BASIC: Backbone-Assisted Successive Interference Cancellation

Wenjie Zhou[†], Tanmoy Das[†], Lu Chen[†], Kannan Srinivasan, Prasun Sinha
 Department of Computer Science and Engineering
 The Ohio State University, Columbus, OH 43210
 {zhouwe, kannan, prasun}@cse.ohio-state.edu, {das.167, chen.4728}@osu.edu
[†]Co-primary authors

ABSTRACT

To meet the growing demand for wireless data, it is time to move away from the age-old paradigm of prohibiting interfering nodes from transmissions. Instead, through proactive management of interference among multiple colliding packets, we can design high throughput wireless systems. This is well explored in the Information Theory community and there are also a few implementational efforts that have been recently reported. The existing solutions are non-trivial to use in real systems as they require either tight time/frequency synchronization or exchange of data between transmitters prior to the transmissions. These requirements are hard to meet in practice especially for uplink transmissions. This paper proposes BASIC, a lightweight multi-user uplink transmission strategy that does not require tight synchronization or exchange of samples among nodes, which makes it an attractive alternative compared to its counterparts. BASIC exploits receiver diversity by controlling the data rates of the clients. A novel greedy algorithm is proposed for data rate selection. We implement BASIC on a software-defined radio platform. Our experiments on a real testbed show that BASIC outperforms TDMA by 48% in terms of overall throughput. Our trace-driven simulations show up to 4.8× gain in throughput with similar flow fairness.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication

Keywords

Wireless Enterprise Networks, Interference Cancellation

1. INTRODUCTION

To meet the rapidly increasing demand for wireless capacity, we need to go beyond traditional strategies that prohibit

interfering transmissions from being simultaneously active. When multiple interfering transmissions are simultaneously active, proactive management of interference becomes essential for successful decoding of these packets. Transmission strategies involving multiple interfering users have been studied in Information Theory [2, 30, 8, 17]. Some of these ideas have been implemented and evaluated using real systems [11, 22, 26, 24, 36]. However, such techniques need significant coordination among the transmitting nodes.

Techniques proposed in Information Theory, such as Network MIMO [2], need the transmitting nodes to not only coordinate their transmissions, but also exchange data with each other before transmitting. Such requirements are easier to meet for downlink transmissions from different APs. Specifically, APs belonging to the same enterprise network have an Ethernet backbone that allows them to exchange data packets before they transmit simultaneously [22]. Further, they can also use the wireless medium [22, 21] or use powerlines [34] to satisfy the synchronization requirements. However, wireless clients do not have these luxuries. Another technique called interference alignment does not need data to be exchanged between the transmitters, but it either requires tight time and frequency synchronization [8, 17, 9], or requires multiple antennas at clients and APs to work [18, 11].

To improve uplink communication efficiency, coordinated multipoint (CoMP) [16] has been proposed for LTE networks. In CoMP, base stations exchange received samples with each other to decode the uplink packets in a MIMO fashion. However, base stations in LTE networks are connected through dedicated high-speed fiber, which provides much higher capacity than Ethernet backhaul in typical enterprise networks. Researchers have shown that exchanging raw samples can lead to unreasonable traffic on the Ethernet [12, 11, 35]. A recently proposed approach for uplink transmissions called BBN [35] removes the synchronization requirement for clients, which is a big step in the right direction, however, it still needs the APs to maintain sample-level synchronization. Such a requirement is still a hindrance to rapid deployment of this technology as it is non-trivial to meet such synchronization requirements. It also requires a large number of APs ($O(N^2)$ to support N uplink transmissions) which puts an additional requirement on the network density. So, a pressing question is - *Can we enable uplink multi-user transmissions in practical systems, i.e., without requiring tight synchronization among APs or clients, without overwhelming the backbone network, and without requiring a high AP density?*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom'16, October 03-07, 2016, New York City, NY, USA

© 2016 ACM. ISBN 978-1-4503-4226-1/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2973750.2973756>

Figure 1: A 2×2 network with 2 clients and 2 APs which are connected through the backbone network. The weight (s_{ij}) of the dotted line is the received signal strength (RSS) at AP_j for transmission from C_i in Watts.

A well studied technique from Information Theory called Successive Interference Cancellation (SIC) supports decoding of multiple transmissions at a single receiver. It does that by first decoding the strongest signal and treating the rest as noise. It cancels (or subtracts) this decoded signal from the ensemble and continues to successively decode the remaining packets. But the achievable gain has been shown to be limited [23]. Our own evaluations (Figure 8) show that SIC has no gains over omniscient TDMA in more than 50% of the cases and only 20% gain in the remaining cases. Our analysis in Section 2 also indicates that the theoretical throughput gain of SIC over TDMA is limited. The main reason for limited gains is that a single receiver is unable to fully capitalize on the diversity of received signal across multiple receivers. However, in a realistic environment, the added dimension of diversity offered by multiple receivers (or base-stations) can be cleverly leveraged to distributedly apply the interference cancellation technique. Based on this observation, we present BASIC, a novel lightweight multi-user uplink transmission technique that does not require tight synchronization and does not impose any restrictions on the AP density. BASIC exploits the inherent receiver diversity and takes advantages of the Ethernet backbone connection between APs which allows them to exchange decoded packets with each other. It decodes multiple simultaneously transmitted uplink packets according to a chosen sequence but in contrast to SIC, each of those packets can be decoded at a different AP. Each decoded packet is forwarded to the succeeding APs where its interference can be removed so that desired packets can be decoded. Thus, a group of APs collaborate to decode a group of simultaneously transmitted uplink packets while leveraging the backbone. A greedy heuristic is proposed to determine the transmission and decoding plan.

The design of BASIC is particularly challenging due to the following reasons. 1) The problem of determining the transmitting clients and decoding plan is combinatorial in nature. For N clients and M APs, we can choose any subset of the N clients to transmit and the packet from each client can be decoded at any of the M APs. So there are $\sum_{i=1}^N \binom{N}{i} M^i$ combinations that are possible. 2) Packet subtraction across multiple nodes is particularly difficult when time-synchronization is not perfect.

To show how BASIC works, we use Figure 1 as an example. Assume each client has a packet to send to the associated AP. BASIC allows both clients to transmit simultaneously. To achieve correct decoding of both packets, the data rate for C₁ is carefully selected such that the packet

can be decoded with a signal-to-noise-ratio (SNR)¹ of $\frac{s_{11}}{s_{21}}$. With this requirement, AP₁ could receive the packet sent by C₁ correctly with the interference from C₂. The decoded packet is then delivered to AP₂ over the backbone. AP₂ then subtracts this packet from the received samples and decodes the packet from C₂ without any interference. To quantify the gains of BASIC, we choose s_{11} and s_{22} to be 20 dB higher than the noise floor, while picking s_{12} and s_{21} to be 10 dB higher than the noise floor. For this example, TDMA schedules C₁ and C₂ alternately with 20 dB SNR at AP₁ and AP₂, respectively. SIC has no gain over TDMA. It allows C₁ to transmit with 10 dB SINR to AP₁. After decoding the packet from C₁, we can subtract its interference from the received samples and decode C₂'s packet with 10 dB SNR. For BASIC, it also allows C₁ to transmit with 10 dB SINR to AP₁. The decoded packet is forwarded to AP₂. After interference cancellation, the packets from C₂ can be decoded with a SNR of 20 dB at AP₂. Since with high SNR, channel capacity is almost linear to channel SNR, BASIC achieves $\frac{20+10}{20} = 1.5$ times the throughput of both TDMA and SIC.

This paper makes the following contributions:

1. We develop a new uplink packet transmission and decoding strategy that does not require any synchronization.
2. We present novel techniques to enable packet subtraction across multiple unsynchronized APs using fine grained frequency estimation and phase error correction.
3. We evaluate our solution using a 20 node USRP testbed that shows considerable gain over TDMA.
4. Our trace-driven simulations show up to $4.8\times$ gain in throughput with similar flow fairness.

2. MOTIVATION: GAINS FROM EXPLOITING DIVERSITY

This section explores the potential gain of BASIC and it gives intuition as to why such gain is possible. To distinguish the version of BASIC in this section from the one in our final implementation, we denoted it as BASIC-OPT. It differs from BASIC in two key aspects: i) it considers all possible decoding orders to identify the best; and, ii) it assumes that packets can be encoded and modulated to exactly achieve any channel capacity. Our analytical results indicate that the throughput of BASIC-OPT over TDMA keeps increasing with the number of AP-client pairs and reaches a median of $2.6\times$ for 4×4 networks.

2.1 An Example

SIC has been shown to improve the throughput [14] in many scenarios. However, if the data rate of each link is capacity achieving, the gain of SIC is marginal [23]. BASIC takes advantage of the diversity in SNR of different clients to different APs. We show how planning the decoding process at multiple nodes can lead to significant gains over traditional SIC at a single AP. In this section, we analyze the network capacity of different schemes with a simple network of 2 APs and 2 clients (2×2 network), as shown in Figure 1.

¹For simplicity, the channel noise is ignored here.

Network Capacity for TDMA: In TDMA, clients C_1 and C_2 transmit alternately. Assuming each client transmits for equal durations, the capacity is:

$$C_{TDMA} = \frac{B}{2} \log_2(1 + \max\{\frac{s_{11}}{n_1}, \frac{s_{12}}{n_2}\}) + \frac{B}{2} \log_2(1 + \max\{\frac{s_{21}}{n_1}, \frac{s_{22}}{n_2}\}) \quad (1)$$

where n_1 and n_2 are noise power levels at AP_1 and AP_2 , and B is the channel bandwidth.

Network Capacity of SIC: Each AP performs SIC independently and the AP with the highest capacity defines the network capacity. Suppose AP_1 first decodes C_1 's packet followed by C_2 's packet. At AP_1 , the capacity for C_1 is:

$$B \log_2(1 + \frac{s_{11}}{s_{21} + n_1}). \quad (2)$$

C_1 's packet is then subtracted from the combined samples at AP_1 . So, the capacity from C_2 to AP_1 is:

$$B \log_2(1 + \frac{s_{21}}{n_1}). \quad (3)$$

So, the sum capacity at AP_1 is:

$$\begin{aligned} C_{AP_1} &= B \log_2(1 + \frac{s_{11}}{s_{21} + n_1}) + B \log_2(1 + \frac{s_{21}}{n_1}) \\ &= B \log_2(1 + \frac{s_{11} + s_{21}}{n_1}). \end{aligned} \quad (4)$$

It can be shown that the sum capacity remains unchanged if AP_1 decodes the packets in the other order (C_2 's packet followed by C_1 's packet). Similarly, the sum capacity at AP_2 is:

$$C_{AP_2} = B \log_2(1 + \frac{s_{12} + s_{22}}{n_2}). \quad (5)$$

So, the network capacity for SIC is

$$C_{SIC} = \max\{C_{AP_1}, C_{AP_2}\}. \quad (6)$$

Network Capacity for BASIC-OPT: Both clients transmit simultaneously in BASIC-OPT. There are multiple decoding choices. It depends on where each packet is decoded and in what order they are decoded. For the 2×2 network, there are four possible decoding orders. When AP_1 first decodes C_1 's packet, transmits the decoded packet to AP_2 over the backbone, and then AP_2 subtracts C_1 's packet and decodes C_2 's packet, the capacity is:

$$C_{11} = B \log_2(1 + \frac{s_{11}}{s_{21} + n_1}) + B \log_2(1 + \frac{s_{22}}{n_2}) \quad (7)$$

When AP_1 first decodes C_2 's packet, transmits the decoded packet to AP_2 , and then AP_2 decodes the other packet, the capacity is:

$$C_{12} = B \log_2(1 + \frac{s_{21}}{s_{11} + n_1}) + B \log_2(1 + \frac{s_{12}}{n_2}) \quad (8)$$

Similarly, when AP_2 decodes first, the corresponding two capacity terms for the two possible options are as follows:

$$C_{21} = B \log_2(1 + \frac{s_{22}}{s_{12} + n_2}) + B \log_2(1 + \frac{s_{11}}{n_1}) \quad (9)$$

$$C_{22} = B \log_2(1 + \frac{s_{12}}{s_{22} + n_2}) + B \log_2(1 + \frac{s_{21}}{n_1}) \quad (10)$$

BASIC-OPT includes SIC as a special case of decoding. So, the network capacity is:

$$C_{BASIC-OPT} = \max\{C_{11}, C_{12}, C_{21}, C_{22}, C_{SIC}\}. \quad (11)$$

(a) Capacity gain over TDMA in 2×2 networks. For BASIC-OPT, the median gain of 2 clients associated with the same AP is 41%, and the median gain for different association is 62%. For SIC, the median gain is about 29% in both cases.

(b) Capacity gain over TDMA in 3×3 and 4×4 networks. For BASIC-OPT, the median gain of 3×3 networks is 110% and the median gain of 4×4 networks is 156%. For SIC, the median gain is 42% and 54%, respectively.

Figure 2: The throughput gain of BASIC-OPT and SIC over TDMA with ideal data rates.

With the above analysis, we can make the following claim: $C_{BASIC-OPT} \geq C_{SIC} > C_{TDMA}$. The first inequality is trivial. The second inequality can be derived as follows:

$$\begin{aligned} C_{SIC} &= B \log_2(1 + \max\{\frac{s_{11}}{n_1} + \frac{s_{21}}{n_1}, \frac{s_{12}}{n_2} + \frac{s_{22}}{n_2}\}) \\ &> \frac{B}{2} \log_2(1 + \max\{\frac{s_{11}}{n_1}, \frac{s_{12}}{n_2}\}) \\ &\quad + \frac{B}{2} \log_2(1 + \max\{\frac{s_{21}}{n_1}, \frac{s_{22}}{n_2}\}) \\ &= C_{TDMA}. \end{aligned} \quad (12)$$

The improvement of $C_{BASIC-OPT}$ and C_{SIC} over C_{TDMA} depends on the RSS of the links. For example, if we pick $s_{11} = 100$ pW (-70 dBm), $s_{22} = 1$ nW (-60 dBm), $s_{12} = s_{21} = 10$ pW (-80 dBm), $n_1 = n_2 = 1$ pW (-90 dBm), we have $C_{TDMA} = 8.3B$, $C_{SIC} = 9.9B$, and $C_{BASIC-OPT} = C_{21} = 13.2B$. BASIC-OPT performs 59% better than TDMA. As s_{12} decreases, the first term in C_{21} increases and the improvement approaches 100%. Our trace-driven results in the next section show that the capacity gain of BASIC-OPT over TDMA increases almost linearly with the network size.

2.2 Trace-driven Analysis

To understand the advantage of utilizing receiver diversity in practice, we collect the RSS values between APs and clients in a large enterprise network. We use the RSS values to generate a number of network scenarios and the capacity formulations from Section 2.1 are used to calculate the capacity of different schemes.

Experiment Setup: The RSS trace collection is conducted on the first floor of a building on campus. We placed a laptop equipped with an Intel Centrino Advanced-N 6205 Wi-Fi adapter at 60 different locations (offices, labs and classrooms) in the building. Then we logged the RSS values from the APs provided by the university by identifying the “ESSID” field in the MAC header.

Results: There are a total of 103 APs detected during the trace collection process. The SNR of all AP-client pairs vary between 2 dB and 63 dB with a mean value of 15 dB. We generate a number of 2×2 networks from the collected trace in the following way. Two locations are selected first. Then we pick two APs that are observed at both locations. A total of 9077 such networks are created.

Figure 2(a) presents the capacity gain over TDMA in these 2×2 networks. We have partitioned the scenarios into two categories termed “Same AP” and “Different AP” depending on whether the two clients select the same AP or different APs for association based on the RSS values. Note that when the clients are associated with different APs, there is more RSS diversity, which favors BASIC. When both clients are associated with the same AP, the median gain for BASIC-OPT is 41%, whereas the median gain is 62% when the clients are associated with different APs. These results show the advantage of leveraging diversity in a real deployment. On the other hand SIC achieves around 29% median gain over TDMA and shows almost no difference in throughput gain when the clients are associated with different APs.

For 3×3 and 4×4 networks created in the same way as for the 2×2 networks, the throughput gains are shown in Figure 2(b). The median throughput gain for BASIC-OPT increases almost linearly, reaching 110% and 156% respectively, indicating that the performance improvement of BASIC-OPT scales with the network size. The throughput gain of SIC increases slowly with 42% for 3×3 and 54% for 4×4 . The reason behind this is that although new clients contribute to the total SNR at an AP linearly, the capacity gain for SIC is in a logarithmic relationship with this linearly increasing SNR.

3. CHALLENGES IN PRACTICE

The analysis in the previous section showed that significant gains can be achieved using BASIC. However, there are multiple challenges in implementing BASIC in reality for achieving the best performance.

1. **Decoding Order and Data Rate Selection:** As shown in Section 2.1, there are four different decoding orders for BASIC and each order may have a different capacity. The number of choices for the decoding order increases exponentially with the network size, making it challenging to compute the best order. Also, we assume the existence of an ideal transmission scheme that achieves capacity. However, in existing 802.11a/g standards, there are only 8 different data rates and

Figure 3: The timeline of the transmissions of APs and clients in the network shown in Figure 1.

none of them is capacity achieving in reality. We need a computationally inexpensive algorithm to determine the decoding order and data rate for each client while achieving close to optimal throughput performance.

2. **Interference Cancellation:** In order to remove interference, we need to reconstruct the samples from the decoded packets. It requires compensating for sampling and frequency offset [10]. In existing interference cancellation techniques, the packet decoding and interference cancellation happens at the same node [14, 10]. So the decoder is able to keep track of the frequency offset estimation error using pilot tones. However, the interfering packets in BASIC may have been decoded at a different AP. Since the APs are not frequency synchronized, the fine grained frequency offset information at one AP cannot be used for reconstruction at another AP.
3. **Fairness Issues:** Throughput optimized schemes tend to make some clients suffer from starvation. A tradeoff between fairness and throughput needs to be considered for the implementation of BASIC in practice.

4. THE DESIGN OF BASIC

In this section, we discuss the design details of BASIC. To decode all packets correctly, we need to determine the decoding order of the packets and select proper data rates for each packet such that interference can be tolerated. The analysis in Section 2.1 shows that the key to solving this problem is the accurate estimation of RSS values between all APs and clients.

Here we define some terms used in BASIC. Let the client set be $Client = \{C_m : m \in [1, M]\}$, the AP set be $AP = \{AP_n : n \in [1, N]\}$, and the RSS values between all AP-client pairs be $RSS = \{s_{mn} : RSS \text{ from } C_m \text{ to } AP_n\}$. Denote all data rates and the minimum SNR required for each data rate as $\Delta = \{(SNR_l, d_l) : l \in [1, L]\}$. Define the following ordered sequence as a *candidate* for BASIC: $\{(C_{i_1}, AP_{j_1}, d_{l_1}), \dots, (C_{i_T}, AP_{j_T}, d_{l_T})\}$, where T is cardinality of the selected subset of clients to participate in the concurrent transmission. Each C_{i_k} in the *candidate* is distinct while AP_{j_k} and d_{l_k} do not necessarily need to be unique. For a candidate, the packet from C_{i_k} is to be decoded at AP_{j_k} with data rate d_{l_k} in presence of noise n_{j_k} and interference from clients decoded after C_{i_k} , i.e.,

$$\frac{s_{i_k j_k}}{\sum_{t=k+1}^T s_{i_t j_k} + n_{j_k}} \geq SNR_{l_k}, \forall k \in [1, T]. \quad (13)$$

The objective of BASIC is to find out the *candidate* that maximizes the sum throughput $\sum_{t=1}^T d_{l_t}$.

4.1 BASIC Overview

In this part, we briefly discuss how the APs estimate the RSS values from the clients accurately and how the APs decode the packets from the clients. To coordinate the APs and clients, one AP is elected as the *Head AP*. It sends commands to the clients through the wireless channel and interacts with other APs using the backbone network. The multi-phase protocol works as follows. The *Head AP* first requires each of the clients to send a preamble in the specified order to estimate channel RSS values. These RSS values are then used to calculate the best *candidate* for BASIC. The *Head AP* then informs the clients to transmit with the given data rates according to the best *candidate*. After the clients finish sending, the APs decode the packets in the order specified in the *candidate*. In the following sections, we explain each phase in detail. We still use the 2×2 network in Figure 1 as an example and assume AP_1 is the *Head AP*. The nodes in BASIC transmit according to the timeline shown in Figure 3.

4.2 Channel Estimation Phase

The uplink transmission slot begins with AP_1 broadcasting a *Poll* message, which contains the IDs of an ordered list of clients ($\{C_1, C_2\}$ in this example). This ordered list informs the selected clients to transmit a preamble sequentially in that order. In this example, C_1 sends a preamble followed by C_2 . When there are too many clients in the network, BASIC cannot schedule all of them to transmit simultaneously because the interference level will be too high for reliable transmissions. In the next section, we will discuss the process for selecting the clients to transmit in a given slot. Here we assume C_1, C_2 are selected. When C_1 and C_2 receive the *Poll* message, they send back a standard 802.11a PHY preamble in the assigned time slot. Since the clients and APs in BASIC are not necessarily synchronized and the propagation delays between all AP and client pairs vary, a small guard interval of $2\mu s$ is inserted between these preambles to protect the transmissions. These preambles allow the APs that overhear them to perform the following two operations: i) Estimate the RSS values on all subcarriers from each client which are used to calculate the decoding order and data rate for the clients; and ii) Estimate the channel properties, including frequency offset and sampling offset, from each client for interference cancellation in the *Decoding Phase*.

4.3 Data Rate Selection Phase

In this phase, AP_1 collects the RSS values from AP_2 and uses a data rate selection algorithm to find the best *candidate* for BASIC. Our analysis in Section 2.1 selects the optimal decoding order for BASIC. However, solving the decoding order problem is a non-trivial combinatorial problem. In this section, we propose a greedy polynomial-time algorithm for data rate selection that is described in Algorithm 1.

In order to provide each client a fair transmission opportunity, we propose a *credit* based client selection algorithm. Each client maintains a weight which is equal to the number of time slots that is allocated to the client for transmission and we choose the set of clients to transmit according to their credits.

The basic idea of Algorithm 1 is to select a client that maximizes the achievable throughput at each step. Since there

Algorithm 1: Maximum SINR greedy algorithm

```

1 Input: (a) a client index set,  $Client = [1, M]$  and their
    $Credits = \{credit_1, \dots, credit_M\}$ , such that
    $credit_i \leq credit_j, \forall i < j$ ; (b) an AP index set,
    $AP = [1, N]$ ; (c) the RSS between all AP-client pairs;
   (d) all possible data rates available in increasing order
   and minimum SNR required:
    $\Delta = \{(SNR_l, d_l) : l \in [1, L]\}$ ; (e) the noise level array
   at each AP,  $N = \{n_i : i \in [1, N]\}$ ; (f) 90% percentile
   value of residual interference from Section 6,
    $Residual = \{residual_i : i \in [minsnr, maxsnr]\}$ .
2 Output: A candidate for BASIC:  $\{(C_i, AP_j, d_k) : C_i \in Client, AP_j \in AP \text{ and } d_k \in \Delta\}$ .

3  $candidate \leftarrow []$ ;
4  $clientSet \leftarrow []$ ;
5  $tmpCset \leftarrow []$ ;
6  $maxTp \leftarrow 0$ ;
7 for  $cNew \in Client$  do
8    $tmpCset \leftarrow tmpCset \cup \{cNew\}$ ;
9    $tmpCand \leftarrow []$ ;
10   $tmpTp \leftarrow 0$ ;
11  for  $ap \in AP$  do
12     $S[ap] \leftarrow \sum_{c \in tmpCset} RSS[c][ap] + n[ap]$ ;
13  while  $tmpCset \neq \emptyset$  do
14     $sinr[c][ap] \leftarrow \{\frac{RSS[c][ap]}{S[ap] - RSS[c][ap]} : c \in tmpCset, ap \in AP\}$ ;
15     $maxSinr \leftarrow \max_{c \in tmpCset, ap \in AP} \{sinr[c][ap]\}$ ;
16     $(newClient, newAp) \leftarrow (c, ap) : sinr[c][ap] = maxSinr$ ;
17     $rate \leftarrow \Delta[\lfloor maxSinr \rfloor]$ ;
18    if  $rate = 0$  then
19       $tmpTp \leftarrow 0$ ;
20      break;
21     $tmpCand.push\_back((newClient, newAp, rate))$ ;
22     $tmpCset \leftarrow tmpCset \setminus newClient$ ;
23     $tmpTp \leftarrow tmpTp + rate$ ;
24    for  $ap \in AP$  do
25       $S[ap] \leftarrow S[ap] - RSS[newClient][ap]$ 
26       $+ residual[RSS[newClient][ap]/n[ap]]$ ;
27  if  $tmpTp > maxTp$  then
28     $candidate \leftarrow tmpCand$ ;
29     $maxTp \leftarrow tmpTp$ ;
30 return  $candidate$ 

```

is a close relationship between SINR and data rate, the algorithm picks the client with the maximum SINR across all APs. The algorithm consists of two loops. In the outer loop, a new client with the highest *credit* priority is added to the transmitting set (lines 7-8). Then it calculates the sum RSS at each AP (lines 11-12). In the inner loop, it appends the client with the maximum SINR to the client decoding order list (lines 14-21). If the data rate for a client is 0, it indicates that the current set of clients should not transmit concurrently (lines 17-20). The total RSS at each AP is updated since BASIC is able to remove the interference of the decoded packets (lines 24-26). The *candidate* is updated if the current set of clients can achieve a higher total throughput (lines 27-29).

Henceforth this greedy algorithm is known as MaxSINR. We have compared the performance of MaxSINR and an exhaustive search based algorithm which evaluates all possible subsets of clients for decoding and chooses the best one. Evidently, this exhaustive search based algorithm re-

Figure 4: Throughput Comparison of MaxSINR and Exhaustive Search. Network contains 50 clients.

quires exponential computation time. Figure 4 compares the throughput of exhaustive search based algorithm and MaxSINR. The network contains 50 clients and the number of APs are varied. Achievable throughput of MaxSINR is only 6% worse than the throughput of exhaustive search based algorithm. It is possible that the total capacity of the exhaustive search based algorithm could be much higher than the total capacity of MaxSINR algorithm. However, 802.11 standard uses discrete data rates causing a small difference between the achievable throughput of MaxSINR and exhaustive search based algorithm.

4.4 Data Transmission Phase

The Head AP AP_1 sends a *Trigger* message to the clients with the data rate information in the beginning of this phase. Then, the clients transmit simultaneously according to the chosen data rates. For the sake of simplicity, we define the transmission time to be of a fixed duration. The clients can perform packet aggregation and splitting to fill the whole transmission duration. In this phase, all APs store the received samples from the clients.

4.5 Decoding Phase

In this phase, assume that packet x_1 from C_1 is decoded first at AP_1 correctly. Then x_1 is forwarded to AP_2 where it is subtracted from the received samples. This interference cancellation process has been studied widely in many previous works [14, 10, 11]. We refer the reader to the existing literature for various techniques used for interference cancellation. However, there is a fundamental challenge in the interference cancellation in BASIC that is different compared to the existing techniques. As discussed in [10, 11], we need to compensate for both frequency and sampling offsets to reconstruct the samples from the decoded packets. Both of these can be obtained from the preamble in the *Channel Estimation* stage. The impact of residual sampling offset on interference cancellation is similar throughout the packet. However, the impact of residual frequency offset keeps increasing with the number of samples in the packet. Although the Wi-Fi preamble in the *Channel Estimation* phase provides an estimate for the frequency offset, it is not accurate enough to reconstruct the packet correctly. Assume the residual frequency offset is Δf . Let the actual received samples for the interference packet be $S_{actual} = \{a_1, \dots, a_n\}$, where n is the total number of samples. Due to frequency offset, the samples reconstructed using the preamble estimation are $S_{rebuilt} = \{e^{\frac{2\pi\Delta f j}{B}} a_1, \dots, e^{\frac{2\pi\Delta f n j}{B}} a_n\}$, where B

is the bandwidth. Using Taylor series for exponential function and ignoring the higher orders, the residual interference strength for the k^{th} sample can be written as follows:

$$r_k = |(1 - e^{\frac{2\pi\Delta f k j}{B}})a_k|^2 \approx |\frac{2\pi\Delta f k j}{B}a_k|^2,$$

which is quadratic with the index k .

According to [27], the frequency estimation accuracy using the Wi-Fi preamble could be as low as 0.1 ppm at 25 dB SNR, which indicates 500 Hz residual frequency offset at the 5 GHz ISM spectrum. For a transmission duration of 1 ms, this inaccuracy leads to a phase shift of π for the last sample of the packet, which is the *reverse* of the actual sample. Obviously, the interference cancellation will fail with this estimation. Existing interference cancellation techniques decode and subtract the packet at the same node, which allows them to keep updating the frequency offset estimation. This can not be applied to BASIC because the decoding and cancellation can happen at different APs. In BASIC, we solve the residual frequency offset problem in two steps: fine frequency estimation and phase error correction based on the following observations. 1). *The noise level in the correlation value of two similar sample sequences keeps decreasing with the length of the sample set.* 2). *Since the samples from different transmitters are independent, the interference from other transmitters can be treated as noise.*

Fine Frequency Estimation: The frequency estimation error varies between 0.1 ppm and 1 ppm with Wi-Fi preamble [27] under different SNR settings. To achieve better estimation accuracy even under low SNRs, we use the first $2K$ samples from $S_{rebuilt}$ to do correlation with the corresponding $2K$ samples from S_{actual} . These samples are divided into two parts and the correlation values of both halves are calculated in the following way:

$$COR_1 = \sum_{i=1}^K \frac{e^{\frac{2\pi\Delta f i j}{B}} a_i}{a_i} = \frac{e^{\frac{2\pi\Delta f j}{B}} (1 - e^{\frac{2\pi\Delta f K j}{B}})}{1 - e^{\frac{2\pi\Delta f j}{B}}},$$

$$COR_2 = \sum_{i=K+1}^{2K} \frac{e^{\frac{2\pi\Delta f i j}{B}} a_i}{a_i} = \frac{e^{\frac{2\pi\Delta f (K+1) j}{B}} (1 - e^{\frac{2\pi\Delta f K j}{B}})}{1 - e^{\frac{2\pi\Delta f j}{B}}},$$

$$\text{angle}(\frac{COR_2}{COR_1}) = \text{angle}(e^{\frac{2\pi\Delta f K j}{B}}) = \frac{2\pi\Delta f K}{B} \text{ mod } 2\pi,$$

where the function $\text{angle}(x)$ returns the phase of x . Assume $K = 2000$. When $|\Delta f| < 10000$ Hz (2 ppm at the 5 GHz band) with $B = 20$ MHz bandwidth, $\text{angle}(\frac{COR_2}{COR_1}) = \frac{2\pi\Delta f K}{B}$. This gives us a more accurate frequency estimation as

$$\Delta \hat{f} = \frac{B \times \text{angle}(\frac{COR_2}{COR_1})}{2\pi K}, \quad (14)$$

Phase Error Correction: Although the above technique further decreases the frequency estimation error, a small amount of frequency offset still remains. Also the above technique cannot capture the offset due to frequency drift. The residual frequency offset needs to be accurately derived during reconstruction. Since the frequency offset is represented as phase rotation in the constructed samples, we divide the whole packet into smaller blocks each with M samples and estimate the accumulated phase error for each block. Assume the residual frequency offset after the *Fine*

Frequency Estimation is now $\Delta f'$. Denote the N^{th} block of the reconstructed samples after correction using $\Delta \hat{f}$ as $S'_{rebuilt} = \{e^{\frac{2\pi\Delta f'((N-1)M+1)j}{B}} a_{(N-1)M+1}, \dots, e^{\frac{2\pi\Delta f'NMj}{B}} a_{NM}\}$. The correlation between S_{actual} and $S'_{rebuilt}$ is as following:

$$\begin{aligned} COR &= \sum_{i=(N-1)M+1}^{NM} \frac{e^{\frac{2\pi\Delta f'ij}{B}} a_i}{a_i} \\ &= e^{\frac{2\pi\Delta f'((N-1)M+1)j}{B}} \frac{1 - e^{\frac{2\pi\Delta f'Mj}{B}}}{1 - e^{\frac{2\pi\Delta f'j}{B}}}. \end{aligned} \quad (15)$$

Assume $|\frac{2\pi\Delta f'M}{B}| \ll \pi$ after the *Fine Frequency Estimation*. Then the higher orders in the Taylor Series of $e^{\frac{2\pi\Delta f'Mj}{B}}$ can be ignored. We have $\angle(COR) \approx \frac{2\pi\Delta f'((N-1)M+1)j}{B}$, which is the accumulated phase error for the N^{th} block.

Our experiment results in Figure 6 show that these two steps allow us to keep the residual interference level below 2 dB for the whole packet.

4.6 Communication Overhead

Communication overhead comes from the channel estimation phase. Such a process is usually also required by MU-MIMO implementations [33]. The transmission time for the *Poll* and the *Trigger* messages are both 40 μ s in 802.11a/g. The transition time between transmitting and receiving is 9 μ s; the transition time between the clients preambles is 2 μ s; and, each preamble takes 16 μ s. So the time overhead of the coordination is $(98 + (16 + 2) \times N)$ μ s, where N is the number of clients. When $N = 4$, the overhead is 170 μ s. If the packet duration is set to just 2 ms (1500 Bytes at the lowest data rate, 6 Mbps), the overhead is as high as 8.5%. In order to amortize this overhead, we set the packet duration to 10 ms in BASIC so that the overhead is 1.7% for 4 clients and 2.8% for 10 clients. Note that packet aggregation is a common technique seen in 802.11n and some MU-MIMO designs [25].

Further, client selection can be decoupled from data transmission phase, as is done in 802.11 ac. Instead of estimating channels from all clients, a group of clients are chosen based on QoS requirements or scheduling policies, and estimation is done only for these clients to reduce overhead [28].

In the above overhead calculation, the duration for the *Data Rate Selection* phase is not considered. The reason is that the first three phases can be decoupled and happen at a different time. Although the *Data Rate Selection* phase takes almost 200 μ s to finish due to delay over the backbone network [26], other devices in the surrounding could transmit during this phase. Since the channel coherence time with a walking speed is around 30 ms [31], the RSSs remain almost the same from the time the *Channel Estimation* phase finishes to the time when the *Data Transmission* phase finishes.

5. BASIC IN MULTIPLE COLLISION DOMAINS

Enterprise networks are typically deployed over a large area such as an airport, a library, and a shopping mall. In addition, the Wi-Fi signal also attenuates quickly in the air which may result in multiple collision domains. So, we need to design a technique to divide a large network into smaller

groups of APs and clients and run the BASIC protocol in those groups.

We use the dynamic group formation scheme as mentioned in [4] to divide a large network into smaller groups of nodes. In this scheme a node can be in one of two states: *recovery* and *idle*. In the *recovery* state, a node is part of a group and is actively transmitting or receiving packets. In the *idle* state, a node is not part of any group or not transmitting or expecting any reception. An AP, that wants to form a group and does not have any neighbor that is in the *recovery* state, broadcasts a *join* message. This AP is referred to as *Group Head*. Any node that can hear this message joins the group if: i) it is in *idle* state; ii) all its neighbors are in *idle* state; and, iii) if the node is an AP then the average round trip delay from this AP to *Group Head* is not above a threshold. We use the same threshold value as mentioned in [4]. After a node joins the group it rebroadcasts the *join* message. The BASIC protocol is repeated for multiple time slots within a group so that every client gets an equal opportunity to transmit.

For each group, if we use the scheme from Section 4 and use the same *Head AP* in every time slot, a client will suffer from starvation if it cannot hear the corresponding *Head AP*. To avoid this problem, we use a *Virtual Head AP*, instead of the *Head AP*, to transmit the commands through the wireless channel. The *Virtual Head AP* is selected dynamically for each slot such that its transmission can be heard by the client that has sent the least number of bytes. For the first slot we use the *Group Head* as the *Virtual Head AP*. If a client cannot hear the *Virtual Head AP*, it will not transmit the preamble. The fairness mechanism in BASIC ensures that this user will get more service in later slots to compensate for this missed opportunity.

6. EXPERIMENTS

In this section, we first evaluate BASIC using a software-defined radio testbed with 4 USRP N210s equipped with XCVR2450 daughterboards. We then present results from the ORBIT [1] testbed based on experiments with 20 USRPs. Our experiments are based on the GNURadio IEEE 802.11 implementation introduced by Bloessl et al. [7]. The code contains both transmitter and receiver implementations that work with commercial 802.11a/g/p devices. We extended the code by adding proper channel estimation on each subcarrier and soft-input-soft-output (SISO) decoding module. Although USRP is able to support 20 MHz bandwidth, processing that many samples in real time overwhelms the host computers. To enable real time implementation of BASIC, we set the bandwidth to 1 MHz in our experiments.

6.1 Microbenchmarks

6.1.1 SNR-PRR Relationship

We first test the modified implementation of IEEE 802.11 with two USRPs, one serving as a transmitter while the other as a receiver. These two USRPs are placed 3 meters apart in an office environment. We fix the modulation and encoding scheme and gradually change the transmission gain at the sender to capture the packet reception ratio (PRR) at different SNR levels. For each transmission gain setting, 200 packets are sent and the packet size is set to 1000 Bytes. The receiver keeps track of the SNR of each packet and logs the total number of correctly and incorrectly decoded pack-

Figure 5: The SNR and Packet Reception Ratio of different modulation and coding schemes used in our evaluations.

Figure 6: The residual interference to noise ratio (RINR) under different SNR conditions with 20 MHz bandwidth.

ets. The SNR and PRR relationship is shown in Figure 5. Because of the use of SISO decoding, these curves showed better performance compared to experiment results in [7, 20]. We use the results from this figure to guide our data rate selection in Algorithm 1. Since our decoding process has a cascading nature, an incorrectly decoded packet makes the following packets undecodable. So we pick the SNR of the 99% PRR point as the minimum SNR requirement for each data rate.

6.1.2 Residual Interference Level

In this section, we study the performance of interference cancellation with the frequency offset compensation scheme described in Section 4.5. Since it is difficult to quantify the residual interference level when it is mixed with signals from another packet, we test the performance with only one transmitter and one receiver. We then subtract the packet transmitted from the received signal using interference cancellation techniques.

Figure 6 shows the residual interference to noise ratio (RINR) relationship with the original signal SNR. To better show the cancellation performance, we perform this experiment with 20 MHz bandwidth. When the SNR is below 23 dB, the average residual interference level is less than 1 dB higher than the noise floor. An average of 20 dB cancellation is achieved when the original SNR is higher than 23 dB. In our data rate selection mechanism in BASIC, we do not assume perfect interference cancellation. Instead, we adjust the SINR using the signal level of the subtracted packet according to the result shown in Figure 6.

(a) Low SNR (8 dB)

(b) High SNR (21 dB)

Figure 7: The RINR for different part of the packet under different schemes.

Since the residual frequency offset of each test is different, we evaluate the frequency offset compensation scheme in BASIC by analyzing results from two tests under different SNR settings as shown in Figure 7. The RINR level for different samples in the packet are plotted. BASIC-FREQ only implements the *Fine Frequency Estimation* while BASIC-PHASE only uses the *Phase Error Correction*. BASIC-NONE refers to the cancellation scheme with neither *Fine Frequency Estimation* nor *Phase Error Correction*. The figures show that the cancellation performance of BASIC remains flat for the whole packet under all settings. Since BASIC-FREQ only attempts to achieve better frequency offset, its performance degrades with the number of samples due to the residual frequency offset. For BASIC-PHASE, we set $M = 3000$, which results in a clear cancellation pattern every 3000 samples. However, without a fine frequency estimation, the assumption that $\angle(COR) \approx \frac{2\pi\Delta f'NM}{B}$ does not hold. So the lowest RINR for BASIC-PHASE is still higher than BASIC. Without any of the frequency compensation techniques, the RINR of BASIC-NONE shows a clear pattern due to frequency estimation error and the highest RINR is even higher than the original signal level when the phase of the rebuilt sample is actually the reverse of the received sample.

6.2 Testbed Results

We first evaluate the throughput performance of BASIC in a 2×2 network with 4 USRPs. This local testbed is in an office with six cubicles and other furnitures. We place the nodes arbitrarily within the room to create 100 topologies in total. A total of 1000 packets with 1000 Bytes each are sent for each setting. In TDMA, we pick the best data rate for each client based on its maximum SNR to the APs. We

Figure 8: The throughput gain distribution of BASIC and SIC over TDMA in a 2×2 testbed.

Figure 9: The throughput gain distribution over TDMA in a 3×3 testbed.

divide the network scenarios into two categories as before, i.e. “Same AP” and “Different AP” based on the RSSs between the APs and clients. The throughput gain of BASIC and SIC over TDMA is shown in Figure 8. When the clients are associated with different APs, the median throughput of BASIC is $1.5 \times$ the throughput of TDMA, whereas an average of $1.24 \times$ throughput is achieved when the clients are associated with the same AP. This indicates that BASIC leverages diversity in RSS. SIC, on the other hand, shows no throughput gain over TDMA in more than 60% of the tests, reconfirming the results reported in [23].

We also studied the performance of BASIC in the ORBIT testbed [1]. This testbed has 20 nodes at fixed locations. Given the placement of the nodes, the 4×4 (or larger size) topologies that we can create are quite similar in terms of SNR. Hence, we restricted our experiments to 3×3 networks, and explored the performance with larger networks in Section 7 with simulations. We also implement the following schemes for comparison:

1. **TDMA**: At each data transmission slot, TDMA picks a client in a round-robin manner. The data rate of the chosen client is selected according to the maximum RSS value across all the APs so that the chosen client can get the best data rate possible.
2. **SIC**: We perform SIC but in a single AP, hence it does not take assistance from the backbone. We calculate

the best achievable aggregated data rate at each AP, and then choose the AP where the aggregated data rate is the highest.

3. **Majority Voting (MV)** [32]: For an incorrect packet, each bit is determined using maximum voting from the bits decoded by all APs.
4. **MRD** [19]: Each packet is divided into several blocks. The assumption is that each block is received correctly by at least one AP. Then all combinations of these decoded blocks are tested to see if the checksum is passed or not.
5. **SOFT** [32]: Instead of collecting the decoded bits from all APs, the confidences of each bit are combined with their variances as the weight.
6. **Epicenter** [12]: Instead of exchanging received samples between APs, coarse estimation of the samples are collected to reduce the communication overhead. Specifically, a higher density constellation is used to quantify the received samples.

For each experiment, we randomly picked 3 transmitters and 3 receivers from the testbed. 100 experiments are performed in total. Figure 9 shows the throughput gain of different schemes over TDMA. There is almost no gain for MV and MRD since the assumption that at least one AP receives a block correctly is difficult to meet. SIC, as predicted in [23], achieves little gain over TDMA. The throughput gains of SOFT and Epicenter are marginal (1% and 7%, respectively). There are three reasons for that. First, since there are only 3 APs, the AP with the strongest SNR dominates the combination process, which indicates that the bit error rate after combination does not differ a lot. Second, the relative throughput gain decreases at high data rates due to control overhead [12]. We are not limiting our data rate to small values in the experiments. Third, we used TDMA as the baseline and picked the best data rate possible while the baseline scheme is 802.11 in SOFT [32] and Epicenter [12]. The average throughput gain of BASIC over TDMA is 48%, which is not as high as shown in Figure 2(b). The first reason is that we used discrete data rates in the experiment while Figure 2(b) assumes capacity achieving data rates. The second reason is that our cancellation scheme cannot achieve more than 20 dB interference cancellation, which limits our throughput gain at high SNRs. More complex digital cancellation technique as reported in [6] can achieve 48 dB interference cancellation. We believe that BASIC can achieve higher throughput gain with this technique and plan to implement it in the future.

7. TRACE-DRIVEN SIMULATION

The number of nodes in the testbed limits the network size to evaluate BASIC. In this section, we turn to trace driven simulation in ns-3 and study the performance of BASIC in both single and multiple collision domains.

7.1 Simulation Setup

Besides implementing BASIC in ns-3, we also implemented the following schemes for comparison:

1. **TDMA**: As in Section 6.
2. **SIC**: As in Section 6 except that we choose the best of TDMA and SIC to maximize throughput of SIC.

3. **Symphony** [4]: Symphony takes advantage of the fact that not all APs are able to hear the transmission from a specific client in multiple collision domains. Symphony reduces the number of collisions by suppressing a subset of clients in each time slot. In the end, some APs are able to receive packets without interference. These packets are forwarded to other APs where their interference can be cancelled. Note that Symphony requires multiple collision domains to work while BASIC works in both single and multiple collision domains. In addition, due to repeated transmissions from clients, Symphony is not as energy efficient as BASIC. In our simulations, Symphony uses the highest rate possible for each link based on the SNR of the individual links.

Since ns-3 does not have sample level simulation capability, we did not implement MV, MRD, SOFT and Epicenter. Also, as shown in Section 6, these schemes perform similar to TDMA. To obtain realistic performance from the simulator, we feed the RSS values from our testbed into the simulator. We also include the residual interference trace from Section 6.1.2.

7.2 Single Collision Domain

In single collision domain, APs and clients are placed in such a way that every AP can hear every client in the network. Performance of TDMA, SIC and BASIC is evaluated in this section.

7.2.1 Throughput Performance

Figure 10(a) compares the throughput of different protocols when the number of APs remains the same. It can be seen from the figure that the throughput of all the three schemes remains similar as the number of clients increases. The throughput of BASIC is almost $2.5\times$ the throughput of TDMA and approximately $2.3\times$ the throughput of SIC. This gain can be attributed to the fact that BASIC can decode multiple clients in a single slot in contrast to TDMA. BASIC outperforms SIC because packets from the clients are decoded at different APs to maximize the total throughput.

Figure 10(b) compares the throughput when the number of clients remains the same. The throughput of BASIC increases almost *linearly*, however the throughput of SIC and TDMA increases slightly. As the number of APs increases, the diversity in the network also increases and it helps BASIC to schedule more clients and reach up to $4.8\times$ the throughput of TDMA.

7.2.2 Jain's Fairness Index

Figure 10(c) compares the fairness of different protocols. Fairness of SIC is very low because we only try to maximize its throughput. Fairness of BASIC is slightly lower than TDMA on average. We take a closer look at the throughput of each client. The throughput improvement of each client is plotted in Figure 10(d). As can be seen, none of the clients in BASIC has lower throughput than TDMA, which indicates that the throughput gain of BASIC does not come from the starvation of any client.

7.3 Multiple Collision Domains

In multiple collision domains an AP may not be able to hear all the clients in the network. The network size is $1000\text{m}\times 1000\text{m}$ and the communication range is set to 200m .

The performance of BASIC, Symphony and TDMA are evaluated. SIC is not evaluated in this setup because as seen in single collision domain it is performing poorly.

7.3.1 Throughput Performance

Symphony requires multiple slots to decode the same set of packets whereas BASIC controls the clients data rate so that all the packets can be decoded in a single slot. Moreover, in Symphony a client chooses the data rate based on the highest SNR to any one of the APs that may cause decoding failure. Channel estimation can be erroneous in Symphony as the PN sequences from multiple nodes are received at the same time. Figure 11(a) compares the throughput of different schemes when the number of APs remains the same. The throughput of all the three protocols remain similar as the number of clients increases. The throughput of BASIC is around $1.5\times$ higher than the throughput of Symphony and $3\times$ better than the throughput of TDMA.

Figure 11(c) compares the throughput of different schemes when the number of clients remains the same. The throughput of BASIC increases *linearly*, however the throughput of Symphony changes very slowly while the throughput of TDMA does not change significantly. As the number of APs increases BASIC exploits the diversity better than other protocols by controlling the data rates of the clients and choosing proper decoding schedule. This characteristic is also evident in single collision domain. Moreover, throughput improvement of BASIC is even better than the single collision domain, e.g., for 30 APs and 50 clients the throughput of BASIC in multiple collision domains is $1.5\times$ better than that of single collision domain. This is the result of better SNR diversity in multiple collision domains.

7.3.2 Jain's Fairness Index

Figures 11(b) compares the fairness when the number of APs is fixed. The fairness of BASIC is slightly lower than Symphony for smaller number of clients. But as the number of clients increases the fairness of BASIC also increases. Symphony has better fairness because all the clients get equal chance to transmit in a group but in case of BASIC it may not be true. Figure 11(d) compares the fairness when number of clients is fixed. Again, fairness of BASIC is slightly lower than the fairness of Symphony. However, fairness of BASIC increases with the number of APs. The improved SNR diversity helps BASIC to schedule more clients thus improving the fairness.

8. RELATED WORK

Successive Interference Cancellation: Interference cancellation schemes have been used in the cellular networks for a long time [3]. The mobile cellular devices share a common receiver, the base station, which is able to control the transmission power and data rate of the mobile devices. The use of SIC in Zigbee networks has been studied in [14]. It is shown that SIC is an effective method to combat both the hidden and exposed terminal problems. In [23], the authors studied the performance gain of SIC over TDMA assuming an ideal data rate. It is concluded that the improvement from SIC is marginal in most cases because the restrictive data rate and SNR requirement limit the applicable scenarios for SIC. AutoMAC [13] takes advantage of rateless coding techniques to realize SIC in the uplink transmissions. Since rateless coding allows the clients to achieve throughput close

Figure 10: The performance in single collision domain.

Figure 11: The performance in multiple collision domains.

to capacity, it is able to utilize the full power of SIC. Our work differs from the above in that the interfering packet can be decoded by another AP, largely relaxing the SNR requirement to use SIC.

Backbone Assisted Schemes: In the enterprise networks, the underutilized backbone network provides a side channel for the APs to coordinate with each other. In SOFT [32] and Epicenter [12], the decoded information about the received packet from all APs are collected together to increase the decoding possibility. These schemes work for one transmission at a time instead of concurrent transmissions.

In these works, the authors notice the importance to keep the backbone from getting overwhelmed. Generally, exchanging decoded packets (instead of raw samples) is cheaper as discussed in [35].

Concurrent Transmissions: Interference Alignment (IA) is one way to enable concurrent transmissions [8, 17, 18, 9]. However, IA requires synchronization of mobile clients. Extending existing AP synchronization schemes to mobile clients is not easy. Vidyut [34] needs powerline connections, which is not available to the clients. Because of the mobility of clients, the propagation delays between them keep changing, which requires considerable overhead to estimate such changes with SourceSync [21].

Interference Alignment and Cancellation (IAC) [11] exploits the presence of multiple antennas. BASIC works with both single antenna and multiple antennas. Even with multiple antennas, we believe BASIC will outperform IAC when the channel SNRs are diversely distributed, which is true in most practical settings. In addition, smartphones have limited number of antennas due to their size, which will limit the performance of IAC. BASIC's performance is related to the number of AP/client pairs, which continues to increase in today's dense enterprise networks.

Another simple scheme is that each AP only decodes the packet with the maximum RSSI and no backbone exchange is required. However, this scheme has a really stringent requirement: each AP should have a close by client. We evaluate this scheme with our collected trace in Section 2. It rarely performs better than SIC.

Multi-user MIMO (MU-MIMO) is also a way for concurrent transmissions. SAM [29] enables uplink MU-MIMO without synchronization requirements. However, extending SAM to Network MIMO schemes [2] needs APs to exchange received samples to decode the uplink packets jointly. Such a large amount of high-fidelity digital samples could easily overwhelm the backbone network [11, 33]. For example, in 802.11ac with 160 MHz channel, samples of two receiving streams require 10 Gbps of bandwidth, which is beyond the commonly used Gigabit Ethernet [33]. Also, MU-MIMO is known to have problem of MIMO rank degradation or channel hardening [5], making scalability even harder.

MegaMIMO [22] uses the backbone to share the downlink packets from each AP to facilitate distributed MIMO. CENATUR [26] uses conflict map to schedule the downlink transmissions to overcome hidden and exposed terminal problems. These schemes are all specific to downlink traffic and do not work for the uplink traffic.

The authors of BBN [35] proposed a two slot concurrent uplink transmission scheme to realize uplink distributed MIMO over the backbone. However, the APs need to be synchronized tightly and the number of APs required is quadratic to the number of clients. Moreover, with 6 clients and 16 APs, the throughput gain of BBN over TDMA is 3 $\times$. As shown in Figure 10(a), BASIC just need 10 APs to achieve 3 $\times$ throughput of TDMA. This indicates BASIC is able to outperform BBN when the number of APs is small.

9. DISCUSSION AND CONCLUSION

In this paper, we present BASIC, a novel concept of using the backbone for distributedly decoding packets at multiple receivers using interference cancellation which does not require synchronization and has low energy overhead. Our analysis shows that BASIC can achieve 48% better throughput than TDMA in a real testbed. In simulation BASIC provides $4.8\times$ better throughput than TDMA.

In this section, we discuss some practical issues with BASIC and how we can further improve its performance.

Wideband Transmissions: Traditionally, researchers have believed that SNR is not a good indicator for data rate adaptation algorithms [31]. The main reason is that different subcarriers in the Wi-Fi OFDM system experience different fading and have different SNR levels due to multipath effects. Since we only use 1 MHz as the bandwidth in our experiment, we did not observe large variations in the SNRs on different subcarriers. To make BASIC work over a larger bandwidth with frequency selective fading, we can introduce the effective SNR [15] concept and use the effective SINR for the data rate selection instead.

Power Control: It is shown that transmission power control improves the performance of SIC [23]. Power control provides another dimension for us to adjust the SNR diversity in the network, which also favors BASIC. We leave this as a future research direction.

Co-existence with Legacy Wi-Fi Devices: Current Wi-Fi systems use carrier sensing to resolve channel contention between devices. It also uses network allocation vector (NAV) as a virtual carrier-sensing mechanism. NAV is a data field in the MAC header that indicates the transmission duration of the current packet, during which the channel is supposed to be busy. To co-exist with legacy Wi-Fi networks, in BASIC, the *Head AP* contends for channel access using carrier sensing and sets the NAV field in the *Poll* and *Trigger* packets to proper values to forbid legacy devices from transmitting. Also if a client detects the channel to be busy while receiving the commanding packet from the AP, it will not transmit anything to avoid interference at legacy nodes.

PHY ACKs: In the current design of BASIC, we did not discuss the ACKs for the transmissions. Since the packet decoding happens on the backbone in a cascaded fashion, it takes longer time than the SIFS duration for instant ACKs as in 802.11a/g. As a future work, we plan to implement the block-ACK scheme used in 802.11n/ac, which allows the APs to send delayed ACKs for several packets together.

Integrating with MIMO: MIMO can significantly improve the performance of wireless systems by transmitting multiple streams simultaneously. BASIC can be extended to work with MIMO. For ease of exposition, let us consider a scenario where all clients and APs have 2 antennas each. The clients can transmit 2 streams instead of 1 using standard MIMO techniques. We first decode 2 packets at the first AP, send them to the second AP to do subtraction. Then, we can decode the next 2 packets. When MIMO is used, the data rate selection algorithm needs to be modified to select appropriate data rates for each stream.

Acknowledgements

This work was partially supported by NSF grant #CNS-1302620.

10. REFERENCES

- [1] Open-Access Research Testbed for Next-Generation Wireless Networks (ORBIT), accessed Sept. 2015. <http://www.orbit-lab.org>.
- [2] A. GOLDSMITH, S. JAFAR, N. J., AND VISHWANATH, S. Capacity limits of MIMO channels. *IEEE J. Select. Areas Commun.*, vol. 21, no. 5 (2003), 684–702.
- [3] ANDREWS, J. G. Interference Cancellation for Cellular Systems: A Contemporary Overview. *Wireless Commun.* 12, 2 (2005), 19–29.
- [4] BANSAL, T., CHEN, B., SINHA, P., AND SRINIVASAN, K. Symphony: Cooperative Packet Recovery over the Wired Backbone in Enterprise WLANs. In *Proc. of ACM MobiCom* (2013).
- [5] BHARADIA, D., AND KATTI, S. FastForward: fast and constructive full duplex relays. In *ACM SIGCOMM Computer Communication Review* (2014), vol. 44, pp. 199–210.
- [6] BHARADIA, D., McMILIN, E., AND KATTI, S. Full Duplex Radios. In *Proc. of ACM SIGCOMM* (2013).
- [7] BLOESSL, B., SEGATA, M., SOMMER, C., AND DRESSLER, F. An IEEE 802.11a/g/p OFDM Receiver for GNU Radio. In *Proc. of ACM SRIF* (2013), pp. 9–16.
- [8] CADAMBE, V. R., AND JAFAR, S. A. Interference Alignment and the Degrees of Freedom for the K User Interference Channel. *IEEE Transactions on Information Theory* (2008).
- [9] CHEN, B., YENAMANDRA, V., AND SRINIVASAN, K. Interference alignment using shadow channel. In *Proc. of IEEE INFOCOM* (2015), IEEE, pp. 2128–2136.
- [10] GOLLAKOTA, S., AND KATABI, D. Zigzag Decoding: Combating Hidden Terminals in Wireless Networks. In *Proc. of ACM SIGCOMM* (2008).
- [11] GOLLAKOTA, S., PERLI, S. D., AND KATABI, D. Interference alignment and cancellation. In *ACM SIGCOMM Computer Communication Review* (2009), vol. 39, ACM, pp. 159–170.
- [12] GOWDA, M., SEN, S., CHOUDHURY, R. R., AND LEE, S.-J. Cooperative Packet Recovery in Enterprise WLANs. In *Proc. of IEEE INFOCOM* (2013), IEEE, pp. 1348–1356.
- [13] GUDIPATI, A., PEREIRA, S., AND KATTI, S. AutoMAC: Rateless Wireless Concurrent Medium Access. In *Proc. of ACM MobiCom* (2012), pp. 5–16.
- [14] HALPERIN, D., ANDERSON, T., AND WETHERALL, D. Taking the Sting out of Carrier Sense: Interference Cancellation for Wireless LANs. In *Proc. of ACM MobiCom* (2008), pp. 339–350.
- [15] HALPERIN, D., HU, W., SHETH, A., AND WETHERALL, D. Predictable 802.11 Packet Delivery from Wireless Channel Measurements. In *Proc. of ACM SIGCOMM* (2010), pp. 159–170.
- [16] IRMER, R., DROSTE, H., MARSCH, P., GRIEGER, M., FETTWEIS, G., BRUECK, S., MAYER, H.-P., THIELE, L., AND JUNGnickel, V. Coordinated multipoint: Concepts, performance, and field trial results. *Communications Magazine, IEEE* 49 (2011).
- [17] JAFAR, S. A. *Interference Alignment: A New Look at Signal Dimensions in a Communication Network*. Now Publishers, 2011.
- [18] LIN, K. C.-J., GOLLAKOTA, S., AND KATABI, D. Random Access Heterogeneous MIMO Networks. In *Proc. of ACM SIGCOMM* (2011), pp. 146–157.
- [19] MIU, A., BALAKRISHNAN, H., AND KOKSAL, C. E. Improving Loss Resilience with Multi-radio Diversity in Wireless Networks. In *Proc. of ACM MobiCom* (2005).
- [20] PEI, G., AND HENDERSON, T. Validation of OFDM Error Rate Model in ns-3. *Boeing Research Technology* (2010), 1–15.

- [21] RAHUL, H., HASSANIEH, H., AND KATABI, D. SourceSync: a Distributed Wireless Architecture for Exploiting Sender Diversity. In *Proc. of ACM SIGCOMM* (2010), pp. 171–182.
- [22] RAHUL, H., KUMAR, S., AND KATABI, D. MegaMIMO: Scaling Wireless Capacity with User Demand. In *Proc. of ACM SIGCOMM* (2012), pp. 235–246.
- [23] SEN, S., SANTHAPURI, N., CHOUDHURY, R. R., AND NELAKUDITI, S. Successive interference cancellation: A back-of-the-envelope perspective. In *Proc. of ACM Hotnets* (2010), ACM, p. 17.
- [24] SHA, M., XING, G., ZHOU, G., LIU, S., AND WANG, X. C-MAC: Model-Driven Concurrent Medium Access Control for Wireless Sensor Networks. In *Proc. of IEEE INFOCOM* (2009), pp. 1845–1853.
- [25] SHEN, W.-L., LIN, K. C.-J., GOLLAKOTA, S., AND CHEN, M.-S. Rate adaptation for 802.11 multiuser MIMO networks. *Mobile Computing, IEEE Transactions on* 13, 1 (2014), 35–47.
- [26] SHRIVASTAVA, V., AHMED, N., RAYANCHU, S., BANERJEE, S., KESHAV, S., PAPAGIANNAKI, K., AND MISHRA, A. CENTAUR: Realizing the Full Potential of Centralized Wlans through a Hybrid Data Path. In *Proc. of ACM MOBICOM* (2009), pp. 297–308.
- [27] SOUROUT, E., EL-GHOURY, H., AND MCNEILL, D. Frequency Offset Estimation and Correction in the IEEE 802.11a WLAN. In *Proc. of IEEE VTC* (2004).
- [28] SUNDARESAN, K., KHOJASTEPOUR, M., CHAI, E., AND RANGARAJAN, S. Full-duplex without strings: Enabling full-duplex with half-duplex clients. In *Proc. of ACM MobiCom* (2014), ACM, pp. 55–66.
- [29] TAN, K., LIU, H., FANG, J., WANG, W., ZHANG, J., CHEN, M., AND VOELKER, G. M. SAM: enabling practical spatial multiple access in wireless LAN. In *Proc. of ACM MobiCom* (2009), ACM, pp. 49–60.
- [30] VERDU, S. Multiuser Detection. *Cambridge University Press* (1998).
- [31] VUTUKURU, M., BALAKRISHNAN, H., AND JAMIESON, K. Cross-layer Wireless Bit Rate Adaptation. In *Proc. of ACM SIGCOMM* (2009), pp. 3–14.
- [32] WOO, G. R., KHERADPOUR, P., SHEN, D., AND KATABI, D. Beyond the Bits: Cooperative Packet Recovery Using Physical Layer Information. In *Proc. of ACM MobiCom* (2007), pp. 147–158.
- [33] YANG, Q., LI, X., YAO, H., FANG, J., TAN, K., HU, W., ZHANG, J., AND ZHANG, Y. BigStation: enabling scalable real-time signal processing in large mu-mimo systems. In *ACM SIGCOMM Computer Communication Review* (2013), vol. 43, pp. 399–410.
- [34] YENAMANDRA, V., AND SRINIVASAN, K. Vidyut: Exploiting Power Line Infrastructure for Enterprise Wireless Networks. In *Proc. of ACM SIGCOMM* (2014), pp. 595–606.
- [35] ZHOU, W., BANSAL, T., SINHA, P., AND SRINIVASAN, K. BBN: Throughput Scaling in Dense Enterprise WLANs with Beamforming and Nulling. In *Proc. of ACM MobiCom* (2014), pp. 165–176.
- [36] ZHOU, W., LI, D., SRINIVASAN, K., AND SINHA, P. Domino: relative scheduling in enterprise wireless lans. In *Proc. of ACM CoNEXT* (2013), ACM, pp. 381–392.