A Practical Approach to Strengthen Vulnerable Downlinks using Superposition Coding

S. Vanka*, S. Srinivasa, and M. Haenggi
Department of Electrical Engineering
University of Notre Dame
Notre Dame, IN 46556, USA

***Abstract*— We propose and experimentally demonstrate a novel approach to improve the packet delivery efficiency on a vulnerable downlink (e.g., from a transmitter to a far-away receiver) using superposition coding, a multiuser transmission scheme that forgoes orthogonal transmission and deliberately introduces interference among signals at the transmitter. On a software-radio platform that uses off-the-shelf point-to-point channel codes, we show that a transmitter serving multiple links can use simple two-user superposition codes to dramatically improve (compared to time division multiplexing) the packet delivery efficiency on its most vulnerable links. Interestingly, our results suggest that superposing signals of far-away users on to those of high-traffic users yields the maximum benefits—implying that the degrees-of-freedom gain in doing so can more than compensate for the increased interference from signal superposition.**

***Keywords*– Superposition Coding, Software Defined and Cognitive Radio, Multiuser Systems, Universal Software Radio Peripheral.**

I. Introduction

A recurring feature of many rate-optimal multiuser communication strategies is that they allow signal superposition: different transmissions on a common communication medium can interfere at one or more receivers. Each receiver then recovers its message(s) of interest by optimally exploiting its knowledge of the codebooks of all the interfering transmitters. Well-known instances include optimal schemes for certain types of one-to-many ("broadcast"), many-to-one ("multiple access") and relay-aided communication [1]. These "superposition coding" techniques stand in contrast to more traditional Time/Frequency/Code Division multiplexing schemes, in that they *deliberately allow significant inter-user interference*.

Nevertheless practical approaches to leverage the coding gains from superposition coding (SC) techniques remain largely unexplored, especially for the broadcast scenario where a base station (BS) has potentially several users to communicate to. In this paper we take the first step towards filling this gap. Specifically, we view SC-based approaches as coding schemes that can efficiently provision the BS's power and bandwidth across different links. By restricting each user's signal to its alloted time slot/frequency band/spreading code, orthogonal coding schemes cannot fully exploit the varying degrees of noisiness among the links to various downlink users, and thus require a higher spectral efficiency to support a given bit-rate on a link. This causes the well-known "near-far problem" [2], where the BS over-provisions the links to its nearby users to the detriment of those farther away (or vice-versa). To demonstrate how SC-based approaches can mitigate this problem, we develop simple experimental procedures that build on our recent work [3] to show that two-user superposition codes for the downlink can achieve dramatic improvements in packet delivery efficiency in static wireless environments. We note here that our approach is quite different from other experimental work on multiuser systems, e.g., [4], [5], that focus entirely on the gains in throughputs or transmission rates.

*Corresponding author. Email: (svanka, ssriniv1, mhaenggi)@nd.edu. The authors are with the Emerging Wireless Architectures Lab, Department of Electrical Engineering, University of Notre Dame.

II. Problem Setup

A. A Case for Two-User Superposition Codes

Consider a BS serving several active users[1]. Given the user density in typical urban cellular networks, it is always possible to pick two active users N (the "near" user) and F (the "far" user), as shown in Fig. 1. The key observation here is that N being geographically closer to the BS has a "stronger" (less noisy) link to the BS than F; thus any packet that can be decoded at F can most probably be decoded at N as well (but not vice versa). The use of good superposition codes can exploit this channel ordering.

With two-user SC the BS transmits a weighted sum of the waveforms resulting from individually-coded user packets (or more precisely, codewords) (see Fig. 1). Thus both links enjoy the *combined* degrees of freedom available to N and F, while sharing the transmit power. The idea is to encode F's data such it can be decoded in the presence of interference from N's signal. N can decode its message via Successive Decoding (SD): since it has a stronger link to the BS, N can first decode, and subsequently regenerate and cancel F's contribution to its received signal. N can then decode its own packet. In many information-theoretic models for one-to-many communication it is known that SC combined with SD at N[2] can offer substantial improvements in spectral efficiency to over orthogonal schemes [1]. In fact, SC achieves the capacity for a scalar Gaussian broadcast channel (see, e.g., [6] for a tutorial overview).

[1]The set of active users and each user's share of the BS's resources is determined by the users' traffic profiles and the BS's scheduling policy.

[2]Hereafter, with SC the adoption of SD at N is understood.

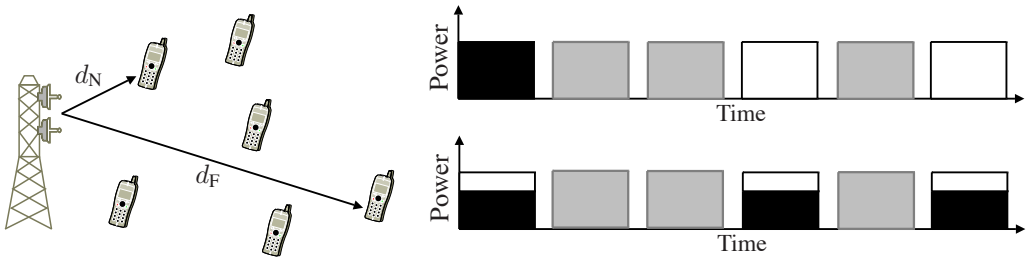

Figure 1. Illustration of two-user SC.(Left) The users N and F picked are at distances d_N and d_F respectively with $d_N < d_F$. (Right) Typical transmission timelines with and without SC. The gray slots represent transmissions to other active users which can remain unchanged. With Time-Division (TD, top), N and F are served in different slots (black and white). With SC (bottom), the BS transmits a linear combination of individually-coded user waveforms.

At *finite* blocklengths, SC can provide a *coding gain* over orthogonal multiplexing schemes such as those based on Time Division (TD)[3]. This gain can be leveraged either as an increased spectral efficiency for fixed link reliabilities[4], or as more reliable communication for fixed spectral efficiencies (for a fixed BS power and bandwidth). In particular, we can use SC towards making the far-link more reliable without degrading the transmission rate or the reliability of the near-link. In the following subsection we examine this idea in greater detail.

B. Improving Link Reliability using SC

We begin by introducing some notation and terminology. The BS has access to a set $\mathcal{C}$ of $M > 1$ *single-user* (point-to-point) channel codes, which we call its *code library*. We will (also) refer to a code's spectral efficiency as its *rate*[5]. Irrespective of the code, each *packet* supplied by the link layer is encoded as one *codeword* of *blocklength* $L < \infty$ [channel uses]. Define the *Packet Error Rate (PER)* $\epsilon < 1$ as the probability that the intended receiver cannot decode its packet. The BS transmits with an average power P [W] over a bandwidth W [Hz] (normalized to 1 without loss of generality (w.l.o.g)). When all of this power is assigned to one link (as in TD), the user enjoys a *single-user SNR* γ. In the following, denote the single-user SNR and PER of a user u by γ_{u} and $\mathrm{PER}_{\mathrm{u}}$ respectively, for $\mathrm{u} \in \{\mathrm{N}, \mathrm{F}\}$. We assume $\gamma_{\mathrm{N}} > \gamma_{\mathrm{F}}$ without loss of generality.

To communicate at a rate r_{N} with N and a rate r_{F} with F, the BS has at least two choices. With Time Division (TD) it can assign a fraction $u \in [0, 1]$ of the total slots to N and the remaining to F. In each slot, the entire power is assigned to the packet being transmitted. To sustain the desired rates, the BS must encode N's packets using a code with rates r_{N}/u and those of F at a rate $r_{\mathrm{F}}/\bar{u}$ where $\bar{u} \triangleq 1-u$. Thus TD eliminates interference between N and F at the cost of increasing the encoding rates of individual packets.

On the other hand, with SC, BS individually codes N's and F's packets *exactly* at r_{N} and r_{F} respectively, and transmits a superposition of these waveforms in every slot. To meet the power constraint, a fraction α of the available power is assigned to N's waveform (and a fraction $\bar{\alpha} \triangleq 1 - \alpha$ to F). Thus SC allows the encoding of individual packets at exactly the desired rate, albeit at the cost of interference between the waveforms.

[3]Hereafter we take TD to be the reference orthogonal scheme.

[4]The link reliability is measured as the probability of a successfully decoding a codeword.

[5]Not to be confused with the code-rate of a binary code.

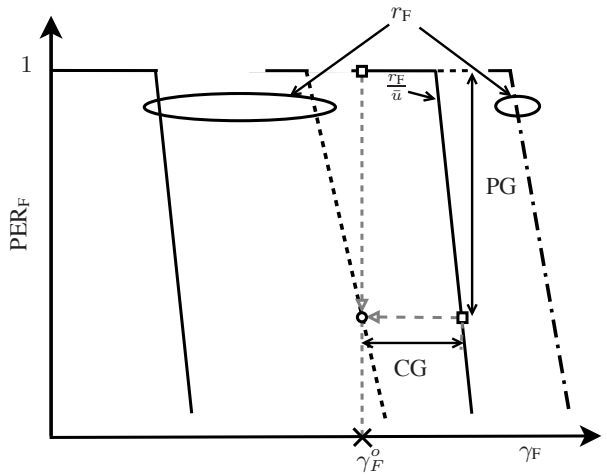

Figure 2. Improving the PER at F using SC. The TD operating points are shown as squares and that of SC as a circle. The coding gain (shown as CG) refers to a reduction in the single-user SNR at F that is made possible by a switch from TD to SC. This gain can be also viewed as a reduction in PER (shown as PG (PER gain)) when this single-user SNR is fixed at γ_F^o.

In particular, for a given γ_{N} and γ_{F}, the PER at F with SC depends on the interference from N's signal, as measured by the Signal-to-Interference Ratio

$$\mathrm{SIR} \triangleq (1-\alpha)/\alpha, \tag{1}$$

We now explain how SC can improve F's PER with the help of Fig. 2, that idealizes the dependence of $\mathrm{PER}_{\mathrm{F}}$ on γ_{F} as a straight-line waterfall curve (for each SIR). Consider two codes with rates $r_{\mathrm{F}}/\bar{u}$ and r_{F},$\bar{u} \in [0, 1]$ from a hypothetical code library. With TD, only the solid PER curves ($\mathrm{SIR} = \infty$) are accessible. Thus for any $\bar{u}$, the PERs with TD are controlled by the waterfall curve of the code with rate $r_{\mathrm{F}}/\bar{u}$. With each finite SIR, each solid curve gives rise to a dashed curve (or a dash-dotted curve, depending on the SIR, which is in turn determined by the choice of N and α). Clearly, only the latter PER curves are accessible to SC.

For any given $r_{\mathrm{N}}, r_{\mathrm{F}}, u$, provided N can cancel most of F's

signal upon decoding it correctly[6], α (and thereby the SIR) can be kept close to the minimum that would be necessary to maintain N's rate and reliability in TD. In this case (dashed curve), good channel codes can help SC leverage the increase in the time slots available to F (more generally, the increased degrees of freedom) to more than compensate for its reduced share of transmit power *and* the interference from N's signal, thereby outperforming TD[7]. There are two ways to measure this performance gain. For the same rate and PER, moving from TD to SC reduces the required γ_{F} by an amount determined by the coding gain (shown as CG). For a fixed operating SNR point (shown as γ_F^o) this transition would provide a PER gain (shown as PG), that translates to a *reliability gain* (RG). Indeed, for a PER of ϵ, a typical packet will require an average of $1/(1-\epsilon)$ transmissions before it is received correctly at F. Thus, given a large number of packets to be sent to F, a BS using TD would require

$$\mathrm{RG} = \frac{1-\epsilon_{\mathrm{SC}}}{1-\epsilon_{\mathrm{TD}}} \tag{2}$$

more transmissions than a BS using SC. If ARQ is used, RG can also be interpreted as a link-level throughput gain. On the other hand, when channel estimation errors prevent N from regenerating F's signal accurately enough despite having decoded it correctly, the residual interference may be large enough to require an increased α to maintain the same rate and reliability as in TD. Apart from increasing N's interference at F, this also reduces F's share $1-\alpha$. In such cases switching to SC may not result in performance gains (dash-dotted curves).

The above framework is quite general, and can be used to compare any superposition scheme with an orthogonal scheme. To derive quantitative results, in the following subsection we will focus on the specific code library and a receiver architecture used in our testbed.

III. Effect of Near-User Interference on the Far-User PER

We first describe the code library and the transmitter and receiver operation. Due to space constraints, our description will only focus on the details that are most relevant to the problem at hand. A more detailed discussion of the testbed's architecture and operation as well as some relevant details about our calibration and measurement procedures can be found in our technical report [7].

For channel coding, the testbed uses off-the-shelf point-to-point channel codes designed using the well-known Bit Interleaved Coded Modulation (BICM) technique [8]. It maps these symbols on to waveforms using Orthogonal Frequency Division Multiplexing (OFDM). The encoder is realized as a binary convolutional encoder followed by an bit-interleaver and a QAM modulator. The decoder is implemented as a maximum-likelihood demodulator, a bit de-interleaver and a Viterbi decoder. The convolutional codes are constructed from a mother code with a generator polynomial $[133, 171]$ using four puncturing patterns to yield code-rates $\{\frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \frac{5}{6}\}$. Each of the four binary codes can be paired with one of three possible interleaver-QAM constellation pairs, one each for gray-coded BPSK, QPSK, or 16QAM. This process results in $4 \times 3 = 12$ BICM codes with rates $\{\frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \frac{5}{6}, 1, \frac{4}{3}, \frac{3}{2}, \frac{5}{3}, 2, \frac{8}{3}, 3, \frac{10}{3}\}$.

In SC mode, the BS first encodes each near- and far-packet using its separate point-to-point BICM encoder into streams of $L = 1536$ symbols each. An adder then sums up N's stream weighted by $\sqrt{\alpha}$ and F's stream by $\sqrt{1-\alpha}$ to produce a single composite symbol stream, which is fed to the OFDM modulator. At the receivers, after standard OFDM pre-FFT processing, both N and F try to recover F's symbol stream first using a maximum-likelihood (ML) demodulator. The demodulator computes the reliability of each of F's code bits in the presence of interference from N's symbols. The de-interleaver and Viterbi decoder process this reliability information in turn to estimate F's data bits. N regenerates F's symbol stream by re-encoding these data bits with a standard point-to-point BICM encoder. After appropriately weighting these symbols by $\sqrt{1-\alpha}$ and the channel coefficient estimates, N cancels F's interference from its received signal and then proceeds to decode its own message using point-to-point BICM decoding.

We now examine the demodulation process at F. Suppose the channel to F has gain h_{F}[8]. Then the demodulator differs from a point-to-point demodulator in that it observes a noisy symbol stream

$$Y_{\mathrm{F}}(n) = \underbrace{h_{\mathrm{F}}\sqrt{\bar{\alpha}P}X_{\mathrm{F}}(n)}_{\text{Signal}} + \underbrace{h_{\mathrm{F}}\sqrt{\alpha P}X_{\mathrm{N}}(n) + Z_{\mathrm{F}}(n)}_{\text{Interference+Noise}}, \tag{3}$$

where $n \in \{1, \ldots L\}$ denotes the symbol index and $Y_{\mathrm{F}}(\cdot), X_{\mathrm{F}}(\cdot), X_{\mathrm{N}}(\cdot), Z_{\mathrm{F}}(\cdot)$ denote the observations, the far- and near-symbol streams, and the white noise sequence respectively. The detection rule clearly depends on the distribution of the perturbation term, that in turn depends on N's constellation. We assume this is known to the demodulator[9]. Now each interfering symbol from N perturbs the original far-symbol (the *parent* point) to a randomly chosen *daughter* point. For each parent point, define the set of all possible daughter points to be its *potential daughter cluster* ("cluster" for short). The shape of this cluster depends on N's constellation, and its spread increases with a decrease in SIR. The demodulator at F infers the most probable parent point of the observed (noisy) daughter point by identifying the most probable cluster to which an observation belongs. Identifying successively less probable clusters helps refine its reliability estimate of each detected code-bit. Analogous to the single-user case, the reliability of the k^{th} bit in the n^{th} symbol is approximated using the max-log-MAP approximation (see [3, Sec. III-B2] for details). The probability of the dominant error events is controlled by the inter-cluster separation

$$d_{\mathrm{eff}} \triangleq \sqrt{\bar{\alpha}P}|h_{\mathrm{F}}| \min_{\substack{p_1,p_2\in\mathcal{X}_{\mathrm{F}},p_1\neq p_2\\ d_1,d_2\in\mathcal{X}_{\mathrm{N}}}} \left| p_1 - p_2 + \frac{d_1-d_2}{\sqrt{\mathrm{SIR}}} \right|. \tag{4}$$

Here $\mathcal{X}_{\mathrm{u}}$ denotes the constellation points of user $\mathrm{u} \in \{\mathrm{N}, \mathrm{F}\}$. For the point-to-point case (when these "clusters" are just

[6]We did not observe imperfect decoding in our experiments.

[7]Indeed, SC systems built with simple off-the-shelf codes can exhibit this property, see Section IV-B.

[8]The analysis that follows can be generalized to each subcarrier in a frequency selective channel due to the use of OFDM.

[9]In practice sending this information entails a small overhead, which we neglect in this paper.

points) the constellation minimum distance can be recovered by allowing SIR $\to \infty$ in (4). It is evident from (4) that the choice of $\mathcal{X}_{\rm N}$ determines the effect of interference at F, suggesting the possibility of modifying its geometry to mitigate this interference. For example, when both N and F are BPSK-modulated, a *rotated*-BPSK constellation for N is preferable to a standard BPSK constellation.

From (4) it is clear that $\mathcal{X}_{\rm N}$ affects $\text{PER}_{\rm F}$. To gain more insight into this problem by constraining both F and N to be BPSK-modulated[10], which we denote as BPSK/BPSK (Signal/Interference). Now $d_{\rm eff}$ has the closed-form

$$d_{\rm eff}^{\rm BPSK/BPSK} = 2P|h_{\rm F}| \min\left(\sqrt{1-\alpha}, \left|\sqrt{1-\alpha}-\sqrt{\alpha}\right|\right), \quad (5)$$

which is plotted in Fig. 3 for $\alpha \in [0,1]$. For a given P and $h_{\rm F}$, its dependence on α can be understood geometrically. Each cluster consists of two daughter points, one each to the left and to the right of the parent point. For $\alpha < \frac{1}{2}$, the cluster interiors[11] do not overlap; hence the nearest points from neighboring clusters lie on the *opposite* sides of their parent points. Increasing α brings these points closer to one another, making them overlap with the origin for $\alpha = \frac{1}{2}$. As α is increased beyond $\frac{1}{2}$, these overlapped points separate, increasing their mutual distance. Here, the value of α that maximizes $d_{\rm eff}$ in (5) can be obtained as the solution to the saddle point equation $\sqrt{1-\alpha} = \sqrt{\alpha} - \sqrt{1-\alpha}$ in $\alpha \in (\frac{1}{2}, 1)$, which is $\alpha = \frac{4}{5}$. For $\alpha \geq \frac{4}{5}$, the nearest points lie on the *same* side of their parent points; their mutual distance is therefore the same as that of their parent points, that in turn vanishes as $\alpha \to 1$. Based on this behavior we will refer to $\alpha = \frac{1}{2}, \frac{4}{5}$ as the *inflection points* for BPSK/BPSK. Arguing similarly, it is easy to see that BPSK/QPSK will have two inflection points (evaluated to be $\alpha = \frac{2}{3}, \frac{8}{9}$) and that BPSK/16QAM will have six inflection points ($\alpha = \frac{40}{76}, \frac{40}{65}, \frac{40}{56}, \frac{40}{49}, \frac{40}{44}, \frac{40}{41}$). A similar analysis can be done for higher order constellations for F (by directly evaluating (4) numerically as necessary).

The inflection points capture the non-monotonic dependence of $d_{\rm eff}$ (and therefore that of $\text{PER}_{\rm F}$) on α. Given that (4) also depends on $P|h_{\rm F}|$, this behavior may not be apparent at small values of P (when the noise is then large enough to cause decoding failures by itself) or when the number of trials limits the statistical reliability of the estimate (which is the case, for example, when the number of observed error events is quite small at larger values of P). In the next section, we experimentally validate the insights obtained so far.

IV. On-Air Experimental Procedure and Results

We perform our experiments on a testbed developed in-house [7]. The testbed software is built on the well-known open-source GNU Radio platform [9] and interfaces with a Universal Software Radio Peripheral (USRP) hardware board that serves as an analog and RF front-end. Some key parameters of the experiment are listed in Table I.

[10]Although not particularly relevant for certain values of SIR (e.g., at low SIR, $\mathcal{X}_{\rm N}$ is more likely to be a 16QAM constellation, see Section IV-B), this case does help understand some key geometric aspects of $d_{\rm eff}$ while permitting simple closed-form results.

[11]As defined by a cluster's convex hull.

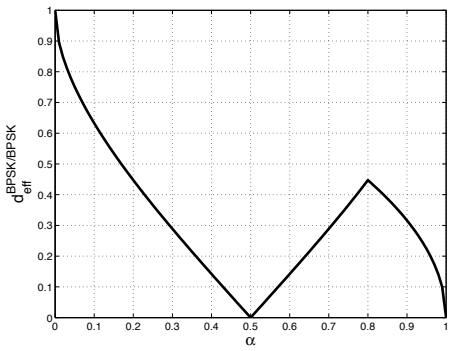

Figure 3. The distance in (5) as a function of α for $2P|h_{\rm F}| = 1$.

Center Frequency	903 MHz
Message Bandwidth	2 MHz
Modulation	16-tone OFDM (8 data, 4 pilot, 4 null)
CP Length	1μs
Average Tx power	-31 dBm

Table I
Parameters used in the experiment.

The experimental setup consists of three USRPs as shown in Fig. 4, where $d_{\rm N} = 0.6$ m, and $d_{\rm F} = 1.2$ m. Note that the goal of the experiment is to fix $\gamma_{\rm N}$ and $\gamma_{\rm F}$—the apparent relationship between $d_{\rm N}$ and $d_{\rm F}$ is due to the table geometry. Similar results can be obtained as long as one link has a much lower SNR than the other. Due to indoor scattering, the disparity between $\gamma_{\rm N}$ and $\gamma_{\rm F}$ can be quite different from what large-scale path-loss models would predict (see Table II) for measured SNR values).

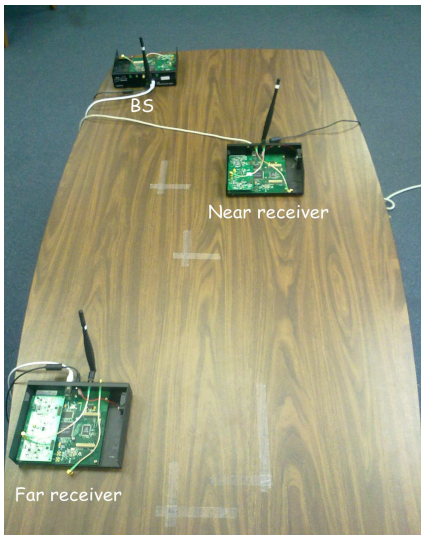

Figure 4. The on-air setup with the three USRPs used to study the efficacy of SC over TD. The BS transmits the private packet to N and the broadcast packet to F. Throughout our experiments, the near and far user distances are fixed at $d_{\rm N} = 0.6$ m and $d_{\rm F} = 1.2$ m.

A. Impact of N's constellation on the PER at F

In light of the discussion in Section III, we first study how $\mathcal{X}_{\text{N}}$ and α affect PER_{F}. Keeping in mind our experiment in Section IV-B, we focus only on BPSK/$\cdot$ and adopt the following procedure: first, we adjust the total transmit power P to obtain a given single-user SNR γ_{F} at F. Second, we transmit $K = 1000$ packets and estimate the PER. We found this value of K to be sufficient for the PER range of interest. The testbed sends a packet every ~ 0.05s, so the channel was kept coherent during a single PER measurement (~ 50 s) . The results are summarized in Fig. 5. We observe the following:

- When γ_{F} is small[12] (noise-limited regime), the noise is large enough to cause decoding errors by itself, thus F's PER is small only at small α.
- For moderate γ_{F}, the PER generally follows the trend in d_{eff} (for instance compare Fig. 3 with the curve for BPSK/BPSK (Fig. 5 (left)). In some cases (see Fig. 5 (right)), many errors can be corrected despite $d_{\text{eff}} = 0$.
- When γ_{F} is large (interference-limited regime), the PER is high only in the vicinity of those inflection points for which $d_{\text{eff}} = 0$; at other values of α, it is very small. Indeed, as long as the signal clusters do not overlap, perfect decoding is possible[13].

Note that F's PER curve is not monotonically increasing with α as one might expect; in fact, it is highly non-monotonic in most cases. Also, depending on $\mathcal{X}_{\text{N}}$ and α, the reliability seen at F can be drastically different.

B. Measuring the Reliability Gain of SC over TD

We now design novel experimental procedures to measure the reliability gain of SC. The basic aim of the experiment is to show that for the same single-user SNR (fixed via the transmit power in a static environment), SC improves PER_{F} over TD without significantly degrading the PER_{N}.

There are five important parameters that determine the performance gain of SC (see Section II-B): (a) F's encoding rate r with TD (b) The share u of F's slots in the slots pooled from N and F, (c) $\mathcal{X}_{\text{N}}$ (as seen from IV-A), (d) F's single-user SNR γ_F, and (e) The near power allocation fraction α (that determines the SIR at F). We now relate these parameters by shrinking this (rather large) parameter space when the BS chooses a nearby N and a distant F, perhaps the most interesting case in practice.

When N is close to the BS, its packets are most likely encoded using spectrally efficient codes (e.g., the 16QAM-5/6 code in our library) in the TD mode. As a result, when the BS tries to maintain the same rate as in TD to a nearby N that contributes most of the combined slots (i.e., small u), it requires a large α (i.e., low SIR) *and* a spectrally efficient ("dense") $\mathcal{X}_{\text{N}}$. The opposite is true for u close to 1. Therefore, choosing a nearby N implies a progressively denser $\mathcal{X}_{\text{N}}$ with increasing α (or equivalently, with decreasing u)[14]. Along similar lines, a distant F means a (relatively) noisy link that operates at a small rate $r' = ur$ to achieve an acceptable PER.

Putting it all together, the experiment would involve measuring PER_{F} for the two approaches that both achieve a rate r': TD, which encodes F's packets at rate r but assigns them only a fraction u of the combined slots, and SC, which can encode at a rate r' by assigning all the slots to F but subjects F's signal to interference from a constellation $\mathcal{X}_{\text{N}}$ that becomes denser as u decreases. We can fix the relatively small rate r' (e.g., BPSK-1/2, the smallest code rate in the library), so that r and $\mathcal{X}_{\text{N}}$ are now controlled by a single parameter u. Each experiment thus involves picking a value of u, and then choosing a pair of operating single-user SNRs $\gamma_{\text{F}}^o, \gamma_{\text{N}}^o$ (implicitly done by choosing the total transmit power P^o and α) that achieve a reliability gain at F. Instead of showing how changing from TD to SC provides reliability gain (as explained in Section II-B), we find it easier to show that changing from SC to TD implies a reliability *loss*.

To this end, we set up the USRPs as shown in Fig. 4. The near-receiver shown therein depicts the N that is paired with F. It is always served using 16QAM-5/6 in the TD mode. Recall that F is always served using $r' = 1/2$ (BPSK-1/2 from the library). In SC mode, for each value of u BS selects N's channel code using its knowledge of N's rate $4 \times 5/6 \times (1-u)$. This fixes $\mathcal{X}_{\text{N}}$, and the BS must now set α and P^o such that both links operate at acceptable PER $\leq 10\%$ (w.l.o.g.).

To achieve this objective, we first define

$$(P_{\text{N}}, P_{\text{F}}) = (\alpha P^o, (1-\alpha)P^o), \tag{6}$$

where P_{N} and P_{F} denote N's and F's signal power respectively. Now, starting from $P_{\text{N}} = P_{\text{F}} = 0$, we adopt the following four-step procedure to find (α, P^o):

1) Keeping $P_{\text{N}} = 0$ and increase P_{F} until the condition $\text{PER}_{\text{F}} \leq 0.1$ is met.
2) With the value of P_{F} from 1), increment P_{N} until[15] $\text{PER}_{\text{N}} < 0.1$.
3) Keeping the ratio $P_{\text{N}}/P_{\text{F}}$ the same, increase P_{N} and P_{F} until $\text{PER}_{\text{F}} < 0.1$. Note that N can still decode its packets at least as reliably as in 2).
4) Use (6) to find (α, P^o).

With this set up we consider the total BS power to be $P^o = P_{\text{N}}^o + P_{\text{F}}^o$ and estimate the single-user SNRs γ_{N}^o and γ_{F}^o at N and F respectively using standard windowed correlator methods (details in [7]).

For TD, we set the transmit power to $P^o = P_{\text{N}} + P_{\text{F}}$. Now N's packets are encoded using 16QAM-5/6, while F's packets are encoded at $r = (1/2)/u$, and the PER[16] at F is

[12] Evidently, the adjectives small, moderate and large are constellation-dependent. For instance, $\gamma_{\text{F}} = 13.73$ dB qualifies as a large value of γ when N's constellation BPSK, but is only a moderate value when N's signal constellation is QPSK (see Fig. 5).

[13] This is only an experimental artifact.

[14] Also note from Fig. 5 that starting from moderate α a dense $\mathcal{X}_{\text{N}}$ results in a worse PER_{F} than a "sparse" (e.g., BPSK) $\mathcal{X}_{\text{N}}$ for a given γ_{F} (ignoring the behavior at the inflection points). Using these points to improve the performance further is a subject of future work.

[15] Due to the impact of imperfect cancellation of F's symbols at N (discussed in detail in [3, Sec. V C. 2)]), P_{N} is, in general, higher than the power that is required to meet that constraint for N alone (when $P_{\text{F}} = 0$).

[16] If r does not exist in the code library, the PER at F may be evaluated using the time-sharing principle. Accordingly, if $r_i < r < r_j$ wherein $r_i, r_j \in \mathcal{C}$, and $\gamma r_i + (1-\gamma) r_j = r$, $0 < \gamma < 1$, we have that the PER at rate $r = \gamma(\text{PER at rate } r_i) + (1-\gamma)(\text{PER at rate } r_j)$.

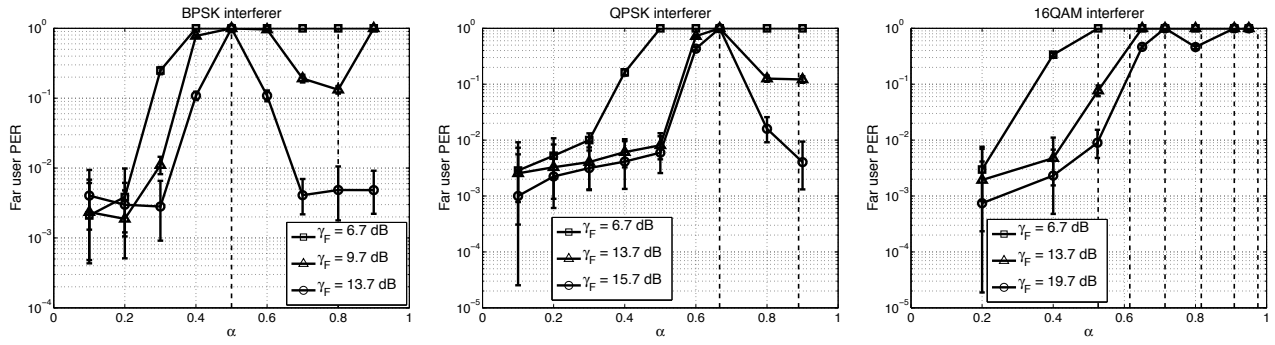

Figure 5. Far user PER (and their 95% confidence intervals) versus the power allocation parameter α at different values of γ for BPSK/BPSK, BPSK/QPSK and BPSK/16QAM respectively. The inflection points are also marked (by dashed lines) in the figure.

	Superposition Coding						**Time Division Multiplexing**				
u	N's rate	F's rate r'	N's PER	F's PER	α	γ_F^0 (dB)	N's peak rate	F's peak rate r	N's PER	F's PER	**RG**
0.1	16QAM-3/4	BPSK-1/2	2.6%	6.8%	0.44	8.8	16QAM-5/6	5	0.1%	Infeasible	N/A
0.2	16QAM-2/3	"	5.3%	6.4%	0.39	7.4	"	2.5	0.1%	100%	∞
0.4	16QAM-1/2	"	3.7%	2.6%	0.24	5.5	"	1.25	0.2%	74.6%	3.83
0.5	QPSK-5/6	"	3.8%	6.6%	0.29	4.5	"	0.5	14.7%	36.7%	1.47
0.55	QPSK-3/4	"	1.6%	5.1%	0.24	4.3	"	0.45	18.9%	38.1%	1.53
0.8	BPSK-2/3	"	3.5%	5.6%	0.2	2.7	"	0.625	77.5%	36.6%	1.49
0.85	BPSK-1/2	"	1.1%	5.1%	0.14	2.6	"	0.5882	80.4%	29.2%	1.34

Table II

ON-AIR EXPERIMENTAL RESULTS. THE LAST COLUMN DEPICTS THE RELIABILITY GAIN (RATIO OF THE NUMBER OF TRANSMISSIONS WITH TD AND SC). FOR EACH EXPERIMENT, WE HAD $\gamma_N^o = \gamma_F^o + 12.8$ DB. THE STANDARD DEVIATION OF THE SNR ESTIMATION METHOD IS ≈ 0.3 DB.

measured. We thus make an *apples-to-apples* comparison: the total transmit power at the BS, the bandwidth and the spectral efficiencies of the two users all remain the same as in the case when SC is employed.

Table II summarizes our results. The values of γ_F^0 and α at the operating points are also listed. The last column depicts F's RG (computed using (2)) that SC offers for F's packets. Note that at every value of u, we obtain reliability gains for F using SC. Interestingly, higher gains are realized at smaller values of u—implying that the *degrees-of-freedom* gain derived in doing so can more than compensate for the increased inter-user interference. Of course, this would also require a larger $\gamma_{\rm F}^o$, as shown. Moreover, in this regime, achieving the equivalent TD rates for F may also be infeasible with our code library (for e.g., when $u = 0.1$). Table II also indicates that a huge benefit in N's reliability is seen at large values of u. Reversing the roles of F and N, our results also suggest that a moderate-to-high rate F needs to be paired-up with a moderate-to-low-rate N.

V. CONCLUDING REMARKS

We have developed a software radio platform to experimentally validate the efficacy of superposition coding over time division multiplexing. Our contributions are two-fold: First, we clearly describe the dependence of N on F's packet error performance, thus motivating the need for joint code optimization. Second, we show how transmitting F's message at a reduced rate in the presence of deliberately-introduced interference can dramatically improve F's PER (and also its link-layer throughput, if ARQ is used).

ACKNOWLEDGMENT

The support of NSF (grant CNS-1016742) and DARPA/IPTO IT-MANET program (grant W911NF-07-1-0028) is gratefully acknowledged.

REFERENCES

[1] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed., John Wiley & Sons, Inc., 2006.

[2] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*, Cambridge University Press, 2005.

[3] S. Vanka, S. Srinivasa, Z. Gong, P. Vizi, K. Stamatiou, and M. Haenggi, "Superposition Coding Strategies: Design and Experimental Evaluation", *IEEE Trans. Wireless Comm.*, Submitted, 2011. Available at www.nd.edu/~mhaenggi/pubs/twc12b.pdf.

[4] J. Zhang, J. Jia, Q. Zhang and E. M. K. Lo, "Implementation and Evaluation of Cooperative Communication Schemes in Software-Defined Radio Testbed," *IEEE INFOCOM*, 2010.

[5] I. Seskar and N. B. Mandayam, "Software Defined Radio Architectures for Interference Cancellation in DS-CDMA Systems," *IEEE Personal Communications Magazine*, Vol. 6, pp. 26-34, 1999.

[6] R. Zhang and L. Hanzo, "A Unified Treatment of Superposition Coding Aided Communications: Theory and Practice," *IEEE Communications Surveys and Tutorials*, No. 99, pp. 1-18, July 2010.

[7] Z. Gong, M. Haenggi, S. Srinivasa, K. Stamatiou, S. Vanka and P. Vizi, "SC: Technical Report", Available at www.nd.edu/~mhaenggi/pubs/techreport_sc11.pdf.

[8] A. G. Fabregas, A. Martinez and G. Caire, *Bit-Interleaved Coded Modulation*, Foundations and Trends in Communications and Information Theory, Vol. 5:No 1-2 pp. 1-153, 2008.

[9] *GNU Radio*, http://www.gnu.org/software/gnuradio/.