

Superposition Coding Strategies: Design and Experimental Evaluation

S. Vanka, *Student Member, IEEE*, S. Srinivasa, Z. Gong, *Student Member, IEEE*, P. Vizi, K. Stamatiou, *Member, IEEE*, and M. Haenggi, *Senior Member, IEEE*

Abstract—We design and implement a software-radio system for Superposition Coding (SC), a multiuser transmission scheme that deliberately introduces interference among user signals at the transmitter, using a library of off-the-shelf point-to-point channel codes. We experimentally determine the set of rate-pairs achieved by this transmission scheme under a packet-error constraint. Our results suggest that SC can provide substantial gains in spectral efficiencies over those achieved by orthogonal schemes such as Time Division Multiplexing. Our findings also question the practical utility of the Gaussian approximation for the inter-user interference in Superposition-Coded systems.

Index Terms—GNURadio, superposition coding, software-defined radio, universal software radio peripheral.

I. INTRODUCTION

A. Motivation and Prior Work

THE problem of communicating with many receivers arises in many “downlink” scenarios such as communication from an access point to stations in WiFi or from a base station in cellular systems. The conventional approach is to set up orthogonal channels to each user by time/frequency/code-division multiplexing. Although this approach eliminates interference between transmissions, it does not in general achieve the highest possible transmission rates for a given packet error rate (or reliability) [1]. In fact, Superposition Coding (SC) [2] is a well-known non-orthogonal scheme that achieves the capacity on a scalar Gaussian broadcast channel.

We motivate the use of SC for the two-receiver case. Consider a cellular downlink with several active users. Given the user density in typical networks, it is always possible to pick two 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 base station (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 idea behind SC is to optimally exploit this channel ordering.

A BS that uses two-receiver SC can transmit *superimposed* F and N packets (or more precisely, the far and near user codewords) in *both* F’s and N’s time slots (see Fig. 1). Thus both links enjoy the *combined* degrees of freedom available to N and F, while sharing the transmit power. For large blocklengths, it can be shown that it is possible to encode F’s packets such that they can be decoded in the presence of interference from N’s packets. Since N has a stronger link to the BS, N can replicate this step to regenerate and thereby cancel F’s signal from its received signal. It can then decode its own packet. This is the well-known *Successive Decoding* (SD) or *Successive Interference Cancellation* (SIC) procedure [1].

We can extend this two-user scheme to any number of users. In fact, SC (combined with SD) achieves the capacity on a scalar Gaussian broadcast channel. This implies that *any* TD-achievable rate-pair (i.e., the pair of *spectral efficiencies* on a Gaussian channel) can also be achieved using SC, with the rate gain over TD increasing with the disparity in the user link qualities.

While information theory sufficiently motivates the use of SC, it is largely silent on practical issues such as finite block length codes, finite encoding and decoding complexity, hardware non-idealities (e.g., carrier frequency offset, phase noise) that one would encounter while designing such a system. This motivates the experimental study of SC.

For rapid prototyping and streamlining the design effort, we adopt a software-defined radio (SDR) [3] paradigm using the well-known open-source GNU Radio platform in conjunction with the Universal Software Radio Peripheral (USRP) hardware board that serves as an analog and RF front-end [4]. A well-known prototyping system [5], [6], it has been recently used in testbed design, including UT Austin’s Hydra [7] and by Bell Labs and Microsoft Research [8].

B. Main Contributions

Building on our previous work [9], the main contributions of this work are as follows:

- 1) We propose a design technique for SC using a finite library of finite-blocklength point-to-point codes developed for finite constellations.
- 2) We design a system architecture and implement it on GNURadio/USRP and determine the experimentally-achieved set of spectral efficiencies for a packet-error constraint. To the best of our knowledge, ours is the first

Manuscript received August 30, 2011; revised January 1 and April 2, 2012; accepted April 14, 2012. The associate editor coordinating the review of this paper and approving it for publication was M. Ardakani.

S. Vanka is currently with Broadcom Corporation (e-mail: svanka@nd.edu). S. Srinivasa is currently with LSI Corporation (e-mail: srrniv1@alumni.nd.edu).

Z. Gong and M. Haenggi are with the Emerging Wireless Architectures Lab, Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556, USA (e-mail: {zgong, mhaenggi}@nd.edu).

P. Vizi is with Morgan Stanley (e-mail: peter.vizi@gmail.com).

K. Stamatiou is with the Department of Information Engineering, University of Padova, Italy (e-mail: kstamat@dei.unipd.it).

Digital Object Identifier 10.1109/TWC.2012.051512.111622




Fig. 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.

attempt at systematically designing and characterizing an SC physical layer that, along with its accompanying hardware, forms a functioning system.

- 3) We study the implications of using a finite constellation along with a demodulate-and-decode receiver architecture on the statistics of the interference-plus-noise term.

C. Paper Overview

The remainder of the paper is organized as follows. In Section II we briefly summarize how SC achieves capacity and discuss some implications of restricting the library of codes to a finite set of finite-blocklength codes. In Section III, while retaining the rate-centric approach, we propose a design technique for SC with such a finite code library and specialize this technique to a library comprised of a well-known family of codes designed using the Bit Interleaved Coded Modulation (BICM) technique [11], and predict the theoretically achievable rate region. In Section IV, we describe the system architecture that uses these BICM codes to implement SC. In Section V, we present an experimental setup that emulates a Gaussian BC and use it to experimentally determine the achievable spectral efficiency pairs for a two-receiver BC under a packet-error constraint. The resulting rate region is the finite-library analog of the information-theoretic rate region. We also discuss some practical issues that arise in the design of superposition-coded systems, including the validity of treating inter-user interference as Gaussian noise. In Section VI we conclude the paper and suggest possible avenues for future work.

II. SUPERPOSITION CODING: FROM THEORY TO PRACTICE

We will briefly summarize relevant results from [1], [2] on achieving the capacity of the (scalar) Gaussian Broadcast Channel (BC) using SC with SD. In addition to making this paper self-contained, this discussion identifies the key architectural building blocks of a superposition-coded system. A closer examination of the blocks allows us to identify some key practical issues in implementing this ideal scheme. We use calligraphic fonts (e.g., \mathcal{C}) to represent sets and sans-serif fonts (e.g., $f(\cdot)$) to denote the encoding/decoding maps. Also, we use $[M]$ to represent $\{1, \dots, M\}$ for $M \in \mathbb{Z}^+$, and occasionally use the short-hand Tx for a transmitter and Rx for a receiver.

A. Achieving the Capacity on the Gaussian BC

Consider a BS that wants to communicate with two receivers N and F. The broadcast nature of the wireless medium is captured by the broadcast channel model $X \rightarrow (Y_N, Y_F)$ where X denotes the channel input and Y_N and Y_F are the channel outputs at N and F¹. Let $(X(n))$ be a sequence of channel inputs indexed by the channel use $n \in [L]$. Clearly $(X(n))$ must encode information relevant to each user. The *capacity region* of this channel is the closure of the set of all possible pairs of transmission rates at which the BS can reliably send two independent information streams, one each to N and F (allowing $L \rightarrow \infty$).

For a Gaussian BC, we have

$$Y_N(n) = h_N X(n) + Z_N(n); \quad Y_F(n) = h_F X(n) + Z_F(n) \quad (1)$$

where N (resp. F) has a complex channel gain h_N (resp. h_F) and $Z_{u, u} \in \{N, F\}$ denote the WGN processes. We assume the BS operates with an average power constraint P [W] and a (baseband) bandwidth W [Hz], and denote the noise power spectral density by N_0 [W/Hz]. From the above, the power constraint per channel use is P/W and $\mathbb{E}[|Z(n)|^2] = N_0 W$. From the definition of N and F, $|h_N|^2 > |h_F|^2$.

One way for $(X(n))$ to encode information is to communicate with each user in turns by partitioning the total number of channel uses into time slots (as in TD). For a given n , $X(n)$ contains information pertaining to just one user. This is the well-studied point-to-point communication problem, for which good practical encoding and decoding schemes exist. However, for a BC it is known that TDM is suboptimal in general; the root cause lies in its inability to fully exploit the fact that $|h_N| > |h_F|$: N has a “stronger” channel to BS, and hence can always decode information that can be decoded at F. This makes the scenario ideal for the SC scheme which achieves every pair of transmission rates in the capacity region.

The key architectural elements of an SC system are:

- 1) A *superposition encoder* f that consists of
 - a) Two point-to-point encoders, $f_N: \{0, 1\}^{[2^{LR_N}]} \rightarrow \mathbb{C}^L$ (which we call the *near-encoder*) and $f_F: \{0, 1\}^{[2^{LR_F}]} \rightarrow \mathbb{C}^L$ (which we call the *far-encoder*), that map their respective inputs (the *near-* and *far-*messages) to complex-valued sequences $(X_N(n))$ and $(X_F(n))$, each of *block*

¹In practical terms, $X(n)$ can be understood as a (coded) symbol stream from the BS, and the Y 's as the corresponding noisy and/or distorted observations of this symbol stream at N and F.

length L . Here R_N and R_F denote the bandwidth-normalized transmission rates (or spectral efficiencies) of N and F (the *near-* and *far-rates* for short).

- b) A summation device that outputs a sequence

$$X(n) = \sqrt{1-\alpha} X_F(n) + \sqrt{\alpha} X_N(n), \quad (2)$$

where a fraction $\alpha \in [0, 1]$ of the power is assigned to N (the *near-fraction* for short).

- 2) A *single-user* decoder $g_F : \mathbb{C}^L \rightarrow \{0, 1\}^{2^{L R_F}}$ that estimates the far packet from the observations $(Y_F(n))$ by treating $(X_N(n))$ as Gaussian noise.

- 3) A *successive cancellation* decoder $g_{F,N} : \mathbb{C}^L \rightarrow \{0, 1\}^{2^{L R_N}}$ that is used to recover N's packet in the following steps:

- a) Decode F's packet using the single-user decoder g_F .
 b) Cancel $\sqrt{1-\alpha} h_N X_F(n)$ from $Y_N(n)$ by regenerating $X_F(n)$ using the far-encoder f_F and the knowledge of h_N and α :

$$\begin{aligned} Y'_N(n) &= Y_N(n) - h_N \sqrt{1-\alpha} X_F(n) \\ &= h_N \sqrt{\alpha} X_N(n) + Z_N(n). \end{aligned} \quad (3)$$

- c) Decode N's packet using the single-user decoder $g_N : \mathbb{C}^L \rightarrow \{0, 1\}^{2^{L R_N}}$.

It is well known that as $L \rightarrow \infty$, for all α there exist f_N, f_F, g_F, g_N such that communication can occur arbitrarily reliably for all pairs of transmission rates satisfying

$$\begin{aligned} R_N &< W \log_2(1 + \alpha \gamma_N) \\ R_F &< W \log_2 \left(1 + \frac{(1-\alpha)\gamma_F}{\alpha\gamma_F + 1} \right) \end{aligned} \quad (4)$$

for $\alpha \in [0, 1]$, where

$$\gamma_N \triangleq \frac{P|h_N|^2}{N_0 W} \quad \text{and} \quad \gamma_F \triangleq \frac{P|h_F|^2}{N_0 W}$$

represent the *near-* and *far-SNRs* respectively. We are interested in the spectral efficiencies

$$r_u = R_u/W, \quad u \in \{N, F\} \quad (5)$$

which we will simply call *rates*. Clearly, making α discrete would also make the rate region boundary discrete. The following subsections elaborate on this issue.

B. Practical Design Issues

As noted above, a discrete α results in a discrete set of rate-pairs defining the corner points of the achievable rate region². The lowest (resp. highest) value of α restricts the minimum (resp. maximum) power that can be assigned to a user for superposed transmission. For finite block lengths there is also a non-zero probability of decoding error.

With these practical constraints factored in, the rate benefits from SC over TD will depend on α , the chosen discrete set of codes and the system implementation. Using a combination of theory, simulations and experiments, we show that it is indeed possible to build efficient superposition-coded systems using off-the-shelf single-user coding and decoding techniques.

²Time-sharing can be used to convexify this boundary

III. DESIGNING A SUPERPOSITION-CODED SYSTEM

Based on our observations in the previous section, we first describe our design approach in its full generality in Section III-A. In Section III-B we illustrate this approach for a well-known single-user coded modulation technique known as Bit-Interleaved Coded Modulation (BICM). Indeed, good practical codes designed using BICM techniques form the basis for error correction in many real-world wireless networks (see, e.g., [11] and the references therein).

A. Practical SC using a Finite Code Library

We assume all codes have a *block length* $L < \infty$ [channel uses] and fix a *target Packet Error Rate (PER)* $\epsilon \ll 1$. The latter is the probability that a user cannot decode its packet³. Define a *code library* as a collection \mathcal{C} of $M > 1$ *single-user* encoder-decoder function pairs⁴ $(f(\cdot), g(\cdot))$ (a “code” for short). These codes are ordered by their rates $r_1 < r_2 < \dots < r_M$ [bps/Hz]. We label each code by its *rate index* $i \in [M]$. We say a *rate index* i is (ϵ -)feasible on a link if the receiver can decode a packet encoded at rate r_i with a PER no larger than ϵ . In the following, denote the user rate, the corresponding rate index and the PER by r_u, i_u and PER_u for $u \in \{N, F\}$.

Consider a BS communicating with N alone. When assigned the full BS power, we assume this link has an SNR γ_N that allows the largest rate in the code library to be feasible⁵. Since lower rates require a smaller link SNR to remain feasible, the BS can backoff from this full power to support these rates. We define $\beta(l) \leq 1$ to be the *minimum* fraction of the full power for a rate index l to remain feasible at N, $l \in [M]$. When communicating with N at rates $r_l < r_M$, its link requires only a fraction $\beta(l)$ of the full BS power. Therefore, without violating the power constraint, the BS can superpose a signal with fraction of the total power $1 - \beta(l)$ to communicate with another user, assuming that N can perfectly cancel this interfering signal before decoding its own message⁶.

We use this superposed signal to transmit F's packets. The link to F has an SNR $\gamma_F < \gamma_N$ when assigned the full BS power. Suppose this allows rates no larger than $r_K, K < M$ to be feasible with full BS power. Clearly γ_F must be chosen with some care. If the far-link is too noisy (γ_F is too small, perhaps because F is too far away from the BS), even a small amount of interference from N's signal (from superposition) can render even r_1 infeasible. If $\gamma_F \lesssim \gamma_N$, there is not enough disparity between the links to N and F to take full advantage of signal superposition⁷. For comparison when BS serves N and F using TD, it uses rate index M to serve N and rate index K to serve F. In this case, the rate-pair is determined by the fraction of time slots assigned to each link.

³We assume one packet is encoded as one codeword. Hence packet error is equivalent to codeword error.

⁴In practical terms, an encoding function is a mapping from the packet bits to the signal waveforms induced by the encoder and the modulator. A decoding function is the detection rule to perform the inverse operation, and is induced by the demodulator and decoder.

⁵This models a user located close to the BS.

⁶In Section V-C2, we find that this cancellation, although not perfect, does not require increasing N's signal power significantly beyond $\beta(l)$ for rate-indices $l \in [M]$.

⁷We study the implications of a poor choice of F in Section V-C1.

Given a library \mathcal{C} we would like to find the set of largest simultaneously achievable rates at N and F. Since \mathcal{C} has only M possible rate indices, all of which are feasible at N, we only need to choose (at most) M values of power assignments to N. For each near-user rate r_l , $l \in [M]$, we maximize the far-user rate r_F such that the PER constraints at N and F are simultaneously satisfied. Given the near-user rate r_l , this can be stated as the optimization problem

$$\max_{\{r_1, \dots, r_M\}} r_F \quad \text{s.t.} \quad r_N = r_l, \text{PER}_N(l) \leq \epsilon, \text{PER}_F(l) \leq \epsilon. \quad (6)$$

Clearly, the feasibility of any far-user rate depends on the allocation of power to N and F. Now, in theory, if interference from F's signal could be cancelled perfectly at N, the power assigned to N could be chosen just large enough to sustain r_l , i.e., no more than $\beta(l)$. However, in practice, due to imperfect cancellation, this power should be slightly larger. Hence, for each l , we could start by assigning a fraction $\beta(l)$ of the power to N (and $1 - \beta(l)$ to F) and increase the power assigned to N step-by-step until the PER constraint at N is satisfied (this procedure can be made more efficient, as explained in Sec. V-B). By this procedure, we determine the smallest near-user power fraction $\alpha(l)$ that ensures that a rate r_l is feasible at N when F is assigned a power $1 - \alpha(l)$. For this power assignment, we now increase the far-user rate until the PER constraint at F is satisfied. We denote by $i_F^*(l)$ the rate index of the solution to (6). Assuming small levels of residual interference after cancelling F, $\alpha(l) \gtrsim \beta(l)$, $l \in [M]$ ⁸.

The solution set of (6) is the set of optimal rate pairs $\{(r_l, r_{i_F^*(l)}) : l \in [M]\}$. Combining the achievable endpoints $(r_M, 0)$ and $(0, r_K)$, we obtain the *rate region* of this code library $\mathcal{R}^*(\mathcal{C}; \epsilon, L, \gamma_N, \gamma_F) \triangleq \mathcal{R}^*$ as

$$\mathcal{R}^* \triangleq \text{Conv}(\{(0, 0), (r_M, 0), \{(r_l, r_{i_F^*(l)}) : l \in [M]\}, (0, r_K)\}). \quad (7)$$

where $\text{Conv}(\cdot)$ denotes the convex hull operator. Fig. 2 graphically summarizes the design procedure. In the following subsection we pose (6) for a code library consisting of point-to-point codes based on BICM.

B. SC using a BICM Code Library

Based on the framework developed so far, we revisit the key elements discussed in Section II-A to map them to specific subsystems in a Superposition-Coded-BICM (SC-BICM) system; we also introduce some terminology specific to BICM. Subsequently we solve (6) numerically for BICM codes in the high-reliability regime ($\epsilon \rightarrow 0$). For reliabilities of practical interest (e.g., $\epsilon \lesssim 0.1$), we solve this problem via Monte Carlo simulations for a specific class of BICM codes we have implemented on our testbed.

1) *The SC-BICM System*: The channel coding in a canonical BICM system [13] is specified by:

- c : Convolutional code with code-rate ρ_c , free distance d_c and number of free-distance error events w_c .
- x : A gray-coded constellation mapper that maps b_x interleaved coded bits mapped to each symbol. The symbols have unit average energy and a minimum intersymbol Euclidean distance of d_x^9 . We will refer to the range \mathcal{S}_x

⁸We find that this is indeed the case in Sec. V-C2

⁹The gain of the subsequent Tx stages controls the overall transmit power.




Fig. 2. A graphical illustration of the rate region obtained by the design process in Section III-A for a hypothetical code library with $M = 9$ codes. The single-user points (filled dark circles) are $(r_9, 0)$ and $(0, r_K)$, for some $K \leq 9$. The gray dashed line joining these two points represents the rate region achievable by TD. The open circles represent the optimal rate-pairs obtained by solving (6) as N's rate is varied from r_1 through r_8 . To obtain \mathcal{R} using (7) we find the convex hull of this optimal set (along with the single-user points and the origin) to obtain a convex polygon with vertices $\{(r_l, r_{i_F^*(l)}) : k \in \{2, 5, 6, 8\}\}$ inside the positive quadrant. The solid black lines represent the rate region boundary in this quadrant.

of x as its *constellation*. For example, the constellation of a BPSK mapper is $\{-1, 1\}$.

- Π : An interleaver matrix that specifies the order in which code bits are read by the constellation mapper. The corresponding de-interleaver is denoted by Π^{-1} .

Combining a convolutional code c with a modulator x results in a spectral efficiency of $r = \rho_c b_x / W$ [bps/Hz]. Thus the near-encoder f_N is the composition $x_N \circ \Pi_N \circ c_N$. A block length L can encode $B = Lr$ information bits per codeword or packet. With a slight abuse of notation, we denote the Viterbi decoder and the constellation demapper by c^{-1} and x^{-1} respectively.

We can now describe the key functional units in an SC-BICM system: the superposition encoder f finds the weighted sum of the outputs of two separate BICM encoders f_N and f_F to construct X which we will now call the *composite symbol stream*. The composite symbol stream is mapped to a waveform and transmitted over the wireless medium. At F, Y_F in (1) represents the noisy observations of X seen by the constellation demapper (that is the first stage in g_F). From Y_F the demapper estimates the reliabilities of each encoded bit of the far-packet, treating the symbol stream $(X_N(n))$ as interference. The remaining steps are the same as in standard BICM decoding.

Since g_F is part of the successive cancellation decoder $g_{F,N}$, these steps are reproduced at N; thereafter the far-encoder f_F (at N) reconstructs X_F from this decoded packet which is then subtracted from Y_N to yield Y'_N as in (3). In the final step g_N estimates the near-packet from Y'_N using standard BICM decoding.

2) *A Theoretical Estimate of the Rate Region of an SC-BICM System in the High Reliability Regime*: Given \mathcal{C} , (6) can be solved by checking the feasibility¹⁰ of every candidate rate pair via (time-consuming) Monte Carlo simulation. One could

¹⁰To any desired confidence interval.

reduce this computational overhead by reducing the search space of possible far-rates for a given near-rate. In fact, a formula to compute the PER for a given rate pair and near-fraction would obviate the need for simulations.

Unfortunately, accurate formulas for the PER are difficult to obtain even for the point-to-point case, although there exist well-known upper bounds that are asymptotically tight in the high-reliability regime ($\text{PER} \rightarrow 0$) [12]. In the high-reliability regime, these upper bounds can be suitably modified, as we will show.

The key difference between SC-BICM and point-to-point BICM lies in the far-demodulator x_F^{-1} that estimates the reliabilities of the far-code bits from observations of the form

$$Y_F(n) = \underbrace{h_F \sqrt{1 - \alpha} X_F(n)}_{\text{Signal}} + \underbrace{h_F \sqrt{\alpha} X_N(n) + Z_F(n)}_{\text{Perturbation}}, \quad n \in [L]. \quad (8)$$

Unlike in point-to-point BICM, the perturbation term is not Gaussian; its statistics depend on N's constellation, which we assume is known to the demodulator¹¹. Whether or not x_N is known, the finiteness of the interference constellation raises an interesting question about the validity of the Interference-As-Gaussian-Noise (IAGN) model that assumes such Gaussianity. We investigate this question in greater detail in Section V-C3.

When N's constellation is known, it is useful to treat each composite symbol $X(n) = \sqrt{1 - \alpha} X_F(n) + \sqrt{\alpha} X_N(n)$ as a member of a superconstellation with $2^{b_{x_N} + b_{x_F}}$ points (see also Fig. 7). Viewed from the demodulator, interference perturbs each 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 size and shape of this cluster depends on the interferer's constellation. Thus a maximum-likelihood demodulator interested only in the far-packet 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 the demodulator refine its reliability estimate of each detected code bit of F. Analogous to the single-user case, the reliability of the k^{th} bit in the n^{th} symbol can be approximated using the max-log-MAP approximation:

$$L_n^{(k)} \approx N_0 \min_{s \in \sqrt{1 - \alpha} \mathcal{S}_{x_F}^{(k-)} \times \sqrt{\alpha} \mathcal{S}_{x_N}} |Y_F(n) - h_F s|^2 - \min_{s \in \sqrt{1 - \alpha} \mathcal{S}_{x_F}^{(k+)} \times \sqrt{\alpha} \mathcal{S}_{x_N}} |Y_F(n) - h_F s|^2 \quad (9)$$

where $\mathcal{S}_{x_F}^{(k-)} \subset \mathcal{S}_{x_F}$ and $\mathcal{S}_{x_F}^{(k+)} \subset \mathcal{S}_{x_F}$ comprise symbols whose k^{th} bits are 0 and 1, respectively.

In the high-reliability regime, the dominant error events in such a demodulator are events where a daughter point is incorrectly identified with a neighboring cluster. Analogously to the point-to-point case (when these "clusters" are just points), the probability of these error events is controlled by the effective cluster separation $d_{\text{eff}} \equiv d_{\text{eff}}(\alpha, x_F, x_N)$ given by

$$d_{\text{eff}} \triangleq \min_{\substack{p_1, p_2 \in \sqrt{1 - \alpha} \mathcal{S}_{x_F} \\ p_1 \neq p_2}} \min_{d_1, d_2 \in \sqrt{\alpha} \mathcal{S}_{x_N}} |p_1 + d_1 - p_2 - d_2|. \quad (10)$$

¹¹In practice sending this information entails a small overhead, which we neglect in this paper.

Here the superscript (F,N) on the left hand side emphasizes that the parent points are drawn from x_F and the interferer is drawn from x_N . Using arguments similar to those in [12], [13], PER_F can be approximated as

$$\text{PER}_F \approx B_F w_{c_F} Q \left(d_{\text{eff}}^{(F,N)} \sqrt{\frac{d_{c_F} \gamma_F}{2}} \right) \quad (11)$$

in the high-reliability regime, where $Q(x)$, $x \geq 0$ is the Q-function. Assuming perfect cancellation at N,

$$\text{PER}_N \approx B_N w_{c_N} Q \left(d_{x_N} \sqrt{\frac{d_{c_N} \gamma_N}{2}} \right), \quad (12)$$

which can be plugged into (12) to approximate PER_N .

However, for PERs of practical interest (say $\text{PER} \lesssim 0.1$) these Q-function bounds are too loose to predict the correct solutions to (6). In this regime we use these bounds as estimates that reduce the search space of the achievable far-rates, and refine them further via simulation. In the following subsection we illustrate these ideas with a design example. The code library chosen in the example is the same as the one used in our system design.

C. A Design Example

We use a BICM code library with a decoder structure explained in Section III-B2. These BICM codes were implemented in our testbed in a point-to-point setting.

The library consists of all possible pairings of 4 convolutional codes $\{c^{(1)}, c^{(2)}, c^{(3)}, c^{(4)}\}$ with three constellation mappers. There is no interleaving¹². Table I summarizes the details of these convolutional codes. Note that $c^{(2)}, c^{(3)}$ and $c^{(4)}$ are obtained by appropriately puncturing $c^{(1)}$, which is the standard rate-1/2 constraint length 7 convolutional code with the generator matrix [133, 171]. The constellations are $x \in \{\text{BPSK}, \text{QPSK}, \text{16QAM}\}$. These are all QAM constellation with even b_x , so $d_x = \sqrt{6/(2^{b_x} - 1)}$ [10]. With this code library the available set of spectral efficiencies is the sequence $(r_1, r_2, \dots, r_{12})$ obtained by ordering the elements of the set $\{1, 2, 4\} \times \{\frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \frac{5}{6}\} = \{\frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \frac{5}{6}, 1, \frac{4}{3}, \frac{3}{2}, \frac{5}{3}, 2, \frac{8}{3}, 3, \frac{10}{3}\}$ in ascending order.

We now solve (6) via Monte Carlo simulation (using the simplified procedure outlined in Section) for $\epsilon = 0.1$ and $L = 1536$ with perfect receiver CSI and perfect interference cancellation. For our implementation this choice of L strikes a balance between code performance and implementation constraints. Besides illustrating the design procedure for a concrete example, the far-rates from this simulation provide upper bounds for a practical system (where neither of these conditions holds). The simulation procedure closely follows the experimental procedure described in Section V-B. The first step involves obtaining the single-user PER curves similar to those in Fig. 5 by simulating a link operating over a point-to-point Gaussian channel. As with the experiment, these results yield $\{\beta(l)\}$, $l \in [M]$ that are given to another Matlab program that simulates a Gaussian BC and numerically

¹²Code performance can be optimized by suitably tailoring the interleaver Π . However its absence does not change the main message of this example (or indeed, that of the paper), which is to show substantial gains from SC even for finite blocklengths and constellations.

TABLE I
KEY PARAMETERS OF THE CONVOLUTIONAL CODES IN THE CODE LIBRARY.

Code	Rate ρ_c	Free Dist. d_c	#Free Dist. Error Events w_c
$c^{(1)}$	1/2	10	11
$c^{(2)}$	2/3	6	1
$c^{(3)}$	3/4	5	8
$c^{(4)}$	5/6	4	14




Fig. 3. Optimal rate pairs (solution to (6)) for $\gamma_N = 15$ dB and $\gamma_F = 8$ dB for the library of $M = 12$ codes described in Section III-C. The TD rates are obtained by time-sharing between the single-user operating points. The values of α at corner points A-E are 1, 0.2138, 0.1259, 0.0631, 0 respectively.

solves (6) using the procedure described in Section V-B. Fig. 3 shows the simulated rate region for $\gamma_N = 15$ dB and $\gamma_F = 8$ dB. For these parameters, it takes about 4 hours to complete this procedure on Matlab version 7.13 running on a dual-core Linux workstation running at 2.4GHz and with a 2 GB RAM.

IV. IMPLEMENTING A SUPERPOSITION-CODED SYSTEM

A. The Platform

We implemented all physical layer processing steps at BS, N and F by suitably modifying an existing point-to-point wireless testbed. The testbed uses Orthogonal Frequency Division Multiplexing (OFDM), and its design parameters are similar to those in the IEEE 802.11a standard. The testbed runs on GNU Radio (revision 10923) on a Linux PC. GNURadio provides driver functions that interface the PC with the USRP board that functions as the analog frontend and the RF.

B. Packet Structure

Transmissions occur in *frames*. As in WiFi, each frame consists of a preamble followed by a header and a payload. The preamble assists the receiver in frame acquisition and channel estimation. The header encodes the BICM code type and α . The preamble and the header designs are left unchanged

TABLE II
PARAMETERS USED IN THE EXPERIMENT.

Center Frequency	903 MHz
Message Bandwidth	2 MHz
Modulation	16-tone OFDM (8 data, 4 pilot, 4 null)
CP Length	1 μ s

from the single-user case¹³. The only “multiuser” section in the frame is the payload. For lack of space, we will not provide more details here; they can be found in our technical report [16].

The block of $L = 1536$ coded symbols is transmitted over 192 OFDM symbols with 8 subcarriers for the payload in each OFDM symbol. We retain the 4 pilot subcarriers used by the testbed for frequency and phase tracking and 4 null subcarriers for spectral shaping. This brings the total number of subcarriers to $8 + 4 + 4 = 16$. The message bandwidth of 2 MHz is limited by the USRP, and the cyclic prefix is made commensurate with the (relatively) flat frequency response of the channel for this bandwidth. We summarize the system parameters in Table II.

C. Single-User Characterization

Creating experimental conditions that ensure a time-invariant wireless propagation loss can be complicated: although the indoor radio channel and the USRP boards have a reasonably flat response over a 2 MHz bandwidth, the propagation loss is quite sensitive to changes in the environment (e.g., those caused by motion). One approach would be to compensate for such changes by appropriate power control at the BS. Doing so would require SNR feedback from the users on a control channel; designing such feedback links would be worth the effort only if SC could provide substantial rate-gains in a *perfectly power-controlled* environment. We focus on the latter question in the paper.

While manual power adjustments work at smaller time scales lasting a few minutes [18], they are quite cumbersome for longer experiments (e.g., measuring the PER for all the rates in the library, or solving (6), see Section V-A).

To circumvent this problem, we emulate perfect power control by connecting the Tx and Rx with a coaxial cable (see Fig. 4), resulting in controllable experimental conditions and reproducible outcomes¹⁴. With this coaxial setup, we measure the point-to-point PER as a function of the transmit power P (which implicitly determines the SNR γ) for each rate index in the code library using parameters from Table II.

Starting from a value of P chosen such that the received power (at the antenna port, as measured by a spectrum analyzer) is at a fixed level above the theoretically predicted thermal noise floor, we change the transmit power in 1 dB

¹³Consequently, known techniques (e.g., the use of windowed correlators as in [15]) and low-rate header encoding (as in IEEE 802.11a) can be used.

¹⁴Indeed, the same practical considerations motivate the use of channel emulators.




Fig. 4. The setup used to approximate a Gaussian BC. The USRP boards are connected via cables. A splitter is used to split the transmitted signal to the two receivers, while an attenuator is used to (virtually) create the presence of a far user.

steps and measure the PER for each rate index¹⁵. For comparison we overlay the PER plots for a simulated Gaussian system with the same theoretical noise floor but without additional sources of noise and distortion in Fig. 5. For a rate index l , let $P_{\min}(l; \epsilon) \equiv P_{\min}(l)$ be the smallest experimentally obtained power level that makes a link ϵ -feasible for rate index l . Denote the corresponding power level by $\tilde{P}_{\min}(l)$ on the simulated Gaussian link.

Observe that the slopes of the ideal and experimental waterfall curves are similar up to $\approx 1\%$ PER. At $\approx 10\%$ PER, the combined non-ideality of the hardware and implementation result in a maximum power loss of $10 \log \frac{P_{\min}(l)}{\tilde{P}_{\min}(l)} \leq 3.5 \text{ dB}$, $i \in [12]$ ($= 3 \times 4$) from the ideal results. Thus we take this “coaxial channel” to be a reasonable approximation of a Gaussian channel in our experiments.

Using these results we obtain the smallest feasible near-fractions $\beta(l)$ for N as

$$\beta(l) = \frac{P_{\min}(l)}{P_{\min}(12)}, \quad l \in [12]. \quad (13)$$

We use $(\beta(l))$ as starting points to find $\alpha(l)$ in the rate region experiment in Section V-B.

D. Modifications to Implement an SC-BICM System

The SC encoder f is implemented as two instances of a point-to-point BICM (sub-)encoder and a combiner. The encoded data is then transmitted via a standard OFDM modulator. The receiver is implemented as a single GNU Radio signal processing block, and contains a successive decoding block to decode the near user packet. More details on the transceiver operation including their block diagrams may be found in our technical report [16].

¹⁵This mimics a system with a 1 dB granularity in power control. Also note that the effective signal distortion seen at the receiver is usually higher than that due to thermal noise alone, due to additional sources of noise and distortion such as imperfect receive implementation and hardware imperfections. In our experiments an initial power level of 20 dB above the theoretical value was found to be adequate.




Fig. 5. Single-user PER versus SNR, along with the 95% confidence intervals for all the 12 rates in the library. The solid lines depict the experimental results; the ideal curves for BPSK-1/2 and 16QAM-5/6 are also shown (dashed lines). Observe that at 10% PER the implementation loss is in the range 2dB – 3.5dB.

E. The Runtime Requirement

The GNU Radio runs on a general purpose computer. In our experiment, we use the desktop with two quad-core Intel Xeon CPU E5520@2.27 GHz with 8 GB RAM. The transmitter code takes approximately 2.1 ms for the transmitter to generate a single packet. For the far receiver, the decoding procedure takes approximately 8.7 ms. For the near receiver, it needs approximately 19.7 ms to decode its packet. The near user needs to do the decoding twice and encoding once so that the complexity is more than doubled compared to the far user. The time is measured by recording the timestamps before and after encoding(decoding) a packet. The near user needs to do the decoding twice and encoding once so that the complexity is more than doubled compared to the far user. In the experiment, the transmitter has an idle gap of 100 ms between two packets. Hence the receiver has enough time to decode. More detailed discussion about implementation complexity is available in [16].

V. EXPERIMENTAL RESULTS

A. Emulating a Gaussian BC

The time variation of the propagation loss is a bigger problem for the BC because of the presence of *two* Tx-Rx paths in the BC as opposed to just one in the point-to-point setting of Section IV-C. Moreover, checking the feasibility of candidate rate-pairs while solving (6) requires repetitive PER measurements. Thus we emulate a Gaussian BC using a combination of coaxial cables, a splitter, and an attenuator bank, as shown in Fig. 4. The USRPs shown communicate with three Linux PCs via USB 2.0. The PCs are configured to run the appropriate Tx or Rx code in GNURadio.

The second step is to fix the single-user SNRs γ_N, γ_F . In our experiment we fix them implicitly by choosing a transmit power P_N and an attenuator setting a_F ¹⁶. In the single-user

¹⁶It is possible to obtain an estimate of γ_N by measuring the signal and noise powers in the digital domain (see [16]).

mode, P_N is chosen to be $P_{\min}(12)$ (plus some additional loss due to the splitter) using the single user results from Section IV-C. The measured γ_N was found to be ≈ 18 dB (see [16] for details on SNR measurement). With this value of P_N , we set F in the single-user mode to operate at a desired rate (e.g., QPSK, rate-5/6), and increase the attenuation in $\eta = 1$ dB steps until the far link violates the PER constraint. γ_F is the receive SNR for the largest attenuation setting that supports a 10% (or lower) PER at F (for QPSK rate-5/6, this was found to be ≈ 10 dB). With this approximation to the Gaussian BC, we are now ready to obtain the rate region for this system by solving (6) experimentally.

B. The Rate Region Experiment

Given P_N, a_F and a code library \mathcal{C} with rate-indices $l \in [M]$, we use the above experimental setup to solve (6) to find $(l, i_F^*(l))$ using the following procedure:

- 1) Initialize:
 - a) $\alpha(l) = \beta(l)$ where $\beta(l)$ is calculated from (13).
 - b) $i_F(l) = \arg \max_l \{P_{\min}(l) < (1 - \beta(l))a_F P_N\}$.
- 2) Calculate the PER for $(l, i_F(l))$ for stream weights $\sqrt{\alpha(l)}, \sqrt{1 - \alpha(l)}$.
- 3) If
 - a) (feasible at N) AND (feasible at F): set $i_F^*(l) = i_F(l)$, $k \mapsto l + 1$, go to Step 1.
 - b) (feasible at N) AND (infeasible at F): $i_F(l) \mapsto i_F(l) - 1$, go to Step 2.
 - c) (infeasible at N) AND (feasible at F): $\alpha(l) \mapsto \alpha(l) \times 10^{\frac{\eta}{10}}$, go to Step 2.
 - d) (infeasible at N) AND (infeasible at F): $\alpha(l) \mapsto \alpha(l) \times 10^{\frac{\eta}{10}}$, $i_F(l) \mapsto i_F(l) - 1$, go to Step 2.

Here η is the step size for the α parameter (in this paper, we set $\eta = 1$ dB).

C. Results

We use the procedure in Section V-B to study three interesting problems: (a) How does the measured rate region change with a_F (i.e., the far-link is made stronger or weaker)? (b) how much does imperfect interference cancellation at N affect the rate region?, and (c) what are good models to account for N's interference at F, and, in particular, how useful is the popular Interference-As-Gaussian-Noise (IAGN) model in predicting its impact? We discuss these problems in the following.

1) *Changing the Strength of the Far-Link:* To study this problem, we find the rate region for two possible far-link SNRs: $\gamma_F = 5$ dB and $\gamma_F = 10$ dB, which correspond to single user rate indices $K = 3$ and $K = 8$ respectively. The near-user SNR for both cases is kept at $\gamma_N = 18$ dB. These scenarios are emulated by using suitable attenuator values $a_F = 9, 4$ respectively¹⁷. The results are shown in Fig. 6. Here, we used a transmit power $P_N = -43$ dBm and step size $\eta = 1$ dB. We clearly see dependence on the choice of γ_F . With $\gamma_F = 5$ dB, there is not enough disparity between the near- and far-links to fully benefit from superposition (F is “too close” to BS). This is in fact predicted by theory [1]. On the other hand, we

see the effect of a finite code library when $K = 3$ (F is “too far” from BS): since its single-user rate is too small to begin with, interference from N's symbols rapidly degrades its link quality so as to make any far-rate infeasible.

Therefore, the far-user modulation and rate pair may be appropriately chosen based on the rate that the near user's traffic demands. For instance, when the near user's spectral efficiency is 1 (QPSK-1/2), choosing BPSK-3/4 for the far user provides a rate gain of about 28% over TD (as compared to a gain of about 21% over TD for QPSK-5/6), while when the near user's spectral efficiency is 3, it is preferable to choose QPSK-5/6 for the far user (over BPSK-3/4).

2) Impact of Imperfect Interference Cancellation at N:

The deviation of $\alpha(l)$ from its ideal value $\beta(l)$ is a measure of the residual far-user interference seen at N due to the imperfect cancellation of F's symbols (even when the far-packet is decoded correctly). The β 's are calculated from the single-user results using (13). The α 's are determined from the experiment. For $K = 8$, we find that when $l = 6$, $\alpha(l) = 0.13 \geq \beta(l) = 0.1$. At small near-rates the desired symbol stream (of near-symbols) has a much lower power than the interference symbol stream (of far-symbols). In this regime, even small regeneration errors manifest as large residual interference, necessitating an increase of $\alpha(l)$ beyond $\beta(l)$. For example, $1 - \beta(1) \approx 90\%$ of the transmit power is assigned to F. Even if only 10% of this power remains after cancellation, it is still about the same as the signal power. As $\beta(l)$ increases, so does the near-rate, thereby making the near-link susceptible to even smaller levels of residual interference, as well.

The root cause of this problem lies in the small (but non-zero) estimation error in the channel frequency response and a small error in compensating the carrier frequency and phase offsets. Although this level of inaccuracy may result in relatively small losses in a single-user system (as shown in Fig. 5), a multiuser system is much less tolerant to these errors as our results show. Despite this inaccuracy, we find that SC can still provide rate gains using reasonably well-designed single-user building blocks.

3) Modeling the Near-User Interference at the Far User:

We study the performance of F's (maximum-likelihood) demodulator for the three choices of the interfering (N's) signal's constellations—BPSK, QPSK, 16QAM—and for two different interferer strengths $\alpha = 0.2$ and $\alpha = 0.8$. F's rate is maintained at BPSK-1/2. Indeed, as explained in Section III-B2, the performance of the demodulator is dictated by the superconstellation with $2^{b_N + b_F}$ points; its error probability critically depends on the effective minimum distance (10). Fig. 8 depicts the far user PERs versus the SINR at F which we define as

$$\text{SINR} \triangleq \frac{(1 - \alpha)\gamma_F}{1 + \alpha\gamma_F}. \quad (14)$$

We observe the following:

- In the weak interference regime (for $\alpha = 0.2$, i.e., SIR ≈ 6 dB), it is seen that BPSK is the *worst interferer*. For small α , the error probability arises primarily from the separation between the clusters (see Fig. 7 (left)).
- In the strong interference regime (e.g., $\alpha = 0.8$, i.e., SIR ≈ -6 dB), it is seen that 16QAM is the worst

¹⁷Note that a_F is not simply equal to $\gamma_N - \gamma_F$ owing to the two different boards and cables used, which had to be calibrated separately.




Fig. 6. Experimentally obtained rate region for the library of $M = 12$ codes using the setup shown in Fig. 4 for two different choices of F . (Left) The values of α at corner points A-D are respectively 1, 0.126, 0.079, 0. (Right) The values of α at corner points A-C are respectively 1, 0.126, 0.




Fig. 7. Depiction of the superconstellation points in the weak (left) and strong (right) interference cases. Evidently, when α is small, the PER depends on the cluster separation, while when α is large, the PER is determined by the cluster density.




Fig. 8. Far user PER versus SINR for different near-user constellations in the weak (left) and strong (right) interference cases.

interferer. When α is large, the clusters overlap and the error probability may be attributed to the high density of the clusters (see Fig. 7 (right)).

Validity of the IAGN model: The rate constraint on the far-user in (4) is of the form $R = W \log(1 + \text{SINR})$. However, *conditioned on N's codebook*, this interference is clearly not Gaussian. When the channel codes and the detection process at F are allowed to be arbitrarily complex (e.g., allowing Gaussian signaling, ML decoding), the *combined* effect of interference and noise can indeed be shown to be Gaussian for $L \rightarrow \infty$ when the interference (N's signal) cannot be jointly decoded with the desired signal (F's signal) [17], which is certainly the case for the Gaussian BC. It is not clear if this remains valid when the choices of channel codes is restricted (e.g., to those with finite L using finite constellations) and/or the receiver architecture (e.g., to demodulate-and-decode). Under such constraints, it is of fundamental interest to check the validity of the Gaussian approximation. We clarify that the interference considered here is from N's symbols only, and not from other transmitters in the network.

Our results show that for SC-BICM systems with a demodulator structure described in Section III-B2, treating the interference-plus-noise term as Gaussian perturbation can be quite inaccurate. For a small near-user power (i.e., small α), the cluster centers are well separated and each daughter point is likely to be close to its parent (see Fig. 7 (left)). In this case, a BPSK interference distribution can place *every* daughter point at the farthest possible distance ($\sqrt{\alpha}$) from its parent. For Gaussian interference with the same power, 68% of all daughter points will lie *within* a ball of radius $\sqrt{\alpha}$ centered at the parent point¹⁸. Thus for a given noise level, Gaussian interference would result in fewer demodulation errors compared to BPSK interference: a Gaussian interference-plus-noise term *underestimates* the number of decoding errors.

For large near-user power, the cluster centers become closer and the clusters begin to overlap (see Fig. 7 (Right)), making the precise form of the interference distribution important. For a given interferer power, higher-order interferer constellations (such as 16QAM) result in more densely packed clusters, which upon overlapping result in a smaller effective minimum distance $d_{\text{eff}}^{(F,N)}$ from (10). Viewed from this perspective, it is clear that a Gaussian interferer would be the worst than either BPSK, QPSK, or 16QAM in that it results in infinitely dense clusters. Thus, in the low SIR regime, a Gaussian interference-plus-noise term *overestimates* the number of decoding errors. These trends are apparent from Fig. 8.

VI. CONCLUDING REMARKS

A. Summary

We have presented a software-radio implementation of Superposition Coding using off-the-shelf single-user coding and decoding blocks. We experimentally determine the set of achievable rate-pairs for this system under a packet error constraint. Our results suggest that SC can provide substantial gains in spectral efficiencies over those achieved by orthogonal schemes such as Time Division Multiplexing. Our findings

also question the validity of treating inter-user interference as Gaussian noise to measure system performance in practical systems, and thereby the validity of the IAGN model for these systems. To the best of our knowledge, this is the first such attempt.

B. Discussion

In our investigation, we made some assumptions to simplify the design process and performance characterization. We now discuss how our framework can be modified to investigate the value of SC even when these assumptions are relaxed.

- 1) The channel ordering exploited for two-user SC in this paper can be extended to three or more users, although finding the rate-region experimentally can be more cumbersome due to the presence of two power variables α_1 and α_2 .
- 2) The code library that was held fixed in our design can in general be a design parameter by itself. This becomes relevant when an SC system is to be built from scratch rather than on top of an existing single-user system as in this paper. Nonetheless, the current framework can be adopted to evaluate the performance of a given library.
- 3) The BS determines α from its knowledge of the individual link SNRs. In practice, these can be obtained via a feedback link. For every such SNR pair, the BS can choose an SC scheme by adopting the procedure described in this paper.
- 4) The “rate-centric” view in this paper focuses on the spectral efficiency gains from SC (for a given reliability), and has its roots in information theory. Viewing SC as a multiuser coding scheme, it is also possible to adopt a “reliability-centric” view wherein one measures the reliability gains (due to SC's coding gains) for a given pair of spectral efficiencies. Although both viewpoints are equivalent, they differ greatly in their experimental complexity. We study the latter problem in our recent submission [18] via on-air experiments.

C. Future Work

By experimentally demonstrating the benefits of SC using off-the-shelf single-user techniques, we have shown that SC is a potentially valuable transmission scheme. Our work can be extended in several directions. For example, the code library selection can leverage both advanced coding techniques such as turbo or LDPC codes and advanced receiver architectures based on iterative interference cancellation (see, e.g., [19], [20] and the references therein). Also, the principle of superposition can be applied to other cases where signal superposition is known to be theoretically optimal, e.g., in multiple-access channels (models for cellular uplinks) and for certain classes of relay channels. In these cases, the superposition process occurs at the *receiver* rather than at the transmitter in the BC; thus signals that interfere have different propagation paths. This difference opens up a new set of problems in system design, such as node synchronization, channel code selection and receiver design. Allowing multiple antennas at the TX and/or RX adds another dimension to the design space. Another possible line of investigation is to analyze the implications

¹⁸About 68% of values drawn from a normal distribution are within one standard deviation away from the mean.

of SC for higher layers in the network stack. For example, the problem of scheduling multiple users with an SC-enabled physical layer involves many interesting tradeoffs [21].

VII. ACKNOWLEDGMENT

The support of NSF (grants CNS-1016742 and 0830651) and DARPA/IPTO IT-MANET program (grant W911NF-07-1-0028) is gratefully acknowledged.

REFERENCES

- [1] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
- [2] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd edition. John Wiley & Sons, Inc., 2006.
- [3] J. Mitola III, "Software radios: survey, critical evaluation and future directions," *IEEE Aerospace Electron. Syst. Mag.*, vol. 8, no. 4, pp. 25–36, 1993.
- [4] *GNU Radio*. Available: <http://www.gnu.org/software/gnuradio/>.
- [5] X. Li, W. Hu, H. Yousefzadeh, and A. Qureshi, "A case study of a MIMO SDR implementation," in *Proc. 2008 IEEE Military Commun. Conf.*, pp. 1–7.
- [6] E. Bayraktaroglu, C. King, X. Liu, G. Noubir, R. Rajaraman, and B. Thapa, "On the performance of IEEE 802.11 under jamming," in *Proc. 2008 IEEE Conf. Comput. Commun.*, pp. 1265–1273.
- [7] K. Mandke, S.-H. Choi, G. Kim, R. Grant, R. C. Daniels, W. Kim, R. W. Heath, and S. M. Nettles, "Early results on Hydra: a flexible MAC/PHY multihop testbed," in *Proc. 2007 IEEE Veh. Technol. Conf. – Spring*, pp. 1896–1900.
- [8] R. Alimi, L. Li, R. Ramjee, H. Viswanathan, and Y. R. Yang, "iPack: in-network packet mixing for high throughput wireless mesh networks," in *Proc. 2008 IEEE Conf. Comput. Commun.*, pp. 66–70.
- [9] R. K. Ganti, Z. Gong, M. Haenggi, S. Srinivasa, D. Tisza, S. Vanka, and P. Vizi, "Implementation and experimental results of superposition coding on software radio," *2010 IEEE Conf. Commun.*
- [10] J. G. Proakis, *Digital Communications*, 4th edition. McGraw-Hill, 2000.
- [11] A. G. Fabregas, A. Martinez, and G. Caire, "Bit-interleaved coded modulation," *Foundations and Trends in Commun. and Inf. Theory*, vol. 5, no. 1–2, pp. 1–153, 2008.
- [12] G. Caire, G. Taricco, and E. Biglieri, "Bit-interleaved coded modulation," *IEEE Trans. Inf. Theory*, vol. 44, pp. 927–946, May 1998.
- [13] G. Caire and E. Viterbo, "Upper bound on the packet error probability of terminated trellis codes," *IEEE Commun. Lett.*, vol. 2, pp. 2–4. Jan. 1998.
- [14] J. Heiskala and J. Terry, *OFDM Wireless LANs: A Theoretical and Practical Guide*. Sams, 2001.
- [15] T. M. Schmidl and D. C. Cox, "Robust frequency and timing synchronization for OFDM," *IEEE Trans. Commun.*, vol. 45, pp. 1613–1621, 1997.
- [16] Z. Gong, S. Vanka, S. Srinivasa, P. Vizi, K. Stamatiou, and M. Haenggi, "Superposition coding strategies: design and experimental evaluation," Tech. Rep. EWA2011-1, Emerging Wireless Architectures Laboratory, University of Notre Dame, Sep. 2011. Available: http://www.nd.edu/~mhaenggi/pubs/techreport_sc11.pdf.
- [17] F. Baccelli, A. El Gamal, and D. Tse, "Interference networks with point-to-point codes," *IEEE Trans. Inf. Theory*, vol. 57, no. 5, pp. 2582–2596, May 2011.
- [18] S. Vanka, S. Srinivasa, and M. Haenggi, "A practical approach to strengthen vulnerable downlinks using superposition coding," *IEEE Conf. Commun.*, Ottawa, Canada, June 2012.
- [19] R. Zhang and L. Hanzo, "A unified treatment of superposition coding aided communications: theory and practice," *IEEE Commun. Surv. Tuts.*, no. 99, pp. 1–18, July 2010.
- [20] J. Andrews, "Interference cancellation for cellular systems: a contemporary overview," *IEEE Wireless Commun. Mag.*, pp. 19–29, Apr. 2005.
- [21] P. Vizi, S. Vanka, S. Srinivasa, M. Haenggi, and Z. Gong, "Scheduling using superposition coding: design and software radio implementation," *2011 IEEE Radio Wireless Week*.



Sundaram Vanka (S'08) received the B.Tech. and M.Tech. degrees in Electrical Engineering from the Indian Institute of Technology, Madras in 2003. From 2003-07 he was with Redpine Signals Inc., a fabless semiconductor and wireless systems company with offices in Hyderabad, India and Santa Clara, California, where he developed baseband algorithms and VLSI architectures for IEEE 802.11abgn and 802.16e-compliant chipsets. He is a co-inventor of 5 US patents and 2 patent applications for MIMO-OFDM receiver and channel

decoder design.

In Aug. 2012, he graduated with the Ph.D. degree in Electrical Engineering from the University of Notre Dame, Indiana, USA. He has been awarded a Graduate Fellowship for 2010-11 by the university's Department of Applied and Computational Mathematics and Statistics. His research interests include wireless communications and networking, with an emphasis on the use of stochastic geometric techniques in the analysis and design of wireless networks, and wireless system prototyping. He is currently a Senior Staff Engineer at Broadcom Corporation.



Sunil Srinivasa received the B.Tech degree in electrical engineering from the Indian Institute of Technology (IITM), Madras, in 2004, and the M.S. and Ph.D. degrees in electrical engineering from the University of Notre Dame in 2007 and 2011, respectively. He is currently a senior systems design engineer at LSI Corporation in Milpitas, California, where he works on the development of high-speed serial links. His research interests include wireless communications and networking, and information theory.



Zhenhua Gong (S'10) received the B.S. degree in electrical engineering from Shanghai Jiao Tong University, Shanghai, China, in 2007. Since 2007, he has been with the Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN, where he is working towards the Ph.D. degree. He is interested in the analysis and design of wireless sensor networks. Specifically, he focuses on the stochastic geometry for the performance analysis of mobile networks under interference constraints. Besides the theoretical research, he is also involved in the implementation of multi-user techniques in wireless networks.



Peter Vizi received his Master of Sciences in Computer and Electrical Engineering degree from Peter Pazmany Catholic University, Faculty of Information Technology, Budapest in 2007. Between 2007 and 2011 he was a graduate student at the Multidisciplinary Doctoral School of Peter Pazmany Catholic University. His research interest includes Quality of Service communication and multicast over packet switched networks, wireless sensor networks and algorithm desing for graphics processing units. He is currently an Associate Software Developer at

Morgan Stanley.



Kostas Stamatiou (S'98, M'09) received his Diploma in Electrical and Computer Engineering from the National Technical University of Athens in 2000 and his M.Sc. and Ph.D. degrees in Electrical Engineering in 2004 and 2009, respectively, from the University of California San Diego. In 2010, he spent a year as a post-doctoral scholar at the Department of Electrical Engineering of the University of Notre Dame. He currently holds a researcher appointment at the Department of Information Engineering at the University of Padova. His general

research interests lie in the areas of wireless ad hoc networks, underwater communications, queuing theory and energy harvesting sensors.



Martin Haenggi (S'95, M'99, SM'04) is a Professor of Electrical Engineering and a Concurrent Professor of Applied and Computational Mathematics and Statistics at the University of Notre Dame, Indiana, USA. He received the Dipl.-Ing. (M.Sc.) and Dr.sc.techn. (Ph.D.) degrees in electrical engineering from the Swiss Federal Institute of Technology in Zurich (ETH) in 1995 and 1999, respectively. After a postdoctoral year at the University of California in Berkeley, he joined the University of Notre Dame in 2001. In 2007-08, he spent a Sabbatical Year at the

University of California at San Diego (UCSD). For both his M.Sc. and his Ph.D. theses, he was awarded the ETH medal, and he received a CAREER award from the U.S. National Science Foundation in 2005 and the 2010 IEEE Communications Society Best Tutorial Paper award.

He served as a member of the Editorial Board of the *Elsevier Journal of Ad Hoc Networks* from 2005-08, as a Guest Editor for the *IEEE JOURNAL*

ON SELECTED AREAS IN COMMUNICATIONS in 2008-09, as an Associate Editor for the *IEEE TRANSACTIONS ON MOBILE COMPUTING (TMC)* from 2008-11 and for the *ACM Transactions on Sensor Networks* from 2009-11, and as a Distinguished Lecturer for the IEEE Circuits and Systems Society in 2005-06. He also served as a TPC Co-chair of the Communication Theory Symposium of the 2012 IEEE International Conference on Communications (ICC'12), and as a General Co-chair of the 2009 International Workshop on Spatial Stochastic Models for Wireless Networks and the 2012 DIMACS Workshop on Connectivity and Resilience of Large-Scale Networks. Presently he is a Steering Committee Member of TMC. He is a co-author of the monograph "Interference in Large Wireless Networks" (NOW Publishers, 2009) and the author of the textbook *Stochastic Geometry for Wireless Networks* (Cambridge University Press, 2012). His scientific interests include networking and wireless communications, with an emphasis on ad hoc, cognitive, cellular, sensor, and mesh networks.