GRaTIS: Free Bits In The Network

Dola Saha, Aveek Dutta, Dirk Grunwald and Douglas Sicker

University of Colorado

Boulder, CO 80309–0430 USA

Email: {Dola Saha, Aveek.Dutta, Dirk.Grunwald, Douglas.Sicker}@colorado.edu

Abstract—Recent work has examined techniques to estimate the "best" modulation rate for data networks such as 802.11a/g. While accurate rate estimation yields better rate-selection decisions and increased throughput, those methods must still choose between a handful of modulation rates. Each modulation rate is effective in a range of actual signal-to-noise ratios (SNRs) but the limited number of practical rates means that transmitters are often forced to "step down" to a lower data rate despite having a higher SNR than the minimum required for that lower rate.

In this paper we describe, evaluate and implement a practical multiuser communication scheme that exploits these discrete "steps" in modulation rates to transmit two packets in the time normally needed to transmit a single packet, increasing aggregate throughput precisely when it is most needed - when the network is busy and suffers from rate unfairness. Because the method transmits a group of packets simultaneously, we call this scheme Group Rate Transmission with Intertwined Symbols, or GRaTIS. In addition to up to 120% improvement in network throughput achieved by GRaTIS, the technique is backward compatible with 802.11 and doesn't require complex DSP algorithms as required by competing methods.

Index Terms—Wireless communication, Protocol architecture, Signal processing systems, Algorithm/protocol design and analysis, Reconfigurable hardware.

1 INTRODUCTION

Modern wired, wireless and optical communication systems use different *modulation schemes* to balance data rates against error rates. Each modulation scheme encodes a varying number of bits in a physical representation of the data. Low rate modulations are used in a noisy channel and high-rate modulations are used when the SNR is higher. For example, wireless systems such as 802.11 or WiMAX use the following modulations: BPSK (2 states, 1 bit), QPSK (4 states, 2 bits), 16QAM (16 states, 4 bits) and 64QAM (64 states, 8 bits). These modulates are augmented with different redundancy codes to achieve a fixed number of "transmission rates".

Those transmission rates are robust under varying SNRs. For example, the QPSK-3/4 rate requires an SNR of 8.0dB to deliver 18Mb/s throughput. The 16QAM-$1/2$ rate, delivering 24Mb/s, could be used if the SNR was 12.5dB, however if the SNR falls between 8dB and 12.5dB we get a "better" signal that reduces packet drops but results in little net throughput improvement. In this paper, we show that it is possible to exploit the higher SNR of this primary node to encode another message for a second receiver, increasing the aggregate network bandwidth. In doing so we introduce multiple data rates to provide an even gradation of SNR across a group of receivers – we call this Group Rate Transmission with Intertwined Symbols, or GRaTIS. Extending the prior example, consider the network organization in figure 1(a). If the link between Charlie and Beta has an SNR of 10dB, and the link between Charlie and Alpha has an SNR of 20.5dB, using GRaTIS we can send an 18Mb/s message to one node *and simultaneously send another 18Mb/s message to a second node resulting in an effective throughput of 36Mb/s.*

(a) Example Of Spatial Rate Diversity

(b) SNR Profile 1

(c) SNR Profile 2
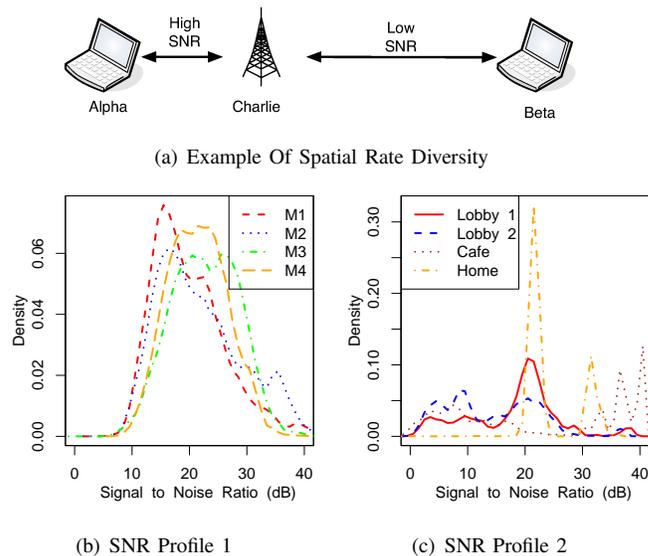
Fig. 1. Variation of SNR due to spatial diversity in 802.11a/g networks. Profile 1: Measured indoors by 4 packet sniffers at SIGCOMM 2008 [36] Profile 2: Measured indoors in common areas around a university cafe and lobbies and also in home networks.

Our method depends on *SNR diversity* between the receivers – in other words, we exploit the fact that most networks have nodes that experience different SNR values. Fortunately, most networks exhibit considerable SNR diversity. Figure 1(b) shows the distributions of SNR at a number of locations measured at a SIGCOMM conference in 2008 and figure 1(c) shows the same variation, which are measured by us around

a university campus. Figure 1(c) also shows a SNR profile that is typical to a home network shared by two users with high volume video streaming. The diversity in SNR occurs because of the spatial layout of nodes, room geometries and interference from other sources.

A broader set of related work is discussed in §8, but it is useful to summarize how GRaTIS relates to similar ideas. GRaTIS is similar to *hierarchical modulation*, which is used in digital broadcasting, but differs because it is applied to non-broadcast data and combines data targeted for specific nodes. CDMA networks use *multiuser detection* methods in the uplink and exploit the gains from orthogonal spreading codes and closed loop power control. CDMA systems use successive interference cancellation (SIC), which has been examined in a broader networking context [15]. SIC removes an interfering signal in order to reveal a secondary signal, and is part of the complex signal processing used in CDMA base-stations. By comparison, GRaTIS can address either infrastructure or *ad hoc* traffic (although our analysis emphasizes an infrastructure downlink), is simple to implement and depends on the wireless stations having *diverse* SNR levels, simplifying control. In GRaTIS, the original signal is encoded to be received correctly by the two receivers and doesn't suffer from channel estimation errors and the concerns that have been raised about the broader utility of SIC [38]. SIC is also used in *superposition codes*, which is another method for entwining two messages [27], [4], but superposition coding requires significantly enhanced SNR for both receivers of the message, limiting its utility and it is also difficult to implement. By comparison, GRaTIS requires only a modest increase in SNR that provides the required diversity to achieve network wide gain and is easy to implement on conventional signal processing pipelines.

The benefits of GRaTIS are the larger number of *group rates*, which provide increased opportunities for improving performance but doesn't impact performance when not used; it can be made backward compatible with existing wireless networks; it is easily implemented on conventional wireless signal processing pipelines; and, it complements the gains with advanced rate adaptation techniques [16], [43], [37] and physical layer techniques such as FARA [6] and network coding [19]. Using different analysis techniques we show that this scheme is both practical, profitable and implementable. We summarize the key contributions of this paper as follows:

- We reinterpret the constellations already available for conventional wireless links and provide *group rates*, which result in higher network throughput with no hardware changes.
- We perform a standard analysis of packet error rates for this scheme to ascertain the applicability in real networks.
- We implement GRaTIS on a 802.11a/g compatible software defined radio (SDR) prototype to show that the technique is easy to implement and makes use of existing hardware modulation and de-modulation methods. In the prototype, much of the added processing is handled by simple software, rather than complicated fixed-function hardware and DSP algorithms.
- We use the SDR nodes in a testbed setup to measure the SNR requirement for over-the-air transmission of GRaTIS encoded multiuser data packets.
- We apply the results of the testbed to analyze various 802.11a/g traces from SIGCOMM conferences and other small/mid sized wireless data networks to determine how the diverse range of SNR of the receivers can be used to obtain substantial gain in network throughput. This shows the potential performance improvement when GRaTIS is applied on realistic downlink traffic.
- We also show that GRaTIS outperforms a competing method (superposition coding) both in theory and practice. We validate this through experiments conducted in a testbed of SDR nodes using actual over-the-air packet transmissions.

In §2, we describe the GRaTIS technique in more detail. In §3 we conduct a basic packet-error analysis for the GRaTIS protocol showing under what conditions the technique can be used. We then describe the hardware used to implement the technique in §4.1 and the results from our implementation in §4.2. §5 describes a (potential) gain analysis for GRaTIS using 802.11a/g packet traces. The benefits of sending two messages at once can yield very different results as described in §7 and in §8, we discuss prior work in this domain and place our work among contemporary techniques.

2 GRaTIS: FREE BITS

In this section, we describe GRaTIS. Discrete steps in the SNR requirements for each data rate forces the rate adaptation algorithm to fall back to a lower rate even if the node is reachable at a SNR higher than the minimum required. In this scenario, we use GRaTIS to identify two "layers" of constellations within the standard constellations available in 802.11a/g. These two layers are used to map two packets of two different users to form one single packet, such that the time required to transmit the merged packet in GRaTIS is less than the time required to transmit two separate packets in the best achievable data rate of 802.11a/g. The two layers are designed in such a way that one of the layers can be decoded by a legacy decoder, and we call it the *Base Layer*. The second layer is obtained by extracting a few bits after the legacy demodulation system converts the I/Q samples from the analog domain to the binary domain, and we call it *GRaTIS Layer*. The second packet is transmitted during the transmission of the first packet, without any extra airtime, and comes as free bits to the receiver with a higher SNR – those free bits increase the aggregate throughput of the network.

For example, assume the received SNR of two nodes n_1 and n_2 from a common transmitter are SNR_{n_1} and SNR_{n_2} respectively. Also, there exists a GRaTIS rate, where the SNR requirement for Base and GRaTIS layers are SNR_b and SNR_g respectively, and $SNR_b <= SNR_{n_1}$ and $SNR_g <= SNR_{n_2}$. Consider the best achievable data rate in 802.11a/g are R_{n_1} and R_{n_2}, while that using GRaTIS are R_b and R_g respectively. The common transmitter uses GRaTIS to transmit x bits of data at rate R_b to node n_1, which takes time t_g (total transmission time using GRaTIS). Since in GRaTIS, the GRaTIS layer is transmitted at the same time along with
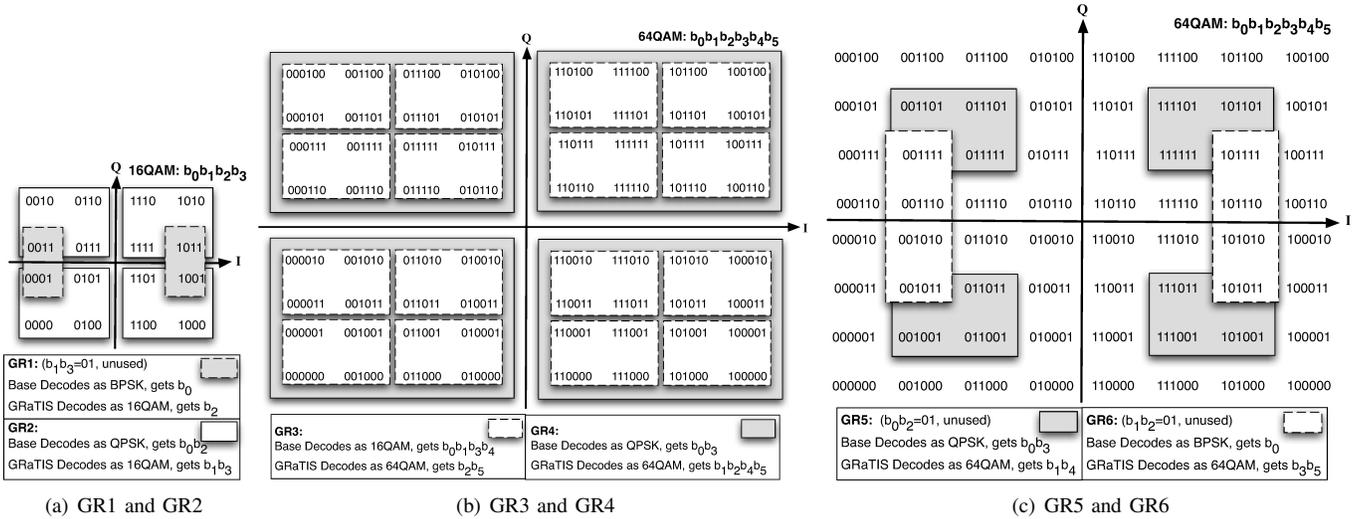
Fig. 2. Encoding and decoding of GRaTIS derived from standard 802.11a/g constellations. Rectangular regions show the transmitted cluster and the corresponding decision boundary for the base layer.

the Base layer, there is no extra time required to transmit the GRaTIS layer. So, the transmitter also transmits y bits of data using rate R_g within the same time t_g. Therefore the total data rate of this transmission is $\frac{(x+y)}{t_g}$.

Now we calculate the achievable data rate if the common transmitter uses 802.11a/g to transmit the same packets. The transmitter uses rate R_{n_1} to transmit the x bits of data to node n_1 in time t_1. After this, it transmits y bits of data at rate R_{n_2} in time t_2. The aggregate data rate for these two transmissions is $\frac{(x+y)}{t_1+t_2}$. The pair of rates R_b and R_g are selected as one of the GRaTIS rates, *iff* $\frac{(x+y)}{t_1+t_2} < \frac{(x+y)}{t_g}$. Or in other words, GRaTIS rates are selected only if there is potential gain in aggregate throughput over the legacy system.

We have developed an encoding and decoding technique that requires minimum change to a stand alone 802.11a/g transceiver and relies on identifying clusters that are a subset of the standard set of 802.11a/g constellations (BPSK, QPSK, 16QAM and 64QAM). Depending on the cluster size and how they are split between the two layers, the error performance of the base layer and the GRaTIS layer varies. Typically, the clusters are derived from higher order constellations (16QAM and 64QAM) so that there is a sufficient amount of bits available to encode the packet for the GRaTIS layer.

2.1 Encoding Packets using GRaTIS

We introduce *six* distinct GRaTIS rates, or methods of combining packets, as shown in figure 2, to increase aggregate throughput of the network. The GRaTIS rates are termed *GR1* through *GR6*. We select the combination of layers as one of the GRaTIS rates if the data rate achievable by merging is more than that of two packets transmitted separately. At the transmitter, two packets are encoded independently up to the modulation subsystem as shown in figure 3(a). Then the bits of two packets, b_b and b_g, are encoded at rates R_b for the base and R_g for the GRaTIS layer respectively, and combined to form a compound symbol that represents one of the constellation points corresponding to a standard modulation in 802.11a/g,

(a) Transmitter Pipeline

(b) Receiver Pipeline
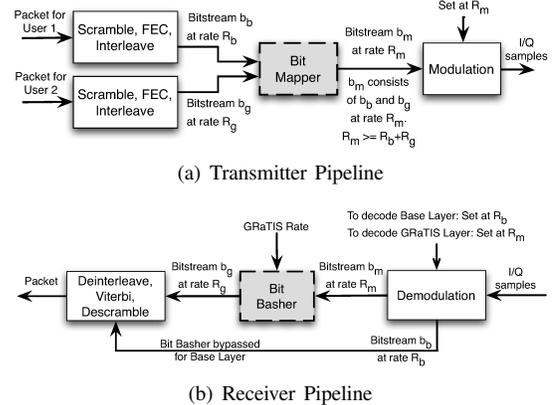
Fig. 3. Transceiver pipeline for GRaTIS – shaded subsystems show additional processes required for GRaTIS.

denoted by R_m. This mapping ensures that the compound symbols are mapped only to the I/Q vectors that are part of a selected GRaTIS cluster. In this way, the modulator remains unchanged, as it is fed with the compound bitstream (b_m), and modulation type (R_m) to which it modulates. Since all the packet merging is done at the bit-level it does not require any change in the signal processing pipeline.

The clusters are selected to optimize properties of the I/Q-plane mapping used to represent information in wireless networks. For example, Figure 2(a) shows the constellation points used in *GR1* and *GR2*, which are derived from a 16QAM constellation. In *GR1*, resultant cluster points are modulated to carry two bits of useful information, one bit for each layer. To reduce the probability of error in the base layer, the points are chosen such that the vectors in the I-plane are greater than that of BPSK mapping, while the deviations in the vectors of the Q-plane is used to carry another BPSK packet in the GRaTIS layer. Out of the 4 bits available for every constellation point in the cluster, bit b_0 is used to encode the base layer, and bit b_2 contains the GRaTIS layer. The other two bits, b_1 and b_3, remain constant at 0 and 1, respectively to map

the compound symbol to the desired cluster. The shaded region shows the transmitted constellation for *GR1*. *GR2* uses all the constellation points of a 16QAM constellation, and provides 2 bits of information per subcarrier, as is done in QPSK, to each of the two nodes. Bits b_0b_2 and bits b_1b_3 are used to encode the information of base layer and GRaTIS, respectively. As a result of such mapping when a cluster point that is closest to an axis crosses the axes due to channel noise, a symbol error occurs for the base layer, as seen in QPSK modulations, but this event does not incur any error in the GRaTIS layer. Hence, this type of mapping provides some extra error protection to the GRaTIS layer and so, the SNR requirement for GRaTIS layer of *GR2* is less than that of 16QAM.

GR3 and *GR4* utilizes 64QAM constellation to encode the two layers as shown in figure 2(b). *GR3* provides a 16QAM data rate to the base layer using bits $b_0b_1b_3b_4$, while *GR4* provides a QPSK data rate to the base layer using bits b_0b_3. The remaining bits are used to encode the data of GRaTIS layer. As in *GR2*, the GRaTIS in these cases have additional protection from error as crossing the base layer boundaries does not introduce any error in the GRaTIS layer. *GR5* and *GR6* uses a cluster derived from a 64QAM constellation as shown in figure 2(c). *GR5* uses bits b_0b_3 for the base layer and bits b_1b_4 for the GRaTIS, providing QPSK data rate to both the packets. Bits b_2b_5 are modulated as 1, to generate the desired cluster as shown in dash-dotted lines. *GR6* uses bits b_0 and b_3b_5 to encode the information of the base and GRaTIS layers, respectively, while modulating bits b_1 as 0 and b_2 as 1 leading to the desired cluster points.

2.2 Decoding Packets using GRaTIS

GRaTIS has two layers, intended for two receivers. The Base layer is encoded in such a way that its decoding is the same as decoding any generic 802.11a/g packet. Decoding the constellations to information bits is done using pre-defined thresholds called *decision boundaries*. For BPSK, this boundary is the Q-axis, whereas for QPSK there are four such boundaries: the four quadrants of the I/Q plane. The number of decision boundaries increases with the increasing number of points in the constellation. For example, in *GR1*, bit b_0 of 16QAM modulation is encoded as the base layer. So, the constellation points generated in the left side of the Q-plane always yields a value of 0, and the right side of the Q-plane is always decoded as a 1. This phenomenon is the same as BPSK modulation. At the receiver side, the base layer of *GR1* can be simply decoded as a BPSK packet, i.e., any received sample to the right side of the Q-plane is mapped to bit 1, and a point to the left side will yield a bit 0.

Figure 3(b) shows the demodulation pipeline for a GRaTIS compatible node. The GRaTIS layer is first decoded as the constellation from which the GRaTIS cluster has been derived (R_m). Then, the bits designated for the GRaTIS layer are extracted using the *GRaTIS rate* information, to form the bit stream for the second packet (b_g) at rate R_g. The rest of the receiver decode pipe remains unchanged. For example, if a packet has been transmitted using *GR2*, then it is decoded as 16QAM packet, to yield 4 bits. To extract the GRaTIS bits, the
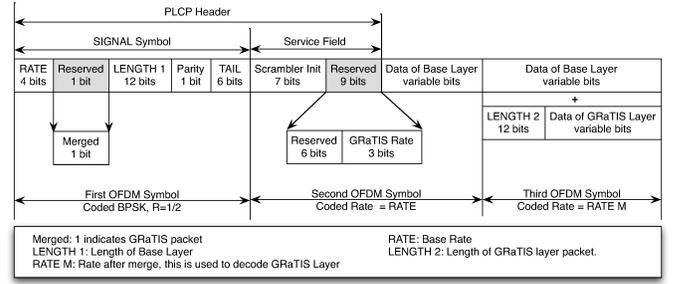
Fig. 4. Modified 802.11a/g PLCP header – shaded fields indicate modifications to support GRaTIS.

receiver extracts 2 bits, in this case b_1b_3. In this way, with the knowledge of the proper constellation mapping, we can decode the extra packet while being completely backward compatible with a legacy node, which is unaware of any GRaTIS layer transmission.

2.3 Medium Access Control for GRaTIS

To successfully decode a GRaTIS packet, a node needs to know the encoding information of the packet. We use the reserved bits of the PLCP Header of IEEE 802.11a/g packet to provide the encoding information as shown in figure 4. The 1-bit *reserved* field in the *SIGNAL* symbol is used to denote whether the packet contains a GRaTIS layer. The RATE field indicates the modulation rate at which the base layer is encoded. We use 3 bits out of 9 reserved bits in the *Service* field to indicate the rate of the GRaTIS layer encoding, which can have values from 1 to 6. As in any unmerged packet of IEEE 802.11a/g, the first symbol is modulated in BPSK with $1/2$ rate coding. The second symbol is modulated in the rate specified in the *RATE* field of SIGNAL symbol. In this symbol, GRaTIS layer information is embedded, which is used to demodulate from the third OFDM symbol onwards, which marks the beginning of the data payload. The first 12 bits of the third symbol in the GRaTIS layer carries the length of the GRaTIS packet that is used to decode the second packet.

The encoding procedure ensures that decoding of the base layer is exactly the same as decoding any generic 802.11a/g packet. The demodulation for the GRaTIS layer changes from the third symbol onwards based on the information received in the 'GRaTIS rate' field in the second symbol of the packet, as shown in figure 4. Based on the GRaTIS layer decoding capability of the clients, an AP can decide which group rate to use to merge packets. An AP will never merge packets of a GRaTIS incompatible client in the GRaTIS layer.

In the 802.11a/g PHY layer, each message must be individually acknowledged. One major hindrance in using multiuser communications such as GRaTIS is the need for those acknowledgments. For this we rely on a *simultaneous acknowledge mechanism* (SMACK) [8] to gather acknowledgments from multiple recipients of the merged packet. SMACK will reduce the overhead of scheduling multiple acknowledgment packets and will reduce the multi-party acknowledgment time. The AP can schedule acknowledgments for the clients that cannot transmit SMACK.

2.4 GRaTIS as a Facilitator

Recent research on accurate channel prediction [43], [37], [16] has enabled finer, more accurate control for the correct data rate for a link. GRaTIS will benefit from these channel estimation techniques, which will essentially help to make correct decision on GRaTIS rate selection, such that there is minimum packet loss and packet re-transmissions are minimized in the network.

In recent years, there has been encouraging work in the physical layer of 802.11a/g, yielding higher aggregate network throughput. GRaTIS is orthogonal to many of these technologies, which can be applied in conjunction with GRaTIS to improve the overall performance of the network. Wireless network coding [20], [22] can be used with GRaTIS in 802.11a/g wireless network to decrease the number of packet transmissions in a network. A frequency aware rate selection [31] mechanism can be applied for GRaTIS rates to transmit lower data rates in frequency selective fading scenarios.

Although we have implemented GRaTIS in a 802.11a/g network, it can be implemented in any wireless or wired network that uses modulation to encode packets. We have implemented and experimented with the protocol in OFDM-based system, and it can be readily applied to any other single carrier or OFDM based wired or wireless protocol. In this paper, we have shown how the merging information can be disseminated in a WiFi-based network. However, to extrapolate this technique in any other protocols, like WiMax or LTE, new methods to disseminate this information is required, and is out of scope of this paper. In the cognitive radio domain, non-contiguous OFDM transmissions will be a necessity, where GRaTIS can be efficiently utilized in the subcarriers chosen for cognitive transmission.

3 GRaTIS: Rate Analysis

In this section we evaluate the bit error rate (BER) and packet error rate (PER) for various group rates in GRaTIS, in the presence of Additive White Gaussian Noise (AWGN) using the standard techniques used to analyze digital communication performance. At the receiver, the noisy constellation points in the I/Q plane are mapped to corresponding bits by using maximum likelihood (ML) decoding. Constellation points are required to be within an area in the I/Q-plane defined by a modulation dependent decision boundary to ensure error free decoding. The BER for such a scheme is given by eq. 1.

$$P_B(E) = \frac{1}{2} \operatorname{erfc}\left(\sqrt{\frac{\Delta_E(i,j)}{4N_0}}\right) \qquad (1)$$

The bit error rate for an arbitrary modulation scheme and ML decoding boundaries is upper bounded by eq. 2, where,

$$P_{WUB}(E) = \sum_{j=1}^{M-1} \sum_{i \neq j} \frac{1}{2M} \operatorname{erfc}\left(\sqrt{\frac{\Delta_E(i,j)}{4N_0}}\right) \qquad (2)$$

$$= \sum_{k=1}^{N} \frac{A_d(k)}{2M} \operatorname{erfc}\left(\sqrt{\frac{\Delta_E(k)}{4N_0}}\right) \qquad (3)$$

TABLE 1
Throughput and SNR requirements for 802.11a/g and GRaTIS rates

| Mod | CR[1] | Data Rate (Mbps) | | SNR(dB) | | | |
| | | | | Base | | GRaTIS | |
		Link	Grp	Th[2]	Exp	Th[2]	Exp
BPSK	1/2	6		3.0	4.5		
BPSK	3/4	9		5.0	6.0		
QPSK	1/2	12		6.0	7.0		
QPSK	3/4	18	n/a[3]	8.0	9.0	n/a[3]	n/a[3]
16QAM	1/2	24		12.5	13.0		
16QAM	3/4	36		17.0	18.0		
64QAM	2/3	48		19.5	24.0		
64QAM	3/4	54		21.0	26.0		
GR1	1/2	6	12	3.5	5.0	17.0	17.0
GR6	1/2	6	18	4.5	5.0	18.5	21.0
GR5	1/2	12	24	7.5	7.5	18.5	23.0
GR2	1/2	12	24	9.0	10.5	11.5	15.0
GR4	1/2	12	36	15.0	15.0	19.0	23.0
GR3	1/2	24	36	17.0	16.5	20.0	25.0
GR1	3/4	9	18	5.5		19.5	
GR6	3/4	9	27	7.0		20.5	
GR5	3/4	18	36	10.0	n/i[4]	20.5	n/i[4]
GR2	3/4	18	36	13.0		14.5	
GR4	3/4	18	54	19.5		21.5	
GR3	3/4	36	54	20.0		22.0	

[1]Coding Rate, [2]Theoretical, [3]Not Applicable, [4]Not Implemented

where,
- $\Delta_E(i,j)$ = Squared Euclidean distance between two distinct constellation points i and j.
- N = Possible different squared Euclidean distances in the *decoded* constellation, where $N \leq M(M-1)/2$.
- M = Total number of *decoded* constellation points.
- N_0 = Additive white noise power.
- $\Delta_E(k)$ = Distinct pairwise Euclidean distance in the *decoded* constellation.
- $A_d(k)$ = Number of signal pairs having squared Euclidean distance of $\Delta_E(k)$.

Using eq. 2 we can compute the BER for any constellation and ML decision boundary. An example BER computation for the base layer of *GR2* has been shown in the appendix. The PER for a packet size of 128 bytes after Viterbi decoding is computed using the maximum free distance for a particular coding rate and its corresponding distance spectrum [30], [11], [45], [14]. The theoretical PER plots for the Base layer and GRaTIS layer is shown in figure 5 and figure 6 respectively.

In the BER computation we consider that the GRaTIS constellations are mapped using Gray code [17] and encoded using a 1/2 rate (except for 64QAM, which is encoded using a 2/3 rate) as well as 3/4 rate convolution code. This analysis has been done to ascertain the operating range of different group rates and their potential benefits when used to merge packets using GRaTIS. The BER performance for the GRaTIS layer is as important as the base layer because it allows the MAC to identify users with suitable SNR that can decode the bits from the enhanced GRaTIS layer. The BER computation for the GRaTIS layer is similar to that of the base layer and can be easily calculated using eq. 2. We discuss the performance of both the layers using a testbed in §4.2. The theoretical and testbed results are listed in Table 1, which shows the SNR

(a) PER for coding rate = $1/2$ (Except $64QAM = 2/3$)

(b) PER for coding rate = $3/4$
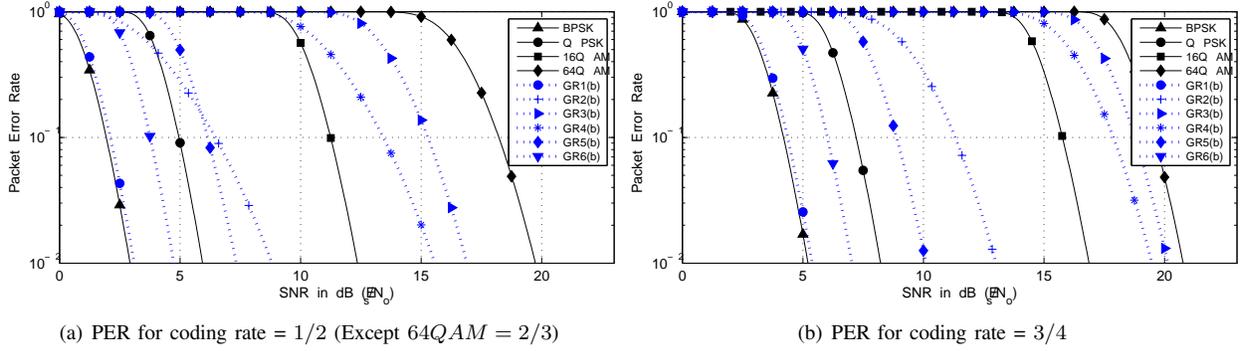
Fig. 5. PER for Base layer compared to legacy 802.11a/g modulations.

(a) PER for coding rate = $1/2$ (Except $64QAM = 2/3$)

(b) PER for coding rate = $3/4$
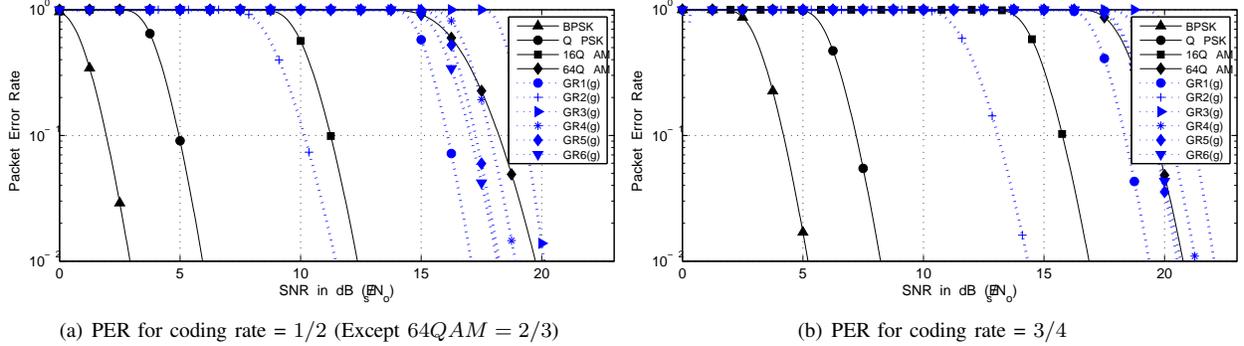
Fig. 6. PER for GRaTIS layer compared to legacy 802.11a/g modulations.

requirements for a 2% packet error rate (PER) along with the SNR requirement for the 802.11a/g standard rates. The SNR - throughput relationship for various group rates are used as a look-up while downlink packets are being considered to be merged. As shown in Table 1, the group rates provide a variety of step-down rates while utilizing SNR diversity in the network to increase the aggregate throughput of the network.

From figure 5, we find that the base layer of *GR1* and *GR6* requires a SNR between QPSK and BPSK. Hence these group rates can be used to modulate signals for nodes whose SNR are lower than that required by QPSK. Similarly, *GR6* and *GR2* offer two step down rates for nodes not reachable with 16QAM but having higher SNR than QPSK. While *GR3* offers similar flexibility by stepping down to an intermediate data rate instead of 16QAM, the benefits from *GR4* can be seen when used in conjunction with the GRaTIS layer: providing a combined data rate equal to that of 64QAM, which none of the two nodes would have been able to achieve with independent packet transmissions.

4 IMPLEMENTATION AND EVALUATION

In this section, we discuss the details of implementing GRaTIS in a reconfigurable radio and the method of evaluation with over-the-air packets, transmitted from the radio in a small testbed scenario in an indoor environment.

4.1 Implementing GRaTIS

The very nature of the technique used in GRaTIS requires minimal changes in a traditional 802.11a/g OFDM based radio pipeline. As discussed in §2, the encoding and decoding of the
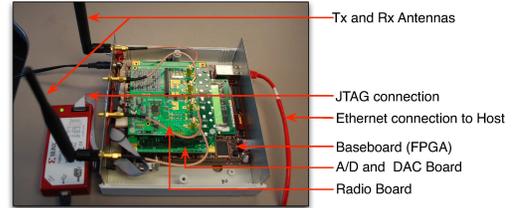
Fig. 7. SDR platform used to implement GRaTIS.

GRaTIS constellations can be easily performed in software by mapping the data bits to the target constellation points. If the proper compound symbols are provided to the baseband modulator, it would produce the target GRaTIS constellation without any requirement to change the I/Q vector lengths to achieve a certain error performance. Keeping the baseband modulator unaltered makes GRaTIS backward compatible and easy to implement in commodity hardware.

Although most of the encoding can be done in software, access to the packetization engine of the MAC layer is required. Since, this abstraction layer is not available to us from commodity hardware, we implemented this technique using a prototype hardware based on previous work [10], [7], [9], as shown in figure 7. The prototype uses a hybrid software defined radio (SDR) based on a Virtex-5 FPGA that can transmit and receive generic 802.11a/g data packets [17]. Using such a platform we can have access to the packetization layer of the MAC where the data bits from two users are combined to form a symbol in the I/Q plane, which is accomplished by the *bit mapper* unit shown in 3(a). Thus, we can selectively transmit only the I/Q samples that correspond to a GRaTIS cluster while others remain unused.

(a) GR1 (b) GR5 (c) GR6
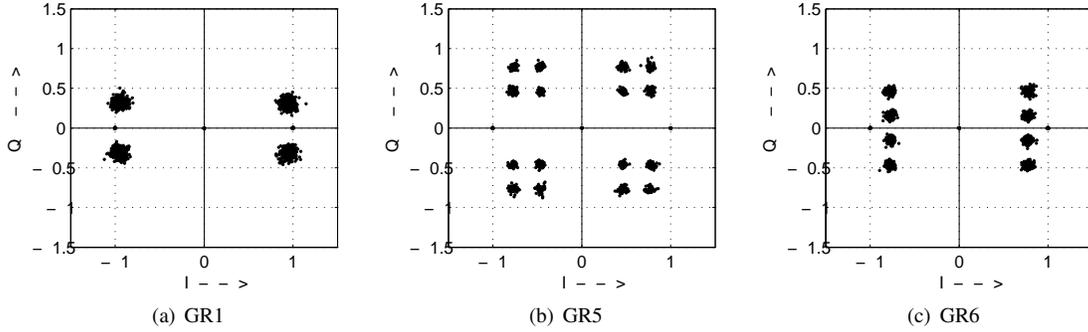
Fig. 8. Various GRaTIS constellations transmitted using the SDR prototype.

Similarly, the receiver of the base layer can be oblivious of any enhanced layer. Only receivers capable of decoding the GRaTIS layer require additional controls to decode packets. Depending on the GRaTIS rate used to encode the second packet, the demodulator extracts the additional bits: typically termed as *bit slicing* or *bit bashing*. Figure 3(b) shows the basic structure of the modified demodulator. The *bit basher* unit is responsible for slicing the bits of the GRaTIS layer. Many of these controls govern how the required bits are sliced from the combined constellation symbols, which is again a bit level operation that can be easily done by simple software controls to the underlying hardware as a part of the MAC de-packetization layer. Other receiver subsystems prior to the demodulator, *e.g.*, the synchronizer and the equalizer, remain unchanged for implementing GRaTIS. The equalizer always aims to restore the original transmitted constellation, while it is the decoder that decides either to interpret it as a standard packet or a combined packet using GRaTIS. Apart from decoding packets, the receiver also reports the average receive SNR of a packet and also performs MAC CRC checks in the hardware to measure packet loss.

Figure 8 shows the various constellations produced by the prototype SDR that support multiuser communication using GRaTIS. Figure 8(a) shows the constellation for *GR1*, which is derived from a 16QAM constellation by not using the other constellation points. Similarly, figure 8(b) and 8(c) shows a modified 64QAM constellation that provides a combined network throughput of 4 bits/OFDM subcarrier (2 for base and 2 bits for GRaTIS layer) and 3 bit/OFDM subcarrier (1 for base and 2 bits for GRaTIS layer) respectively.

As mentioned, the receiver uses the standard OFDM demodulation techniques and largely the same MAC layer. Although the hardware chain is unchanged, the software controlling the binary operations that follow the demodulator do need to be updated; this is usually a software upgrade. This is a distinct advantage over other multiuser decoding techniques that rely heavily on complex signal processing algorithms, involving custom constellations and more general control of the I-Q mapping. Hence, GRaTIS provides an example of harnessing the power of existing resources while innovating new and improved protocols.

4.2 GRaTIS: Putting it to Work

In this section, we utilize our hardware prototype transceiver to measure the actual SNR required for satisfactory operation in a

typical indoor wireless network. A testbed has been setup with three nodes, one transmitter and two receivers in an indoor laboratory environment. The nodes are placed at a distance of approximately 5m on work desks in enclosed cubicles. Channel quality and SNR variation is obtained by controlling the transmit power as well the location of the receiver nodes relative to the transmitter. The average throughput and average SNR required for 2% PER across different node arrangements have been computed. The group rates proposed in §2, with $1/2$ rate convolution coding for both the layers have been compared to the standard rates available in $802.11a/g$.

The SNR is computed from the digital I/Q samples in the hardware. Power is measured as $|r(n)|^2$, where $r(n)$ is the complex signal samples. P_{s+n} denote the power of the signal and the noise combined, averaged over 5 OFDM symbol period, after a packet is detected and the MAC CRC is received correctly. So, it reflects the average power over the data symbols of the OFDM packet. The noise power, P_n, is time averaged over 5 OFDM symbols after the packet is completely decoded. Both of these values are computed in the hardware and sent to the user along every data packet. The SNR in dB is then computed from these values using $10 \log_{10}((P_{s+n} - P_n)/P_n)$.

Figure 9 shows the performance of the base and GRaTIS layer of six group rates, along with standard modulations, BPSK, QPSK, 16QAM and 64QAM. For each modulation, we plot the physical layer throughput and mark the minimum SNR required for a PER of 2%. Maintaining acceptable error rates while maximizing throughput is important, as it might lead to unwanted re-transmissions consuming additional airtime and reducing the throughput of the network. These marks are the SNR requirements for each modulation below which it cannot be used reliably. The base and GRaTIS layers are denoted by suffixes '(b)' and '(g)' respectively.

GR1(b) and *GR6(b)* are two group rates that provide BPSK rate, and have an SNR requirement between BPSK and QPSK. Thus, when a node becomes unreachable in QPSK, we can use these two group rates to combine packets, and send a packet to another node reachable at higher SNR: at BPSK rate to node reachable at 17dB or higher using *GR1(g)* and at QPSK rate to a node reachable at 21dB or higher if using *GR6(g)*. Although the GRaTIS layer needs a higher SNR to operate, which could potentially receive a higher rate packet, the collective throughput by combining the two packets is more than throughput obtained by transmission of two individual
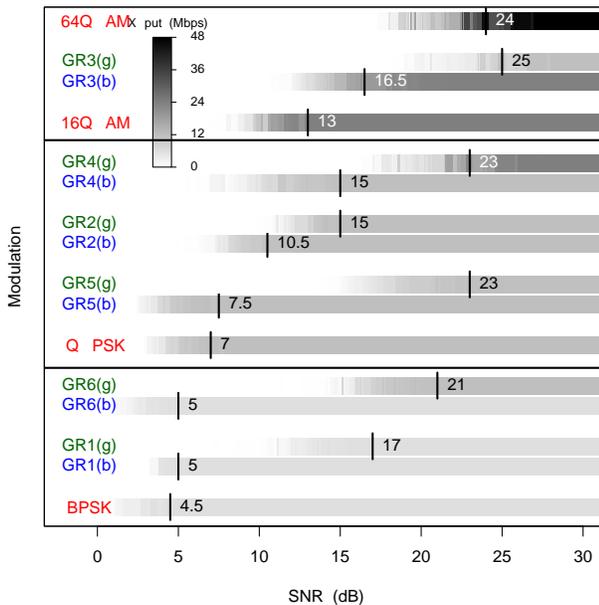
Fig. 9. Link throughput of GRaTIS((b)-Base, (g)-GRaTIS) and 802.11a/g rates with increasing SNR. The rates are grouped according to the increasing base rate. The numbers against each rate denote the SNR required to decode a 128 byte, $1/2$ rate convolution coded packet with 2% PER.

packets. The gain in throughput can be computed as described in §2.

GR2 and *GR5* provides a QPSK data rate to both of the layers, and have SNR requirements in between QPSK and 16QAM. The presence of these two group rates provides more flexibility to choose the GRaTIS rate in a wider range of SNR. If SNR of a node is more than 7.5dB but less than 10.5dB, and that of another node is more than 23.5dB, *GR5* can be used effectively, but *GR2* cannot be used in this SNR range. Similarly, if a node has SNR between 10.5dB and 13dB, while another near node has SNR between $15dB$ and 23.5dB, we cannot use *GR5* for the near node, but can successfully encode the packets using *GR2*.

GR3(b) provides a data rate equal to that of 16QAM, and has an SNR requirement in between 16QAM and 64QAM. Evidently, this group rate can be used whenever the SNR of a node falls below the SNR requirement of 64QAM providing the best effort data rate, 16QAM in this case, for that node. *GR4(b)* provides throughput equal to that of QPSK, but requires more SNR than 16QAM to be decoded correctly. It might seem that this group rate does not provide the best effort rate to the base layer. Careful observations reveals that *GR4(g)* provides a throughput equal to that of 16QAM, and has a lower SNR requirement than 64QAM. So, we can transmit a packet in *GR4(g)* to a node when its SNR falls below 64QAM, and *GR4(b)* then can be used to transmit any extra bits as a GRaTIS layer.

The SNR required to maintain a 2% PER for all the GRaTIS and 802.11a/g are listed in table 1. These experimental results are used to verify that the technique yields practical gains over individual packet transmissions and benefits a network with wide diversity of SNR.

5 GRaTIS: PRACTICAL GAINS

Theoretical and experimental results in §3 and §4.2 respectively show the potential benefits of using GRaTIS in modern networks like 802.11a/g and WiMax. It is sometimes difficult to understand the benefits of a particular wireless optimization from PER plots and bench experiments. Also, testbed implementations, where the network is flooded with UDP packets of the same packet length do not represent a realistic scenario of wireless networks. We want to know -

- How often do stations have sufficient SNR diversity to exploit GRaTIS?
- Does GRaTIS gain in variable packet lengths of the users?
- How useful is the combined coding efficiency over 802.11a/g network?
- Does GRaTIS work in different scenarios, like a conference hall, university cafeteria or home network?

To determine the *system* benefits of GRaTIS, we analyzed captured packet traces from SIGCOMM 2008 dataset [36]. The SIGCOMM traces reported signal and noise power in the Prism header [1], which have been converted to report SNR in dB. We have also captured packet traces in the common areas of our university with more than 50 active users, and in a home network of two users. The captured packet traces are referred to as 'Lobby 1', 'Lobby 2', 'Cafe' and 'Home'. 'Lobby 1' and 'Lobby 2' traces have been captured at the exact same location in a common area to identify the temporal variation of traffic pattern. 'Cafe' denotes another trace, which is also a common area for dining services. 'Lobby 1', 'Lobby 2' and 'Cafe' had more than 50 active users at any given time, while people walked through the area. Our captures use the Radiotap header [2], which also reports the signal and noise power, from which SNR has been computed.

We are interested in knowing the SNR variation in all the scenarios, to access whether GRaTIS can be used to improve the network performance. Benefits of this protocol are maximized when there are clients that are reachable at a wide variety of SNR facilitating packet combination at different GRaTIS rates. Figure 1(b) and 1(c) shows the SNR density variation of uplink data packets in the SIGCOMM and captured packet traces. Within the SIGCOMM traces, we have chosen 4 monitors that have the most data packets to generate the histogram of SNR. The histograms from all the monitors show similar distribution but the SNR is found to be evenly distributed within the set of clients. This presents good opportunity to merge packets that can be sent to clients reachable at different SNR levels. This similarity can be attributed to the confined nature of a conference room, where users were evenly distributed in the room. However, we notice 'Lobby 1' and 'Lobby 2' show significant variation in the histogram, with the maximum reaching around 20dB SNR. The trace for 'Cafe' shows two prominent spikes in the distribution, which are due to the spatial diversity of two very active users in the network. In the 'Home' trace, this diversity is more prominent as there were only two users in the network. All the scenarios show different diversity of the users, but undoubtedly enough variation in SNR to use GRaTIS.
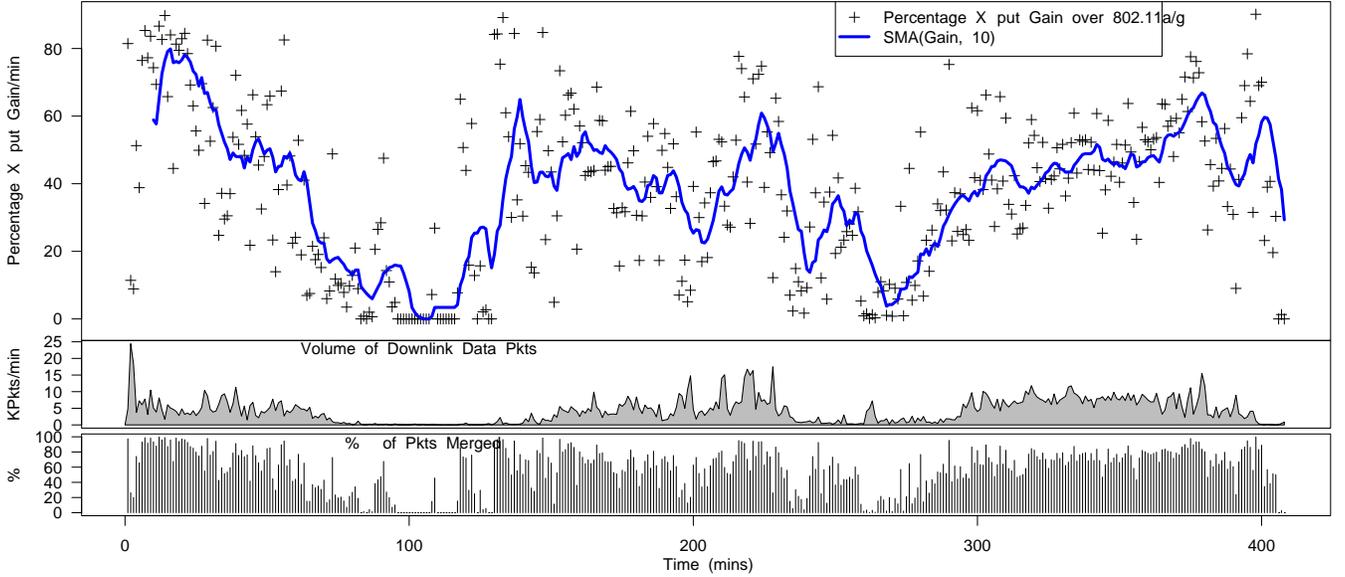
Fig. 10. An example implementation of GRaTIS when applied on the downlink data packets in an infrastructure 802.11a/g network with > 50 users. The first plot shows the percentage gain in aggregate throughput per minute with GRaTIS. A 10 minute Simple Moving Average (SMA) of the gain in also shown. The second plot shows the volume of packets in Kpackets/min while the third plot shows the percentage of packets transmitted per minute using GRaTIS.

We intend to use GRaTIS to merge downlink data packets, which constitutes most of the packet transmission over the air. The AP would be able to merge two packets intended for two clients, operating in different SNR regime, into one GRaTIS packet, if it has information of the SNR of the packets received from AP by each client. The traces did not provide the SNRs at which the clients are reachable from the AP (this requires packet monitoring at each client). However, the traces report a) the *SNRs* at which the monitor is reachable from each client, which gets updated whenever a client transmits a packet and is overheard by the monitor, and b) the actual downlink data packets transmitted by the AP, with the information of *packet size*. Now, if we place an AP at the monitor's location, it would use this information about the SNR to merge packets and transmit the same data packets as the AP is currently transmitting to its clients.

We analyze the traces based on the monitor's view of the network, merging downlink data packets based on the SNRs received from the clients. Two downlink data packets are merged only if the time required to transmit a merged GRaTIS packet is less than the time required to transmit two individual packets using 802.11a/g. The queue length is finite, and only 10 packets have been considered to be available at any time for merging. Off-the-shelf wireless AP queue capacity varies from 39 to 337 packets [26]. In congested network scenarios with more packets available in the queue, there will be more options available, which will improve the performance of GRaTIS. To maintain low latency and fairness in the network, the merging algorithm always transmits the packet in the lowest position of the queue, and searches for any possible combination of GRaTIS among the next 9 packets that will reduce the overall time required for transmission. Often in our evaluation, merging algorithm could not find a suitable combination between the lowest packet in the queue

and 9 packets above it, and at that time, the lowest packet in the queue is transmitted without any merging. The airtime requirement in 802.11a/g not only depends on the packet airtime, but also on the medium access time, which equals the DIFS time and a random back-off time. In our analysis, we used the SIGCOMM 2004 dataset [35] to estimate an average medium access time per packet, which equals $2730\mu s$. We consider this time for a single packet transmission for both 802.11a/g and GRaTIS. We used the experimental results, as reported in Table 1, to compute the airtime usage in both of the cases. Since we have not implemented $3/4$ coding rate for GRaTIS, we do not consider those in the trace analysis. With both the coding rates available, there will be more opportunities to merge packets, for example merging a packet of $3/4$ coding rate with another of $1/2$ coding rate. However, we do consider all the coding rates available for 802.11a/g. This is a conservative approach to show the improvement of using GRaTIS over 802.11a/g, and will give a lower bound on the possible gains.

We selected one random monitor on one of the days at SIG-COMM, Monitor 4 on Aug-19, and compute the throughput gain achieved per min using our method. Figure 10 shows the temporal variation per minute of the percentage gain in throughput if GRaTIS is used. The volume of downlink data packets transmitted per minute, and the percentage of packets merged per minute is also shown. A simple moving average (SMA) of 10 data points shows that the average gain goes up to 80%, with instantaneous gain/min. goes up to 90% over 802.11a/g network. The gain is proportional to the percentage of packets being merged. Most of the time, when there is network traffic, GRaTIS could merge $\approx 70\%$ of the packets. We notice that the volume of packets drops significantly at time 100 mins, which is probably the lunch hour, and most of the packets observed by the monitor have very low SNR.
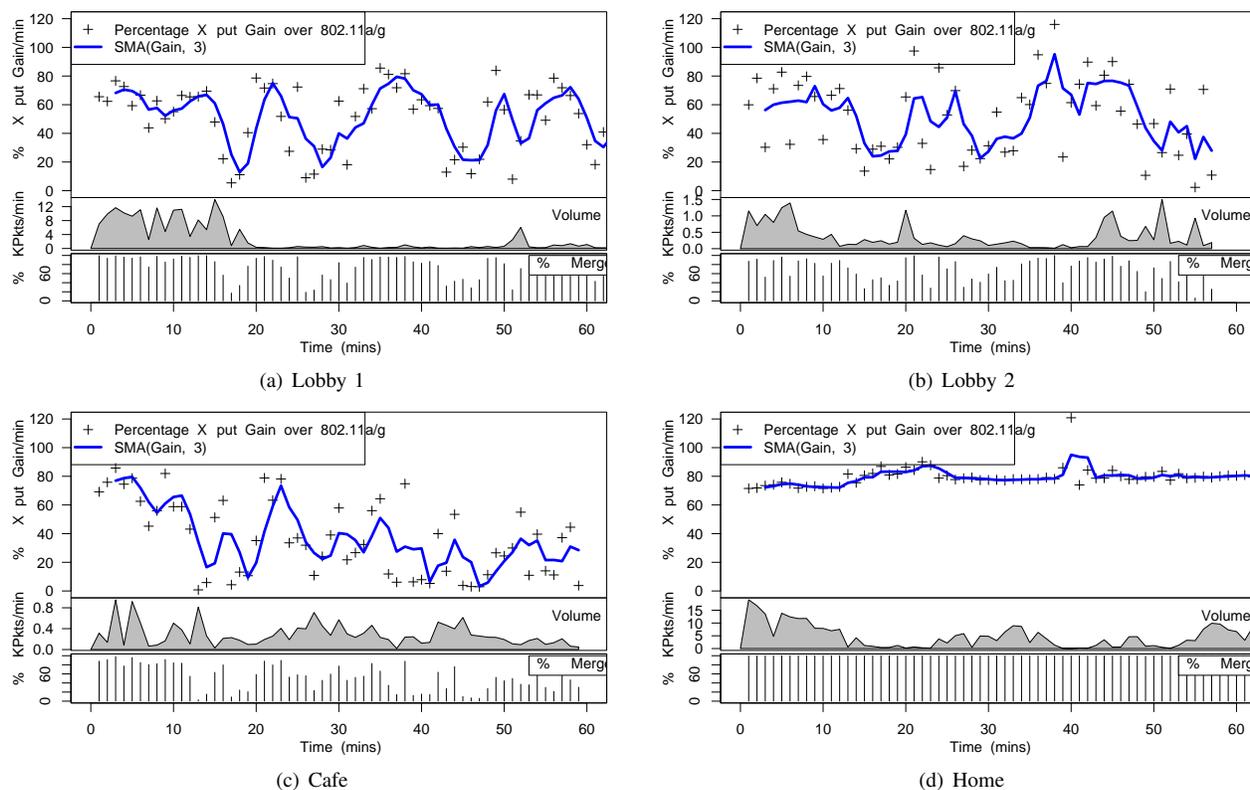
Fig. 11. Gains of using GRaTIS in four different scenarios, based on captured packet traces.

GRaTIS did not find much opportunity to combine packets as there was limited variety in the SNR among the clients, and the volume of packets/min. was extremely low. To ensure that we receive similar gains in other days using the trace from other monitors as well, we computed average gain in each day for each monitor. Results show consistent gains in other scenarios as well.

All networks are not the same, and the traffic pattern varies from one network to another. Hence, we performed similar analysis on our captured dataset to investigate whether GRaTIS can be beneficial in a variety of wireless network scenarios. Figure 11 shows the percentage gain in throughput per minute by using GRaTIS over 802.11a/g network. The trace for 'Lobby 1' shows very high volume for the first 20 mins, which reduces significantly later. However, average gain goes up to 80% when there are enough packets available in the queue to be merged. We captured the trace termed 'Lobby 2' just after completing the capture of 'Lobby 1'. There is significant difference in the volume of downlink data packets. However, GRaTIS still maintains 80% average gain in throughput, with instantaneous gain of up to 120%. We also find that $\approx 70\%$ of the packets have been merged. It is to be noted that GRaTIS will get some savings in DCF and contention times due to aggregating packets, but this gain for concatenating two frames is less than 20%, even in the best case scenario [24]. Nonetheless, aggregated frames can be GRaTIS frames as well, indicating that our technique has potential for even better throughput if we aggregate GRaTIS frames.

The trace termed 'Cafe' corresponds to a very dynamic scenario, where we notice multiple users logging onto the network, using it for few minutes and logging off. Similar to the detailed analysis of the SIGCOMM 2004 trace [34], our analysis shows that the AP spends most of the time in back-off, leading to an overall low throughput on the downlink. Even in this congested scenario, we notice an average gain of 80%. We believe that GRaTIS will perform better in this type of network to reduce the transmission time, and thus more time will be available for data transmission, which will increase individual throughput of any client in the network.

The home network scenario is another common use of 802.11a/g network, where often two or more users in a family share the same wireless network on a regular basis. They often have spatial diversity, leading to variation in SNR. We captured the trace in such a scenario with only two users. The SNR variation, as seen in figure 1(c) is not uniform, and there are only two very high densities in the histogram due to two users. GRaTIS successfully merged more than 99% of the packets on average. However, we notice that most of the time *Gr3* and *Gr4* were used due to the two users being in the operating range of these two GRaTIS rate. The average gain remained constant at 80% with highs of 120%.

This analysis shows that when the clients SNR vary in a diverse range, GRaTIS can be used to combine packets and gain airtime, which can essentially be used to transmit more packets and increase the overall throughput of the network.

(a) Failed1   (b) Failed2

(c) Throughput of Failed1 and Failed2
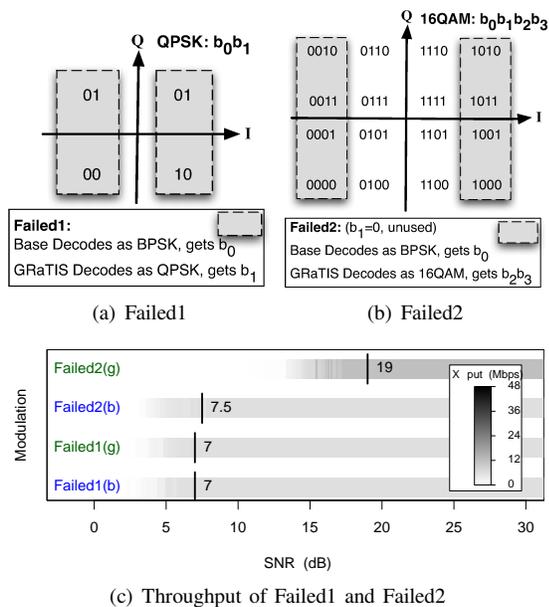
Fig. 12. Example of two "group rates" that don't result in improved performance, showing that group rates must be carefully designed.

6 GRaTIS - A Non-trivial Solution

This paper presents six different types of combinations to generate GRaTIS rates. However, an exhaustive search can be done to find more combinations for any constellation. It is to be noted that this search can be done one time and the results can be stored for reference. Also, a combination from a constellation can be extrapolated to it's corresponding higher constellation, making GRaTIS easy to incorporate for bigger constellations. For example, we can envision GR3 and GR4 for 64QAM as an extension of GR2 for 16QAM. In the same way, it can also be extrapolated for 256QAM as well. However, exhaustive evaluation is necessary to ensure the expected performance gains.

In this paper, we show GRaTIS for two receivers, such that the base layer is transmitted to one receiver, while the GRaTIS layer is transmitted to another receiver. It is possible to generate multiple GRaTIS layers on top of a base layer to transmit to three or more receivers at the same time, as long as each constellation point denote more than 2 bits. In other words, we cannot use QPSK to transmit to three receivers. Nonetheless, it requires careful examination and evaluation to determine whether the generated layers yield better performance.

This paper presents GRaTIS but also opens up a new set of open research problems. While GRaTIS provides more flexibility by providing different intermediate rates and merging multiple packets to enhance the network throughput, the combinations cannot be chosen randomly to ensure better performance. Careful inspection has to be made before selecting any combination. We present two rates *Failed1* and *Failed2* shown in figure 12(a) and 12(b) respectively, chosen randomly from a QPSK and 16QAM constellations. The throughput and SNR requirement of these two GRaTIS rates are shown in
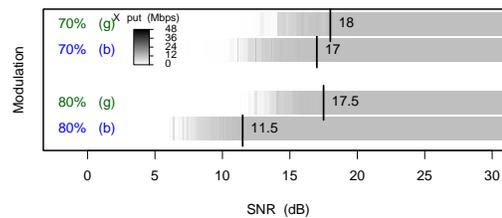
Fig. 13. Over-the-air link throughput for superposition coding with 70% and 80% of total energy allocated to the far node. The numbers against each rate denote the SNR required to decode a 128 byte, $1/2$ rate convolution coded packet with 2% PER.

figure 12(c). Clearly, the base layer of both failed rates provide a BPSK data rate but requires a SNR that can support QPSK data rates. Hence, these two group rates cannot be used to yield improved aggregate throughput as the base layer has a sub-optimal data rate.

In our implementation and analysis discussed in this paper, we have used $1/2$ rate convolution coding. Considering a $3/4$ coding rate we can obtain a similar set of results but at presumably higher SNR, where the base layer and the GRaTIS layer will have a $3/4$ coding rate. By using all the basic rates and their GRaTIS derivatives, we can have almost a linear relationship between SNR and achievable throughput, instead of the conventional step-wise discrete rate allocation. Therefore, from a MAC layer point of view, GRaTIS inspires us to delve into the modalities of a novel rate adaptation algorithm for wireless networks and compare it with related SNR-based rate adaptation algorithms [46], [5], [31], [16]. Also in wireless mesh networks we can forward packets to multiple users from a common point (router/relay node) using GRaTIS. This is particularly helpful when the involved node-pairs cannot over-hear each other's transmission and hence cannot use network coding. We also intend to explore the possibilities of using GRaTIS to improve protocols of higher layers such as opportunistic routing algorithms.

7 Comparing with Superposition Coding

GRaTIS is distinctly different from Superposition Coding (SC) in many ways. SC merges two packets by superimposing one signal on another, which results in mapping the bits in such a way that an error in one layer introduces error in another. GRaTIS uses the bits of existing Gray Code to merge two packets in such a way that the bits of the GRaTIS layer are protected from errors due to symbols crossing either axes due to channel noise. In figure 2(a), it is evident that bits $b_1 b_3$, which encodes the GRaTIS layer of *GR2*, do not change while crossing axes. To compare with our results, we have also implemented SC using our hardware and measured the throughput and PER with the same radio prototype used to evaluate GRaTIS. The SNR requirement for the two layers is shown in figure 13 for 70% and 80% of total energy allocated to the far node. These two modes of SC are the most commonly used energy allocation ratio used in SC [12], [3]. We find from table 1 that *GR2* of GRaTIS outperforms 'SC-70%' in terms of minimum SNR requirement by 6.5dB for the base layer and 3dB for the enhanced or the GRaTIS

layer. However, we noticed that if the energy allocation ratio is $80 : 20$ to the two layers, the mapping thus generated resembles that of 16QAM, which is used to encode *GR2*. Even using the same mapping we notice that *GR2* performs better than 'SC-80%'. This noticeable improvement of GRaTIS over SC is attributed to the unique mapping technique of GRaTIS.

From an implementation stand-point, implementing SC at the transmitter requires changing the I/Q vectors, which needs a high degree of programmability in the radio. Also, the receiver has to perform Successive Interference Cancellation (SIC) to receive the second layer. On the other hand, GRaTIS involves mapping data bits to desired sub-constellations and hence can use the same I/Q vector configurations of the standard 802.11a/g modulation system. It is easier to implement (combine bits) in the existing transmitter pipeline than to super-impose signals. At the receiver end, no complex signal processing algorithms are required – only digital bit manipulation.

GRaTIS provides more combination options compared to SC, provides a fine-grained control over the SNR-throughput space and easily implementable, while outperforming Superposition Coding.

8 RELATED WORK

In this paper we discuss a form of multiuser communication that utilizes information embedded within a modulation constellation. We compare this scheme to other multiuser communications methods currently used in wireless networks. The transmitted signal contains bits for two users and each user can receive, equalize and decode exactly in the same way as done in conventional 802.11a/g packets. We also utilize the existing constellations in 802.11a/g and reinterpret them to encode multiple packets in one data stream. This ensures coexistence with commercial WiFi technologies.

Multiuser Detection: Our implementation and analysis of GRaTIS uses Orthogonal Frequency Division Multiplexing (OFDM), although it is not dependent on OFDM. Simultaneous transmissions can also be detected by the multi-user detection scheme in CDMA. However, in CDMA networks, this technique is expensive to implement and is only implemented by the CDMA base-station and not the mobile devices. The technique requires capturing the waveform data and successively subtracting the signal of individual transmitters [29]. Transmitters use an elaborate power control mechanism to insure the signal received by the base station has a similar SNR for each transmitter. This adds to the latency of CDMA data networks. To reduce the complexity in correlating for different user codes, the authors in [40] use various heuristic methods to obtain a sub-optimal solution. In contrast to these multiuser techniques, GRaTIS requires no control over different users and does not suffer from error propagation due to channel estimation error. While SNR diversity is detrimental to other multiuser techniques, our scheme actually use the difference in SNR experienced by multiple receivers to encode additional data.

Multiuser communication using Orthogonal Frequency Division Multiple Access (OFDMA) is also prevalent in

WiMAX [18] technology. OFDMA has also been introduced in cellular networks as a simultaneous communication mechanism, where separate contiguous sets of subcarriers are assigned for carrying multiuser data [41], [25]. Packets are decoded by multiple users by using a 2-dimensional map in time and frequency. In this way a significant amount of time is saved in contending for the shared wireless channel. However, GRaTIS can still be useful for multi-user systems such as WiMAX although our implementation is based on WiFi. In both WiMAX and WiFi, low rate modulation schemes are used to transmit data to stations with poor SNR; our scheme simply uses varying SNR experienced by multiple receivers to encode additional data.

Hierarchical Coding: Another domain where constellations are used to transmit to multiple users is digital video broadcast. The authors in [32] discuss a layered modulation technique, often termed as *hierarchical modulation* or *multiresolution modulation*. The higher order bits in a dense constellation are used to decode a poorer quality signal that can be decoded by a user with low SNR while users with high SNR will be able to decode all the bits in the symbol. In the DVB digital television standard, 16QAM and QPSK are used for hierarchical modulation. Most of the work involving hierarchical modulation, including [42] and [39], finds its application in a multicast or broadcast environment of digital video.

Hierarchical modulation is used for transmitting the same information (*e.g.*, video signals) to multiple users. GRaTIS can transmit completely different information to multiple users with similar or greater reliability. While prior work in related fields calls for optimizing the signal structure and the signal constellations, our effort is focused on harnessing the strengths embedded in unexplored areas of existing specifications, which makes it novel and yet compatible with coexisting technologies. Also, the wide range of operation of our method (higher number of transmission modes over a range of SNR) puts our scheme in a unique position among its peers and predecessors.

Superposition Coding and Network Coding: Other packet mixing techniques in wireless networks either employ superposition coding [28], [33], [3], [13], [4] or network coding [22], [20], [21], [23] or a combination of the two [27]. Superposition coding relies on iterative decoding by decoding the base layer (lower order bits) first and then re-modulating it to extract the higher layer (higher order bits); this is more complicated to implement, which is why it is difficult to find experimental data for the performance of superposition coding. Also, superposition coding offers less flexibility by limiting the SNR ranges where it can be used with acceptable error performance and often requires very high SNR as the constellation gets denser, which may not be available for any wireless node. In comparison GRaTIS provides a simpler decoder structure and offers more flexibility by providing more data rates for the MAC to select among to merge multiple packets for receivers within realistic SNR ranges. Also under similar operating conditions, GRaTIS provides higher network throughput compared to superposition coding. Network coding relies on the overhearing of a packet which it uses to decode a second packet encoded using network coding techniques. However, network coding fails if the first packet

is not overheard by the intended receiver – under such a case the relay node falls back to multiple packet transmissions, which can be avoided by using GRaTIS. Most 802.11a/g based communication is in the form of an AP with multiple clients. The AP acts as a router between the wireless clients and the wired back-end network. Network coding cannot be used in such environments since the routing medium is not common and packets cannot be overheard. In such cases GRaTIS can still be used to merge downlink packets.

On the other hand, GRaTIS can be combined with network coding. Since network coding simply transmits binary information, either the base or group message can be part of a network coded packet. In fact, GRaTIS increases the applicability of network coding because the network coding message can always be used as a GRaTIS-layer message even if other messages are not queued.

In [44], authors propose mixing bits for relaying purposes, but the mapping of bits to the constellations is different from GRaTIS. The work does not consider the fact that packet mixing will increase the SNR requirements for decoding both the layers. On the contrary, our work is practical and the SNR requirement for each layer has been shown both theoretically and experimentally. The packet trace analysis with variable packet size shows GRaTIS is applicable to most common wireless network scenarios.

9 CONCLUSION

GRaTIS provides an efficient method of simultaneous packet transmission and reception. This increases the network throughput without compromising the throughput of one node while using widespread channel variability to simultaneously transmit an independent packet destined for another node. The GRaTIS packet is indeed extra free bits to the high SNR node, which it would have received after the completion of the first packet using a serialized medium access pattern as in 802.11a/g. We have implemented the protocol in hardware and have shown the ease of implementation if the signal processing is done using a hybrid platform of software and hardware components. The experimental results show several possibilities of use of GRaTIS giving unforeseen gains in throughput in wireless networks. Applying GRaTIS on real-time packet trace analysis reveals that even with a few simple combinations, we can gain significant airtime, which can be further utilized to transmit more packets. Also through our analysis we show that GRaTIS provides more flexibility with better error performance than other contemporary simultaneous packet transmission techniques, making it a suitable candidate for emerging wireless networks.

REFERENCES

[1] Madwifi driver, http://madwifi-project.org/.
[2] Radiotap header, http://www.radiotap.org/.
[3] R. Alimi, L. E. Li, R. Ramjee, H. Viswanathan, and Y. R. Yang. iPack: in-network packet mixing for high throughput wireless mesh networks. In *Proceedings of IEEE INFOCOM*, Phoenix, AZ, Apr. 2008.
[4] S. Bopping and J. Shea. Superposition coding in the downlink of cdma cellular systems. In *Wireless Communications and Networking Conference, 2006. WCNC 2006. IEEE*, volume 4, pages 1978 –1983, 2006.

[5] J. Camp and E. Knightly. Modulation rate adaptation in urban and vehicular environments: cross-layer implementation and experimental evaluation. In *MobiCom '08: Proceedings of the 14th ACM international conference on Mobile computing and networking*, pages 315–326, New York, NY, USA, 2008. ACM.
[6] D. D. Clark, R. Braden, A. Falk, and V. K. Pingali. Fara: reorganizing the addressing architecture. *Computer Communication Review*, 33(4):313–321, 2003.
[7] A. Dutta, J. Fifield, G. Schelle, D. Grunwald, and D. Sicker. An intelligent physical layer for cognitive radio networks. In *WICON '08: Proceedings of the 4th international conference on Wireless internet*, 2008.
[8] A. Dutta, D. Saha, D. Grunwald, and D. Sicker. Smack: a smart acknowledgment scheme for broadcast messages in wireless networks. *SIGCOMM Comput. Commun. Rev.*, 39(4):15–26, 2009.
[9] A. Dutta, D. Saha, D. Grunwald, and D. Sicker. An architecture for software defined cognitive radio. In *Proceedings of the 6th ACM/IEEE Symposium on Architectures for Networking and Communications Systems*, ANCS '10, pages 5:1–5:12, New York, NY, USA, 2010. ACM.
[10] J. Fifield, P. Kasemir, D. Grunwald, and D. Sicker. Experiences with a platform for frequency agile techniques. In *DYSPAN*, 2007.
[11] P. Frenger, P. Orten, and T. Ottosson. Convolutional codes with optimum distance spectrum. *Communications Letters, IEEE*, 3(11):317 –319, Nov. 1999.
[12] R. Ganti, Z. Gong, M. Haenggi, C. Lee, S. Srinivasa, D. Tisza, S. Vanka, and P. Vizi. Implementation and experimental results of superposition coding on software radio. In *Communications (ICC), 2010 IEEE International Conference on*, pages 1 –5, May 2010.
[13] R. K. Ganti, Z. Gong, M. Haenggi, C.-H. Lee, S. Srinivasa, D. Tisza, S. Vanka, and P. Vizi. Implementation and Experimental Results of Superposition Coding on Software Radio. In *2010 IEEE International Conference on Communications (ICC'10)*, Cape Town, South Africa, May 2010.
[14] D. Haccoun and G. Begin. High-rate punctured convolutional codes for viterbi and sequential decoding. *Communications, IEEE Transactions on*, 37(11):1113 –1125, Nov. 1989.
[15] D. Halperin, T. Anderson, and D. Wetherall. Taking the sting out of carrier sense: interference cancellation for wireless lans. In *Proceedings of the 14th ACM international conference on Mobile computing and networking*, MobiCom '08, pages 339–350, New York, NY, USA, 2008. ACM.
[16] D. Halperin, W. Hu, A. Sheth, and D. Wetherall. Predictable 802.11 packet delivery from wireless channel measurements. In S. Kalyanaraman, V. N. Padmanabhan, K. K. Ramakrishnan, R. Shorey, and G. M. Voelker, editors, *SIGCOMM*, pages 159–170. ACM, 2010.
[17] IEEE Computer Society : LAN/MAN Standards Committee. *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*.
[18] IEEE Computer Society : LAN/MAN Standards Committee. *Part 16: Air Interface for Broadband Wireless Access Systems*.
[19] S. Katti, S. Gollakota, and D. Katabi. Embracing wireless interference: Analog network coding. In *in ACM SIGCOMM*, pages 397–408. MIT, 2007.
[20] S. Katti, S. Gollakota, and D. Katabi. Embracing wireless interference: analog network coding. In *SIGCOMM '07: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 397–408, New York, NY, USA, 2007. ACM.
[21] S. Katti, D. Katabi, H. Balakrishnan, and M. Medard. Symbol-level network coding for wireless mesh networks. *SIGCOMM Comput. Commun. Rev.*, 38(4):401–412, 2008.
[22] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft. Xors in the air: practical wireless network coding. *SIGCOMM Comput. Commun. Rev.*, 36(4):243–254, 2006.
[23] M. Kim, D. Lucani, X. Shi, F. Zhao, and M. Medard. Network Coding for Multi-Resolution Multicast. *ArXiv e-prints*, Aug. 2009.
[24] Y. Kim, S. Choi, K. Jang, and H. Hwang. Throughput enhancement of ieee 802.11 wlan via frame aggregation. In *Vehicular Technology Conference, 2004. VTC2004-Fall. 2004 IEEE 60th*, volume 4, pages 3030 – 3034 Vol. 4, sept. 2004.
[25] F. Kojima, H. Harada, and M. Fujise. Adaptive sub-carriers control scheme for ofdm cellular systems. *Vehicular Technology Conference Proceedings, 2000. VTC 2000-Spring Tokyo. 2000 IEEE 51st*, 2:1065–1069 vol.2, 2000.
[26] F. Li, M. Li, R. Lu, H. Wu, C. Mark, and K. Robert. Measuring queue capacities of ieee 802.11 wireless access points. In *Broadband*

Communications, Networks and Systems, 2007. BROADNETS 2007. Fourth International Conference on, pages 846 –853, sept. 2007.

[27] C.-H. Liu and A. Arapostathis. Joint network coding and superposition coding for multi-user information exchange in wireless relaying networks. In *Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008. IEEE*, pages 1 –6, 30 2008-dec. 4 2008.

[28] E. S. Lo and K. B. Letaief. Network coding versus superposition coding for two-way wireless communication. In *WCNC'09: Proceedings of the 2009 IEEE conference on Wireless Communications & Networking Conference*, pages 307–311, Piscataway, NJ, USA, 2009. IEEE Press.

[29] P. Patel and J. Holtzman. Analysis of a simple successive interference cancellation scheme in a ds/cdma system. *Selected Areas in Communications, IEEE Journal on*, 12(5):796 –807, June 1994.

[30] D. Qiao, S. Choi, and K. Shin. Goodput analysis and link adaptation for ieee 802.11a wireless lans. volume 1, pages 278 – 292, 2002.

[31] H. Rahul, N. Kushman, D. Katabi, C. Sodini, and F. Edalat. Learning to Share: Narrowband-Friendly Wideband Networks. In *SIGCOMM '08: Proceedings of the ACM SIGCOMM 2008 conference on Data communication*, pages 147–158, New York, NY, USA, 2008. ACM.

[32] K. Ramchandran, A. Ortega, K. Uz, and M. Vetterli. Multiresolution broadcast for digital hdtv using joint source-channel coding, Jun 1992.

[33] R. Ramjee, J. Shi, Y. Sun, H. Viswanathan, and Y. R. Yang. Extended abstract: Superposition coding for wireless mesh networks abstract.

[34] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorian. Measurement-based characterization of 802.11 in a hotspot setting. In *Proceedings of the ACM SIGCOMM 2005 Workshop on experimental approaches to wireless network design and analysis (E-WIND-05)*, Philadelphia, PA, August 2005.

[35] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, J. Zahorjan, and E. Lazowska. CRAWDAD data set uw/sigcomm2004 (v. 2006-10-17). Downloaded from http://crawdad.cs.dartmouth.edu/uw/sigcomm2004, Oct. 2006.

[36] A. Schulman, D. Levin, and N. Spring. CRAWDAD trace set umd/sigcomm2008/pcap (v. 2009-03-02). Downloaded from http://crawdad.cs.dartmouth.edu/umd/sigcomm2008/pcap, Mar. 2009.

[37] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi. Accurate: Constellation based rate estimation in wireless networks. In *NSDI*, pages 175–190. USENIX Association, 2010.

[38] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi. Successive interference cancellation: a back-of-the-envelope perspective. In *Proceedings of the Ninth ACM SIGCOMM Workshop on Hot Topics in Networks*, Hotnets '10, pages 17:1–17:6, New York, NY, USA, 2010. ACM.

[39] N. Shacham. Multipoint communication by hierarchically encoded data. In *INFOCOM '92. Eleventh Annual Joint Conference of the IEEE Computer and Communications Societies, IEEE*, pages 2107–2114 vol.3, May 1992.

[40] P. H. Tan and L. Rasmussen. Multiuser detection in cdma - a comparison of relaxations, exact, and heuristic search methods. *Wireless Communications, IEEE Transactions on*, 3(5), Sept. 2004.

[41] T. Thanabalasingham, S. Hanly, L. Andrew, and J. Papandriopoulos. Joint allocation of subcarriers and transmit powers in a multiuser ofdm cellular network. *Communications, 2006. ICC '06. IEEE International Conference on*, 1:269–274, June 2006.

[42] K. Uz, M. Vetterli, and D. LeGall. Interpolative multiresolution coding of advance television with compatible subchannels. *Circuits and Systems for Video Technology, IEEE Transactions on*, 1(1):86–99, Mar 1991.

[43] M. Vutukuru, H. Balakrishnan, and K. Jamieson. Cross-Layer Wireless Bit Rate Adaptation. In *ACM SIGCOMM*, Barcelona, Spain, August 2009.

[44] Z. Yang, Y. Luo, and L. Cai. Network modulation: A new dimension to enhance wireless network performance. In *INFOCOM, 2011 Proceedings IEEE*, pages 2786 –2794, april 2011.

[45] Y. Yasuda, K. Kashiki, and Y. Hirata. High-rate punctured convolutional codes for soft decision viterbi decoding. *Communications, IEEE Transactions on*, 32(3):315 – 319, Mar. 1984.

[46] J. Zhang, K. Tan, J. Zhao, H. Wu, and Y. Zhang. A practical snr-guided rate adaptation. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 2083–2091, April 2008.

Dola Saha Dola Saha has a Masters degree in Computer Science in the University of Colorado Boulder. Currently, she is a PhD candidate in the same university. Her research interest is broadly in the area of wireless networks with primary focus on designing complete systems and its practical implementation. She emphasizes on utilizing physical layer concepts to improve higher layer protocols. Apart from MAC-PHY cross-layer research, she has also worked on directional antennas and its application in broader wireless networking domain.

Aveek Dutta Aveek Dutta is a PhD candidate in the department of Electrical Computer and Energy Engineering in the University of Colorado Boulder. His current research is focused on defining novel architectures for Software Defined Radios (SDR). During his initial years as a graduate student he worked on design and implementation of a prototype OFDM transceiver on FPGA. This prototype has been used to innovate new wireless network protocols, OFDMA based PHY layer signaling, Digital Communication and signal processing applications and also as a foundation for future architectures of cognitive radio and SDRs.

Dirk Grunwald Dirk Grunwald is a Professor in the Department of Computer Science with a dual appointment in the Dept. of Electrical and Computer Engineering at the University of Colorado. Dr. Grunwald's research area includes networking, wireless, computer architecture and computer systems design, including the prototyping of systems such as the one described in this paper. Dr. Grunwald is a member of the ACM and IEEE.

Douglas Sicker Dr. Douglas C. Sicker is a professor in the Department of Computer Science at the University of Colorado at Boulder and the Director of the Interdisciplinary Telecommunications Program. Recently, Doug was the Chief Technology Officer at both the National Telecommunications and Information Administration (NTIA) and the Federal Communications Commission (FCC) and served as a senior advisor on the FCC National Broadband Plan. Previously he was Director of Global Architecture at Level 3 Communications and a Division Chief at the FCC. Doug is a senior member of the IEEE, as well as a member of the ACM and the Internet Society.

APPENDIX

In this section we present an example of BER calculation for an overlapped modulated packet. In GRaTIS the transmitted constellation is decoded using a different set of decision boundaries, which yields a constellation of smaller size. For example, a transmitted constellation of 16QAM can be decoded as a QPSK as done in *GR2* base layer. In such a case the BER for the base layer will be between QPSK and 16QAM. For a 16QAM constellation the possible constellation point sets are $\{\pm A, \pm A\}$, $\{\pm 3A, \pm A\}$, $\{\pm A, \pm 3A\}$ and $\{\pm 3A, \pm A\}$. A is a modulation dependent parameter given by $\sqrt{\frac{K_b E_b}{K_{mod}}}$, where K_b is the number of bits per symbol, E_b is the energy per bit and K_{mod} is a scaling factor so that all constellation points have unit energy: for $16QAM$ this scaling factor is $\sqrt{10}$ [17]. Thus $A = \sqrt{\frac{K_b E_b}{10}} = \sqrt{\frac{E_s}{10}}$, where E_s is the energy per symbol. Figure 2(a) shows a 16QAM constellation and the decision boundaries for a QPSK decoding. The constellation points are marked $[0 \dots 15]$. We recall from §2, that for the base layer a jump from one signal point to another within a quadrant will not cause any bit error since all the points in one quandrant are mapped to one QPSK constellation point. Therefore, we start by identifying the pair-wise squared Euclidean distances between transmitted constellation points that will cause a bit error.

$$\Delta_E(0,3) = 36A^2 \tag{4}$$

$$\Delta_E(0,15) = 36A^2 \tag{5}$$

$$\Delta_E(1,2) = 4A^2 = \Delta_E(7,8) \tag{6}$$

$$\Delta_E(1,14) = 36A^2 = \Delta_E(7,4) \tag{7}$$

$$\Delta_E(6,5) = 4A^2 \tag{8}$$

$$\Delta_E(6,9) = 4A^2 \tag{9}$$

Now, we can compute the conditional error probabilities of the points in the top-right quandrant *viz.* $0, 1, 6, 7$ using eq. 1 as follows,

$$P_{WUB}(E|\vec{I}=0) = \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{\Delta_E(0,3)}{4N_0}}\right) + \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{\Delta_E(0,15)}{4N_0}}\right)$$
$$= \operatorname{erfc}\left(\sqrt{\frac{9A^2}{N_0}}\right) \tag{10}$$

$$P_{WUB}(E|\vec{I}=1) = \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{\Delta_E(1,2)}{4N_0}}\right) + \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{\Delta_E(1,14)}{4N_0}}\right)$$
$$= \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{A^2}{N_0}}\right) + \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{9A^2}{N_0}}\right) \tag{11}$$

$$P_{WUB}(E|\vec{I}=6) = \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{\Delta_E(6,5)}{4N_0}}\right) + \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{\Delta_E(6,9)}{4N_0}}\right)$$
$$= \operatorname{erfc}\left(\sqrt{\frac{A^2}{N_0}}\right) \tag{12}$$

Combining eq. 10, 11 and 12 we get,

$$P_{WUB}(E) = \frac{4}{16}P_{WUB}(E|\vec{I}=0) + \frac{8}{16}P_{WUB}(E|\vec{I}=1)$$
$$+ \frac{4}{16}P_{WUB}(E|\vec{I}=6) \tag{13}$$
$$= \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{9A^2}{N_0}}\right) + \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{A^2}{N_0}}\right)$$

Substituting the value of A,

$$P_{WUB}(E) = \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{9E_s}{10N_0}}\right) + \frac{1}{2}\operatorname{erfc}\left(\sqrt{\frac{E_s}{10N_0}}\right) \tag{14}$$

Similar BER can be obtained for any constellation and any decision boundary. We compute the theoretical PER using these results and plot against increasing symbol to noise ratio (E_s/N_0) in figure 5.