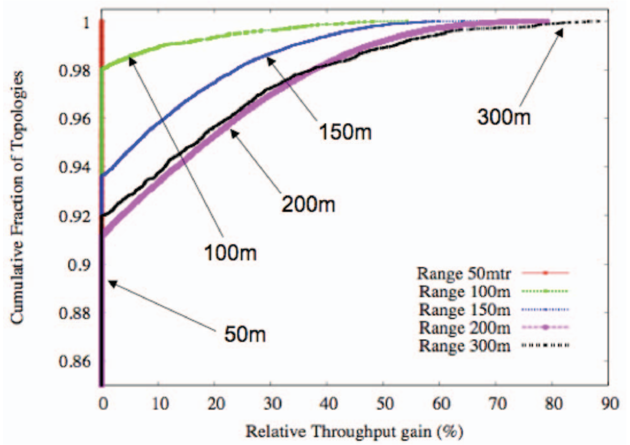

Fig. 6. Two transmissions to different receivers: no gain from SIC in 90% of the cases.

APs coordinate the transmission of download and upload traffic to achieve better overall performance. We now probe whether and where SIC is useful in EWLANs.

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 one AP, say AP_1. This is no different from the two transmitter and one receiver scenario discussed in the 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 high speed wired network, packets can be delivered to C_2 via either of the APs. Using SIC, it is conceivable to achieve 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, and so the total time needed for transmitting one packet from each AP is same as (6). Without SIC, the two packets have to be transmitted sequentially as in upload scenario. However, here 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

$$\frac{2L}{\max(B\log_2(1+\frac{S_2^1}{N_0}), B\log_2(1+\frac{S_2^2}{N_0}))} \quad (10)$$

The relative gain from SIC then is (10)/(6), which is plotted in Fig. 8. We can observe that SIC offers modest gains when the RSSs are such that one is roughly square (twice in terms of dB) of the other. But more importantly, the overall gains with SIC are quite limited in this download scenario.

Upload Traffic: Two Clients to Two APs

We have seen that SIC is beneficial for upload traffic from two clients to one AP. Could SIC also help in scenarios of two clients transmitting to two different 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 EWLAN settings as shown in 7a. SIC is needed in this setting 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 passing the packet to any of the APs, transmission to the closest AP is obviously a better alternative in this setting. Then, each client's signal will be stronger at its respective AP than the other client's, and hence SIC is not needed to receive them.

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, similar to the above case, SIC figures in this setting only when AP_2 is delivering a packet to C_1 or C_2 in concurrence with AP_1's transmission of a packet to C_3 or C_4. Given that packets can be delivered to clients through any of the APs, there is little benefit in choosing a farther AP (with lower RSS) and applying SIC to receive the packets.

4.2 Residential Wireless LANs

Residential wireless LANs in adjacent apartments could be like in Fig. 7b where a (WPA protected) AP serves clients in that home. Unlike in enterprise wireless LANs, in residential wireless LANs, 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 though C_2 is closer to AP_2. Strangely, this restriction provides some opportunities for SIC. In this example, if C_2 performs SIC, $AP_1 \rightarrow C_2$ transmission can be concurrent with $AP_2 \rightarrow C_4$. The client C_2 can decode stronger interfering $AP_2 \rightarrow C_4$ packet, cancel the interference, extract its signal of interest from AP_1. 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 can be 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, there are 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, residential wireless LANs offer some opportunities for SIC in large apartment complexes crowded with APs. Of course, upload traffic from two clients to their AP benefits from SIC as explained in Section 3.1.

4.3 Multihop Mesh Networks

Wireless mesh networks are peer to peer networks of APs, one of which can be a gateway to the Internet.

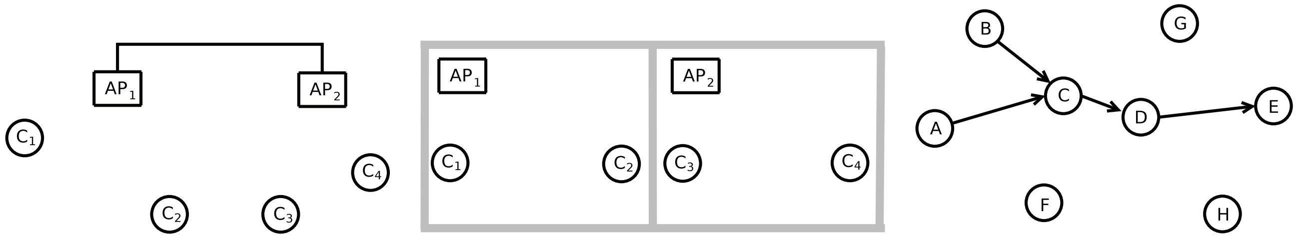

(a) Enterprise wireless LAN (b) Residential wireless LAN (c) Wireless Mesh Network

Fig. 7. Example topologies of different wireless architectures to illustrate the throughput gains from SIC.

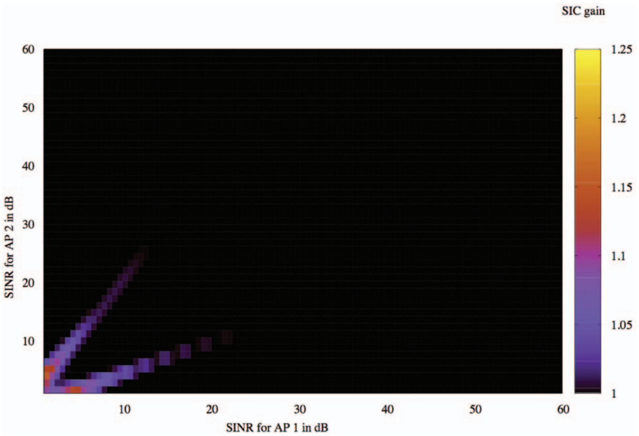

Fig. 8. Download traffic from two APs to one client in an EWLAN: very little benefit from SIC.

Traffic from source clients can be routed across multiple APs to either reach the gateway, or to another AP that is connected to the destination client. Fig. 7c shows an example set of APs connected through a wireless backbone network. Being a more generic setting than the other two, the requisite conditions for SIC exist, though the opportunities and the resulting gains depend on traffic patterns and routing decisions. For example, when both A and B have a packet to transmit to C, depending on their relative proximity to C, SIC would offer significant gain as shown in Section 3.1. Now suppose packets from A to E are routed along the path $A \rightarrow C \rightarrow D \rightarrow E$. Notice that this path consists of a long-hop followed by short-hop and then again a long-hop. It is a perfect recipe for SIC at C. Both $A \rightarrow C$ and $D \rightarrow E$ transmissions can be concurrent, addressing the so called *self interference* problem. However, the long-hop transmissions become the bottleneck as they have to be at a lower bitrate, bringing down the end-to-end throughput. On the other hand, if long-hops are made shorter (A and E move closer to C and D), C may not be able to decode the high bitrate interfering packet from D, and hence can not extract its packet from A with SIC. When the transmitter-receiver pairs belong to different traffic flows and each flow is routed independently, then their interaction would be similar to the case discussed in Section 3.2.

5 Techniques for Empowering 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

It is common for a single AP to serve many clients in a WLAN environment. When multiple clients have packets to the AP, we can reduce the upload time by allowing two clients at a time to transmit concurrently. But not all client pairs produce the same gain with SIC as it depends on the difference in RSSs of the clients at the AP. As mentioned before, an ideal client 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.

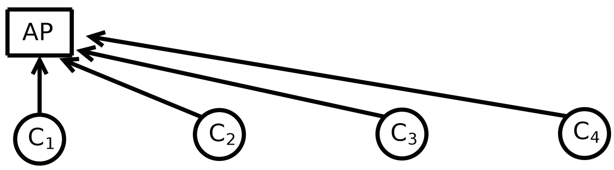

Fig. 9. Topology for illustration of client pairing.

For example, consider the upload setting shown in Fig. 9 where 4 clients have a packet each to send to the AP. Without SIC, they transmit sequentially as shown in Fig. 10a. The bitrate of each client's transmission depends on its proximity to the AP. Fig. 10a 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 of time. Note that, these values are not precise and are meant for illustration only to contrast different approaches and their relative merits. 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 shown in Fig. 10b, Fig. 10c, and Fig. 10d. The corresponding transmission times are

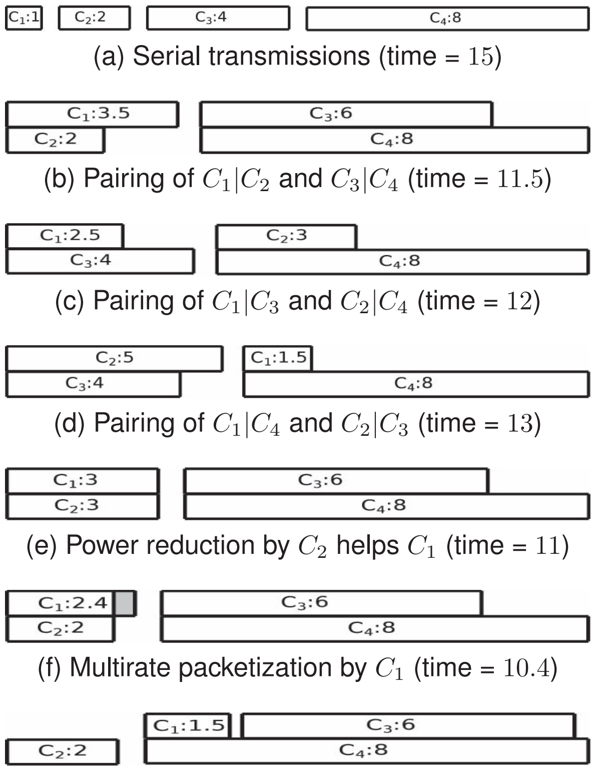

(a) Serial transmissions (time = 15)

(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)

Fig. 10. Illustration of client pairing, power control, packet packing, and multi-rate packetization with SIC.

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

Two clients are considered a perfect pair under SIC if they achieve the same bit rate for their transmissions to the AP. The above approach finds as close to perfect matches as possible among the clients with pending packets. It yields the best possible outcome when the power level used by a client has to be static and can not be controlled. On the other hand, there is more potential for gain with SIC if a client is allowed to dynamically adjust its power level depending on its partner client. Interestingly, gain with SIC 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 bitrate 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 bitrate for the stronger client while decreasing bitrate for weaker client, and thereby achieve the best completion time for both the packets. For example, by lowering power by C_2 and equalizing the bitrates of C_1 and C_2, the overall time can be reduced from 11.5 to 11 units, as shown in Fig. 10e. Thus, dynamic reduction of weaker client's power to equalize the transmission times of both clients is an option to improve the upload performance under SIC.

5.3 Multirate Packetization

Power reduction by a weaker client is a way to pull up the bitrate of stronger client to equalize their transmission times. When power control is not an option, multirate packetization proposed in [15] can be an effective tool to utilize SIC. With multirate packetization, different parts of the packet are transmitted at different bitrates. This permits the stronger client with an effective lower bitrate during the simultaneous transmission for SIC, to transmit at a higher bitrate possible in the absence of interference, once the weaker client with faster bitrate completes the transmission. This is illustrated in Fig. 10f where the packet from C_1 with the slower bitrate is sent at a faster bitrate (maximum possible for the given SNR) after the completion of the packet with higher bitrate. The total time of completion is about 10.4 as opposed to the 11.5 in Fig. 10b 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

Power reduction is suitable when the power levels are similar and the stronger client is the bottleneck. On the other hand, if the weaker client has lower bitrate, power reduction won't help. Increasing the power of the weaker client is not desirable as it will amplify the overall channel interference and may cause a cascading effect. Another alternative to power control is to send a single large packet or multiple packets serially at higher bitrate before the packet at the lower bitrate finishes. This kind of transmission will depend heavily on the traffic patterns and might be useful in some scenarios. Packet packing is also applicable even when the stronger client has lower bitrate. Then the weaker client could send multiple packets before the stronger client finishes transmitting one packet. A more generic version of packet packing would be to allow multiple higher bitrate transmissions from different clients in parallel with a single lower bitrate transmission. An illustration of such packet packing is shown in Fig. 10g, i.e., send multiple packets (C_1 and C_3) serially at higher rates before the packet at the lower rate (C_4) finishes. Such a packing of packets from different clients is difficult today as practical SIC receivers will require some parts of C_1 and C_3 to be in the clear for reliable transmitter-receiver synchronization. Packet C_1 can achieve this by starting before C_4, however, C_3 cannot do the same. Future advancements in SIC may allow such forms of packet packing, providing some gains in favorable settings.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication.

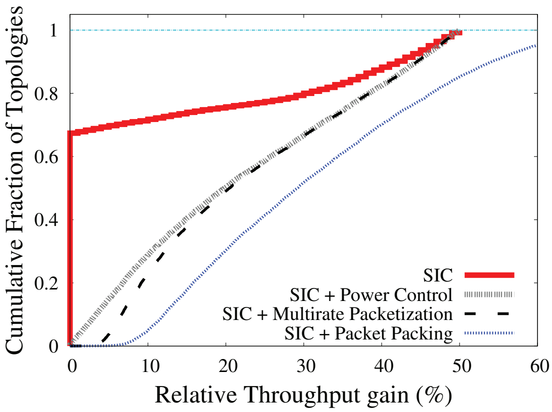

(a) One receiver

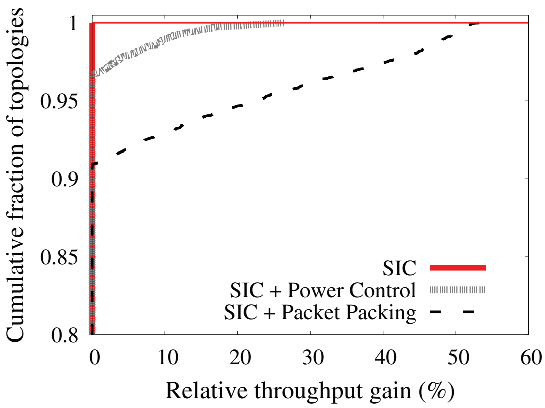

(b) Two receivers in EWLAN

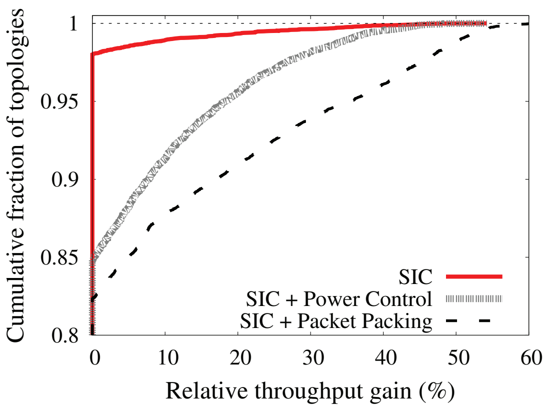

(c) Two receivers in RWLAN

Fig. 11. Throughput gain of SIC when coupled with power control and packet packing.

5.5 Comparison

To evaluate the utility of the above mechanisms, we have performed Monte Carlo simulations. We evaluate SIC, SIC with power control, multi-rate packetization, and packet packing. Multi-rate packetization is not possible in a two transmitter, two receiver scenario since the transmitter of the receiver employing SIC will operate at a higher rate than the other link. We discuss transmitter pairing in the next section as it can be augmented with any of these techniques. Fig. 11 shows that in a two transmitter-one receiver scenario, the gains with SIC alone are modest (20% of the cases gain over 20%) but there are significant gains (over 20% in 40% of the topologies) by using one of the above mechanisms. In the two-receiver cases, SIC alone has almost no gain and very little gains even with these optimizations. This shows that the two receiver scenarios are not amenable to SIC (very few opportunities and very few gains in most cases).

6 SIC-Aware Scheduling

Examination of SIC shows that the performance gains are worthwhile in the scenarios of upload traffic in wireless LANs. We intend to tap into this gain by developing an SIC-aware scheduling algorithm. This section presents the algorithm. We evaluate it on wireless traces collected from our university buildings.

Section 3 shows that MAC layer throughput under SIC is maximized when the bitrates achieved by the two transmitters are equal. Equal bitrates cause the packets to end at almost the same time, eliminating channel wastage (assuming equal sized packets). This calls for carefully choosing backlogged client pairs and controlling the power of transmissions. Specifically, for a given client pair, the transmit power can be controlled to achieve nearly equal bit rates. Further, client pairs can be carefully selected such that all the backlogged clients can be serviced in the minimum time. We present an SIC aware client scheduling algorithm to exploit these opportunities. Our problem statement is as follows:

SIC-Aware Scheduling: *Given a set of backlogged clients C, and their respective maximum bitrates to the AP, find all pairs of clients and their associated transmit powers, such that the total time to upload all the backlogged traffic is minimum.*

We approach the problem by reducing SIC-aware scheduling to Edmond's minimum weight perfect matching algorithm [16]. Edmond's algorithm solves the following problem:

Edmond's Minimum Weight Perfect Matching: *A perfect matching in a graph G is a subset of edges such that each node in G is met by exactly one edge in the subset. Given a real weight c_e for each edge e of G, the minimum weight perfect-matching problem is to find a perfect matching M of minimum weight P.*

We show the reduction next, and briefly sketch the Edmond's solution. The final result is the SIC-aware client schedule.

The reduction begins by forming a graph $G = (V, E)$, where V is the set of all clients in the WLAN. The set E is initially empty. Each client-pair ij is chosen and the minimum time for their joint transmission

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication.

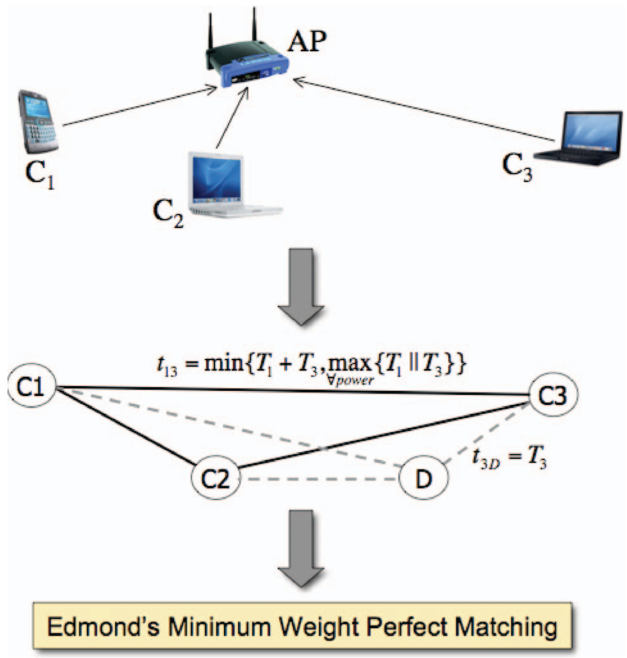

Fig. 12. Translating SIC-aware scheduling into Edmond's minimum weight perfect matching algorithm.

is computed – denote this by t_{ij}. This computation considers the minimum of: i) time for serialized transmissions, and ii) the minimum time for joint transmissions using SIC. The cost of a joint transmission will depend on whether power control is enabled. An edge is added between clients i and j with this cost t_{ij}. To account for odd number of clients (in which case some client will have to transmit a packet alone), we add a dummy client D in the graph. An edge is added from each client i to D. The cost of the edge iD is the transmit time when client i transmits individually. This is computed based on the maximum bitrate client i can achieve, given its relative position to the AP and a maximum transmit power level. The resulting graph G is then fed to the Edmond's algorithm. Figure 12 illustrates the reduction.

Edmond's algorithm is one solution to the minimum weight perfect matching problem. The basic idea is to start with empty (or any) matching solutions, and improve it with repeated M-augmenting paths until the improvement stops [16]. In a graph with n vertices and m edges, the algorithm runs for $O(n)$ iterations. The recursion depth is $O(n)$; and the complexity per recursive call is $O(m)$. Hence, the algorithm runs in $O(n^2m)$ time. We omit the details in the interest of space.

The output of the algorithm is a set of edges, ξ, such that any vertex in G has exactly one edge in ξ incident upon it, and the sum of weights of edges in ξ is minimum. From ξ, we can easily retrieve the pairs of clients that must transmit simultaneously. This is done by pairing the clients that correspond to the vertices of each edge in ξ. When the number of clients in G are odd (i.e., a dummy vertex is introduced), only one edge in ξ will be incident upon the dummy vertex. The dummy's counterpart vertex is chosen to be the one that performs the individual transmission. The AP can now schedule the client-pairs and the individual transmission in any arbitrary order.

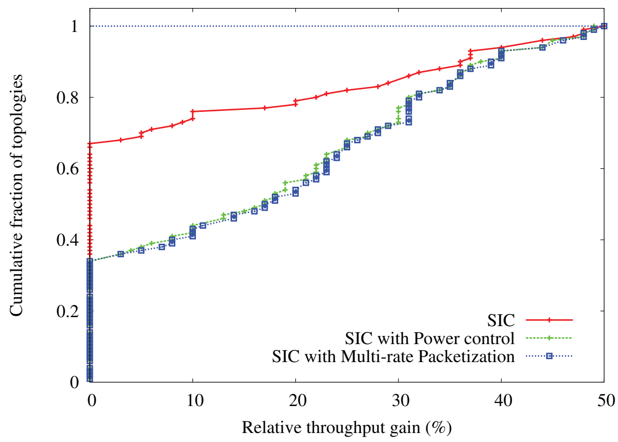

Fig. 13. Trace based evaluation of multiple instances of two transmitters and one common receiver scenario with the link pairing algorithm.

7 Trace-Based Evaluation

The observations about SIC to this point are mostly theoretical. This section focuses on verifying these observations and validating SIC-aware scheduling through trace based evaluation of both upload and download traffic scenarios in a practical setting.

Upload traffic scenario

We collected real world 802.11g link RSSI traces from a busy building in Duke University over 2 weeks. For this, we co-located Soekris boxes with existing access points (APs) in the building. From the traces, we parsed out topology snapshots (every 15 minutes) that provide sets of wireless clients associated to each AP. Using the per-client RSSI at the AP, we quantified the achievable gains with SIC-aware link-pairing. Fig. 13 presents the results, confirming the *prospective gains in pairing wireless client transmissions in real life WLANs.* Evident from the graph, the relatives gains from SIC are enhanced when used in conjunction with power control or multi-rate packetization. The trends are similar to the results shown in Fig. 11a.

Download traffic scenario

We collected separate traces to evaluate throughput gain with SIC in download traffic scenario of two transmitters transmitting packets to different receivers. To collect the traces, we co-located 5 Soekris boxes with existing APs in our department building. We randomly chose 100 locations in adjacent classrooms and offices as client locations. For each client we recorded the SNR from all the 5 APs. We also experimentally found the best bitrate supported by the channel from each AP to this client — the highest 802.11g bitrate at which 90% of packets are received successfully. Similarly, we also found the bitrate supported to a client from an AP under interference from other APs. To calculate this, we turn off carrier sensing on the Soekris APs and schedule two APs to transmit simultaneously to a client. For the stronger AP, we calculate the highest bitrate for which 90% of the packets are successful under interference.

From the traces collected as described above, we evaluate the performance of SIC both when any arbitrary bitrate is allowed and also when only a few discrete bitrates are permitted as in 802.11. First, assuming arbitrary bitrates, we use the recorded SNR values to determine the optimal bitrate between an AP and a client. We compute the relative throughput gain with SIC for each scenario of two transmitter-receiver (AP-client) pairs using the equations derived in Sections 3.2 and 4. Fig. 14(a) shows the CDF of gain from SIC with and without packet packing. It is evident that even with packing SIC offers limited gains. These results are similar to that of simulation presented in Figure 11.

Discrete bitrates

A key assumption in this paper is that each packet is transmitted at the best feasible rate supported by the channel with the intention of understanding the stand-alone benefits of SIC. However, a practical bitrate adaptation scheme is unlikely to operate at the ideal bitrate and instead select one from the few discrete bitrates permitted by the underlying protocol standard such as 802.11. It is expected that SIC can harness the slack and perform better when packets are transmitted at suboptimal bitrates. Therefore, to assess the extent of gain with SIC in case of discrete bitrates, we use the above traces where 802.11g bitrates are determined experimentally.

To compute the relative gain due to SIC with the discrete bitrates, we replace the logarithmic terms in the expressions presented in Section 3.2 with the actual bitrates observed in experiments. Again, we compute the SIC gain for each scenario of two AP-client pairs. The resulting CDF of gain from SIC with and without packet packing is shown in Fig. 14(b). As expected, the performance of SIC improves under discrete bitrates. Still, without packet packing the gain from SIC is not quite significant. However, with packet packing, SIC offers more than 20% gain in 40% scenarios of two AP-client pairs. These results indicate that SIC is beneficial under discrete bitrates provided packing of packets is possible and practical.

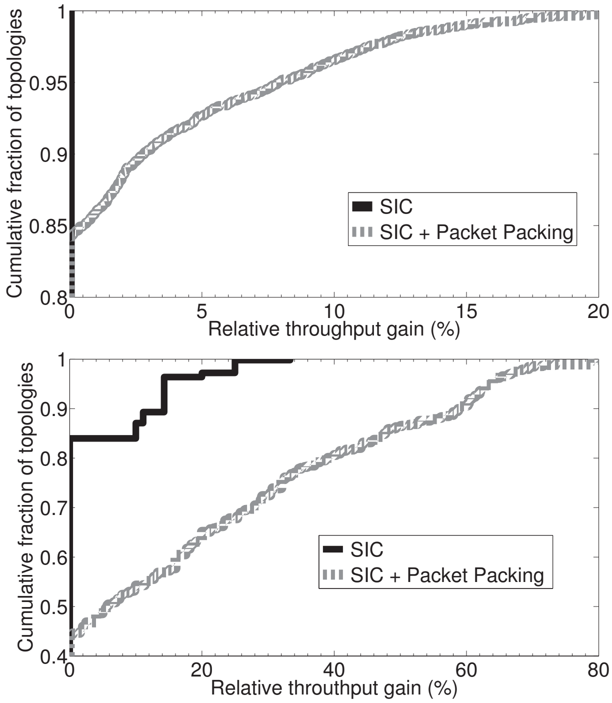

Fig. 14. Trace based evaluation of two transmitter-receiver pairs: (a) allowing any aribitrary bitrate (b) using only discrete bitrates allowed by 802.11g.

8 Related Work

SIC is a form of multiuser detection mechanism [1]. Interference cancellation has been deployed in practical CDMA systems which operate at sub-optimal rates. A large body of literature exists for SIC in the context of cellular networks [17], [18] where the power disparity between receptions is high.

One of the earliest works on the capacities of multiuser channels is [19] which derives achievable rates for each of the broadcast channels in a multi-receiver compound channel. The capacity regions for SIC uplink and downlink scenarios are summarized in [12], which we used extensively for this work. The SIC capacity improvement for ad hoc networks is shown in [13]. The authors show that imperfections in interference cancellation will sharply cut down SIC's usefulness.

Toumpis and Goldsmith showed [14] that the capacity of cellular networks increases significantly with SIC and with power control (if transmission rates are not variable). Most of these studies deal with the physical layer capacity under various models. Our work explores the relation between physical layer capacity and MAC layer throughput.

The practical issues with implementing SIC are given in [20] and [21] with the primary recommendation of power control integration. In [22], the authors show that increasing the power for higher rate CDMA systems in multi-rate systems, is better than adding more lower rate parallel channels. In this work, we show how power control can be used to improve the gains in multi-rate SIC systems. A good summary of existing work on interference cancellation for cellular systems can be found in [23].

Several works have explored the possibilities for multiuser access with power control, coding and rate control [17], [18]. [24] proposed a power control mechanism to achieve high capacity even if the channel estimation was not perfect. However, this was specific to CDMA cellular networks and is not applicable in the current context. Recenly, superposition coding has been explored in [25] to improve throughputs of networks operating with suboptimal rate links. The authors combine superposition coding with network coding to improve the throughput of links with significant power disparity.

In the context of wireless networks, SIC has recently become prominent due to the advances in software defined radio (SDR) [3]. Several works [5], [25], [26] utilized the flexibility of SDR for improving link layer performance. There have been some works exploring multi-rate SIC [27] using matched filter approach. Receivers for interference cancellation using GNU radio/USRP platform have been implemented to deal with exposed [4] and hidden [6] terminal problems. In [28], the authors establish the functionality which must be close to the radio (to reduce latencies) in order to implement efficient MAC protocols on SDR similar to GNU Radio/USRP testbed. Our work tries to connect the physical layer research with the MAC layer research.

Many recent works are utilizing known data for interference cancellation [8], [29], [30]. The data of a signal that has to be cancelled can be known in advance in some special scenarios: through a wired back bone [29], by virtue of transmitting that packet earlier (in the case of self interfereing flows in multi hop wireless networks), multicast retransmissions [31] etc. When the data is known in advance, the wireless signal can be remodeled and cancelled to obtain the unknown signal in some cases. [32] on the other hand, uses multiple receptions on different antennas to solve for the multiple received but overlapping signals. Our work deals with the more general problem of unknown interference cancellation and analyzes the MAC layer improvements for this general case.

The closest to our work though, is [4] which has the first implementation of an SIC receiver in the context of CSMA networks and discusses MAC layer issues like ACKs, buffering/latency and carrier sensing with SIC. The focus of our work is to study the implications of SIC for MAC protocol design, particularly when multiple bitrates are used.

9 Discussion and Conclusion

The gist of this paper is that the advances in rate adaptation limit the scope of possible gains from SIC. 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. The SNR of the signal of interest should also be sufficiently low to allow for decoding of the interfering signal. Furthermore, after decoding the stronger signal, recovering the weaker signal of interest by subtracting the stronger one can also be practically challenging. A SIC receiver needs to model the stronger signal accurately which entails correct channel estimation along with other parameters such as frequency offset and sampling offset. This is challenging especially when the two signals overlap considerably. Moreover, if the stronger signal is significantly stronger than the weaker (which facilitates decoding of the stronger signal), due to ADC saturation issues, recovering the weaker signal becomes difficult. 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 [5], ZigZag [6], CSMA/CN [7] and Full-Duplex [8] 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, gains are feasible through various forms of coordination. When the interference is unknown and links have no common node, complex mechanisms like packet packing are necessary to derive the gains with SIC.

References

[1] S. Verdu, "Multiuser Detection," in *Cambridge University Press*, 1998.
[2] E. Research, "Universal Software Radio Peripheral," *http://www.ettus.com*.
[3] E. Blossom, "GNU Radio Project," *http://gnuradio.org/trac*.
[4] D. Halperin, T. Anderson, and D. Wetherall, "Taking the Sting out of Carrier Sense: Interference Cancellation for Wireless LANs," in *Mobicom*, 2008.

[5] S. Katti, S. Gollakota, and D. Katabi, "Embracing Wireless Interference:Analog Network Coding," in *SIGCOMM*, 2007.
[6] S. Gollakota and D. Katabi, "Zig-Zag Decoding: Combating Hidden Terminals in Wireless Networks," in *SIGCOMM*, 2008.
[7] S. Sen, R. R. Choudhury, and S. Nelakuditi, "CSMA/CN: Carrier Sense Multiple Access with Collision Notification," in *ACM MobiCom*, 2010.
[8] J. I. Choi, M. Jain, P. Levis, and S. Katti, "Achieving Single Channel, Full Duplex Wireless Communication," in *MobiCom*, 2010.
[9] M. Vutukuru, H. Balakrishnan, and K. Jamieson, "Cross-Layer Wireless Bit Rate Adaptation," in *SIGCOMM*, 2009.
[10] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi, "AccuRate: Constellation based Rate Estimation in Wireless Networks," in *NSDI*, 2010.
[11] A. Gudipati and S. Katti, "Automatic Rate Adaptation," in *HotNets*, 2010.
[12] D. Tse and P. Vishwanathan, "Multiuser Capacity and Opportunistic Communication," in *Fundamentals of Wireless Communication*, 2005.
[13] S. W. et. al, "Transmission Capacity of Wireless Ad Hoc Networks with SIC," *IEEE Trans. on Information Theory*, Aug. 2007.
[14] S. Toumpis and A. Goldsmith, "Capacity Regions for Wireless Ad Hoc Networks," *IEEE Trans. on Wireless Communications*, vol. 2(4), Jul. 2003.
[15] A. Misra, S. Rayanchu, D. Agrawal, and S. Banerjee, "Supporting Continuous Mobility through Multi-rate Wireless Packetization," in *HotMobile*, 2008.
[16] J. Edmonds, "Paths, trees, and flowers," *Canadian Journal of Mathematics*, vol. COM-24), 1965.
[17] X. W. et. al, "Iterative(turbo) Soft Interference Cancellation and Decoding for Coded CDMA," in *IEEE Trans. on Communications*, 1999.
[18] A. Agrawal, J. G. A. andJ. M. Cioffi, and T. Meng, "Iterative Power Control for Imperfect Successive Interference Cancellation," *IEEE Trans. on Wireless Communications*, vol. 4(3), 2005.
[19] T. Cover, "Broadcast Channels," in *IEEE Transactions on Information Theory*, 1972.
[20] I. Seskar, T. Kolding, and J. Holtzman, "Implementation aspects for successive interference cancellation in DS/CDMA systems," *Wireless Networks*, vol. 4, no. 6, pp. 447–452, 1998.
[21] K. Pedersen, T. Kolding, I. Seskar, and J. Holtzman, "Practical implementation of successive interference cancellation in DS/CDMA systems," in *Universal Personal Communications, 1996. Record., 1996 5th IEEE International Conference on*, vol. 1. IEEE, 2002, pp. 321–325.
[22] A. Johansson and A. Svensson, "Successive interference cancellation in multiple data rate DS/CDMA systems," in *Vehicular Technology Conference, 1995 IEEE 45th*, vol. 2. IEEE, 2002, pp. 704–708.
[23] J. Andrews, "Interference cancellation for cellular systems: A contemporary overview," *Wireless Communications, IEEE*, vol. 12, no. 2, pp. 19–29, 2005.
[24] J. G. Andrews and T. H. Meng, "Optimum Power Control for Successive Interference Cancellation With Imperfect Channel Estimation," in *IEEE Transaction on Wireless Communications*, 2003.
[25] R. Alimi, L. E. Li, R. Ramjee, H. Viswanathan, and Y. R. Yang, "iPack: In-Network Packet Mixing for High Throughput Wireless Mesh Networks," in *Infocom*, 2008.
[26] G. Woo, P. Kheradpour, D. Shen, and D. Katabi, "Beyond the bits: Cooperative packet recovery using physical layer information," in *Proc. ACM Mobicom*, 2007.
[27] I. B. et. al, "Multirate SIC receiver for UMTS," in *Electronics Letters*, 2003.
[28] G. Nychis, T. Hottelier, Z. Yang, S. Seshan, and P. SteenKiste, "Enabling MAC Protocol Implementations on Software-Defined Radios," in *6th USENIX NSDI*, 2009.
[29] S. Gollakota, S. Perli, and D. Katabi, "Interference alignment and cancellation," in *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*. ACM, 2009, pp. 159–170.
[30] B. Radunovic, D. Gunawardena, A. Proutiere, N. Singh, V. Balan, and P. Key, "Efficiency and fairness in distributed wireless networks through self-interference cancellation and scheduling," in *IEEE WiMesh*, 2010.
[31] C. Qin, N. Santhapuri, S. Sen, and S. Nelakuditi, "Known Interference Cancellation: Resolving Collisions due to Repeated Transmissions," in *IEEE WiMesh*, 2010.
[32] K. Tan, H. Liu, J. Fang, W. Wang, J. Zhang, M. Chen, and G. Voelker, "SAM: enabling practical spatial multiple access in wireless LAN," in *Proceedings of the 15th annual international conference on Mobile computing and networking*. ACM, 2009, pp. 49–60.

Souvik Sen Souvik Sen received the B.E. degree in Computer Science from Bengal Engineering and Science University, in 2007. He received the M.S. degree in Computer Science from Duke University in 2011. He is currently pursuing his Ph.D. degree in computer science at Duke University. His research interests include wireless networking and mobile computing.

Naveen Santhapuri Naveen Santhapuri received his B.Tech degree from Andhra University, India (2002). He did M.S (2005) in Wireless Network Security and Ph.D (2009) in Wireless Networks at the University of South Carolina. Naveen worked as a Post Doctoral associate at Duke University during 2009-2010, doing research in Wireless Networks and Mobile Computing. He is currently working as a Software Developer in New York city.

Romit Roy Choudhury Romit Roy Choudhury is an Associate Professor of ECE and CS at Duke University. He joined Duke in Fall 2006, after completing his PhD from UIUC. He received the NSF CAREER Award in January 2008. His research interests are in wireless protocol design mainly at the PHY/MAC layer, and in mobile social computing at the application layer.

Srihari Nelakuditi Srihari Nelakuditi received his Ph.D. from University of Minnesota, Minneapolis in Computer Science and Engineering. He is currently an Associate Professor at University of South Carolina, Columbia. Earlier, he worked as a software design engineer at Texas Instruments, Bangalore. Srihari Nelakuditi is a recipient of the NSF CAREER Award in 2005. His research interests are in resilient routing, wireless networking, and mobile computing.