Full length article

PHY modulation/rate control for fountain codes in 802.11a/g WLANs☆

CrossMark

Xiaomin Chen [a,*], Vijay G. Subramanian [b], Douglas J. Leith [a]

[a] Hamilton Institute, NUI Maynooth, Ireland
[b] EECS Department, Northwestern University, USA

ARTICLE INFO

ABSTRACT

In this paper we consider the joint performance of fountain codes and 802.11a/g PHY modulation/coding. We consider optimality both in terms of maximizing goodput and minimizing energy, and results are presented for both theoretical and experimental channel models. In contrast to studies in cellular networks, we find that in 802.11a/g WLANs the cross-layer approach of a higher-layer fountain coding with a PHY layer modulation and FEC coding can yield very limited gains, and the PHY modulation/rate that optimizes the uncoded multicast performance is also close to that for fountain-coded multicast traffic over a wide-range of network conditions. This is potentially an important observation as it indicates that in 802.11a/g WLANs cross-layer design for multicast rate control would bring few benefits and PHY layer rate control can be carried out without regard to the use of fountain coding at higher layers.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

In this paper we consider the joint performance of fountain codes and 802.11a/g PHY modulation/rate. Fountain codes have been the subject of much interest in recent years, both in the context of wireless video multicast (e.g. see [1–3]) and of network coding and joint coding/routing (e.g. see [4]). However, although there is a wealth of literature on the subject of fountain codes, there have been relatively few studies of fountain code performance in 802.11 WLANs and even fewer on the cross-layer trade-offs between higher-layer fountain coding and PHY layer

modulation and forward error correction (FEC) coding selection. The use of fountain codes in 802.11 WLANs is nevertheless of considerable interest in view of their ubiquitous deployment and the trend towards their use for multimedia distribution within the home and elsewhere.

Fountain coding is a higher layer technology since it makes use of an erasure channel abstraction where received symbols are labelled to indicate whether they are erased (i.e. corrupted by errors) or error-free. In a packet-switched network, the fountain code symbols are typically packets and a cyclic redundancy check (CRC) checksum is used to mark each packet as either erased or error-free. This abstraction essentially provides an interface to the PHY layer which in turn interfaces with the actual wireless channel. In the erasure channel model the probability that a packet is erased is strongly coupled to the choice of modulation and FEC coding scheme used at the PHY layer. To understand system performance, it is therefore necessary to consider the joint fountain code and PHY modulation/coding performance.

Some prior work has been done on the joint performance of a concatenated higher layer fountain code and

PHY modulation and FEC code. Paper [5] looks at the trade-off in a general single-user wireless setting. In our setting we consider multicast rather than single-user operation and include the specific constraints imposed by 802.11 technology. Paper [6] considers multicast and the use of Raptor codes in 802.11-like WLANs. Their analysis is based on an access mechanism different from the IEEE 802.11 standards, and it does not consider the rate-selection problem in comparison to the uncoded setting. Paper [7] considers a fundamentally different multi-receiver scenario from the standard multicast setting with the objective for a group of M client stations to cooperatively receive a block of N packets rather than for each individual client to receive all N packets successfully. To our knowledge the present paper is the first detailed study on the choice PHY modulation/rate with and without fountain codes in standard 802.11 WLANs.

In contrast to studies in cellular networks [8,9], we find that in 802.11a/g WLANs the PHY modulation/rate that optimizes uncoded multicast performance is also close to that for fountain-coded multicast traffic. This is potentially an important observation as it indicates that in 802.11a/g WLANs cross-layer design for multicast rate-control would bring few benefits and PHY layer rate-control can be carried out without regard to the use of fountain coding at higher layers.

The rest of this paper is organized as follows. Analytical expressions to determine the packet error rate (PER) given a signal to noise ratio (SNR) value for each of the PHY rates provided by 802.11a/g in both AWGN and Rayleigh channels are given in Section 2. Given a fixed set of PER and PHY rate, the goodput performances of multicast traffic with and without fountain coding are evaluated in Section 3. Combining the analysis in Section 2 and Section 3, the choices of PHY rates which respectively maximize goodput and minimize energy are determined using numerical search method in Section 4 and Section 5. Results are presented for both theoretical and experimental channel models. Discussions on the model setting and generalizations are given in Section 6.

2. 802.11a/g PHY modulation and FEC

We begin by briefly reviewing the modulation and FEC choices supported by the 802.11a/g PHY (802.11a and 802.11g share identical modulation and FEC schemes). The 802.11a/g PHY layer is OFDM based. A high-speed binary data stream is encoded by a convolutional encoder and then partitioned into low-speed substreams, with each substream then being modulated and transmitted over an orthogonal subchannel. At the receiver side, after demodulation, a Viterbi decoder performs hard or soft decision decoding for the convolutional coded bit sequence. In terms of the convolutional code rate and the modulation scheme used in each subcarrier, 802.11a/g provides eight PHY rates. Given a physical channel model, we can derive the relationship between the packet error rate (PER) and signal to noise ratio (SNR) for each modulation/coding choice as follows.

2.0.1. AWGN channel

The Viterbi algorithm searches through the trellis of a convolutional code to determine a path that maximizes the probability that the received signals represent the coded bit sequence expressed by this path. For this purpose, the decoder associates a metric for each path. Depending on whether the Viterbi decoder performs a hard or soft decision, the corresponding metric is either a Hamming metric or a Euclidean metric, respectively. An upper bound on the PER in Viterbi decoding is given by [10]

$$P_p < 1 - (1 - P_e)^L \tag{1}$$

where L is the length of packet in bits, and P_e is the union bound on the first-event error probability, which is calculated as follows:

- Viterbi hard decision decoding (HDD):
  When Viterbi HDD is used, analytic expressions for the bit error rate (BER) in an additive white Gaussian noise (AWGN) channel when using each of the various 802.11a/g modulation schemes are described in [11]. As the demodulator is followed by the Viterbi decoder, we need to adjust the demodulation BER to take account of the error correction provided by convolutional coding. With Viterbi HDD, P_e is given by [12]

$$P_{e\text{HDD}} \leq \sum_{d=d_{\text{free}}}^{\infty} \alpha_d \cdot P_2(d) \tag{2}$$

where d_{free} is the minimum free distance of the convolutional code, α_d is the total number of paths with degree d, and $P_2(d)$ is the probability that an incorrect path with degree d is selected, which is given by

$$P_2(d) = \begin{cases} \displaystyle\sum_{k=(d+1)/2}^{d} \binom{d}{k} \cdot \rho^k \cdot (1-\rho)^{d-k} & \text{if } d \text{ is odd} \\ \frac{1}{2} \cdot \binom{d}{d/2} \cdot \rho^{d/2} \cdot (1-\rho)^{d/2} \\ \quad + \displaystyle\sum_{k=d/2}^{d} \binom{d}{k} \cdot \rho^k \cdot (1-\rho)^{d-k} \\ \quad\quad \text{if } d \text{ is even} \end{cases} \tag{3}$$

with ρ being the HDD demodulation BER, which is a function of the receiving SNR.

- Viterbi soft decision decoding (SDD):
  Since the path metric in Viterbi SDD is based on the Euclidean distance, the probability of choosing a wrong path in a pairwise comparison with the all-zero path depends not only on the degree d, but on the modulation scheme used in the communication system, i.e. the series of symbols constructed along that path. The union bound on the first-event error probability in SDD is then expressed as

$$P_{e\text{SDD}} \leq \sum_{d=d_{free}}^{\infty} \sum_{i=1}^{\alpha_d} P_2(d_s^i(d)) \tag{4}$$

where $P_2(d_s^i(d))$ is the probability that the ith d-degree path with $d_s^i(d)$ different symbols from the all-zero path is wrongly chosen.
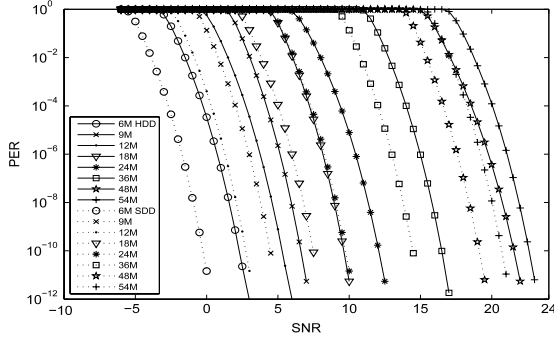
(a) AWGN channel.
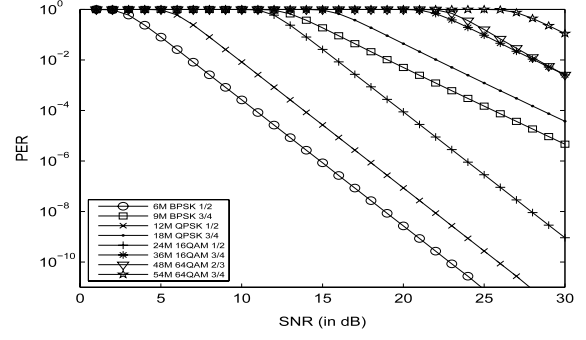
(b) Rayleigh channel.

Fig. 1. 802.11a/g packet error rate (PER) vs. channel SNR for the range of modulation and coding schemes available.

The derivation of $P_2(d_s^i(d))$ for a coherent M-ary QAM demodulator in an AWGN channel is described in [13], given by,

$$
P_2(d_s^i(d)) = Q\left(\sqrt{\frac{\sum_{l=1}^{d_s^i(d)} \|C_l - C_0\|^2}{2N_0}}\right)
$$

$$
= Q\left(\sqrt{\frac{\sum_{l=1}^{d_s^i(d)} D_l^2}{2N_0}}\right) \tag{5}
$$

where C_0 is the all-zero symbol; C_l is the lth symbol that differs from C_0 along the path; N_0 is the noise spectral density; $D_l = \|C_l - C_0\|$ is the Euclidean distance between a non-zero symbol C_l and the all-zero symbol C_0 in the constellation.

In 802.11a/g standards, the convolutional coded bit sequence is converted into constellation points by Grey-coded constellation mapping [14]. For BPSK and QPSK, the union bound on the SDD first-event error probability is given by

$$
P_{eSDD}(\gamma_b)_{BPSK/QPSK} \le \sum_{d=d_{free}}^{\infty} \alpha_d Q\left(\sqrt{2d\gamma_b}\right) \tag{6}
$$

where γ_b is the receiver side SNR per bit.

For 16-QAM and 64-QAM, since the value of $P_2(d_s^i(d))$ is not equal for every d-degree path, an upper bound on it, i.e. the highest probability that a path with degree d is selected in SDD, can be established by determining a combination of symbols that minimize $\sum_{l=1}^{d_s^i(d)} D_l^2$. Since symbols modulating the same number of bit "1"s may have different Euclidean distances to C_0 in the constellation, we actually need to minimize $\sum_{l=1}^{d_s^i(d)} D_{min}^l{}^2$, where D_{min}^l is the minimum Euclidean distance to C_0 among symbols modulating the same number of bit

"1"s as symbol C_l. Therefore, the union bound on the SDD first-event error probability for 16-QAM can be upper-bounded by

$$
P_{eSDD}(\gamma_b)_{16-QAM} \le \sum_{d=d_{free}}^{\infty} \alpha_d Q\left(\sqrt{\frac{4}{5}d\gamma_b}\right) \tag{7}
$$

and the union bound on the SDD first-event error probability for 64-QAM can be upper-bounded by

$$
P_{eSDD}(\gamma_b)_{64-QAM} \le \sum_{d=d_{free}}^{\infty} \alpha_d Q\left(\sqrt{\frac{2}{7}d\gamma_b}\right). \tag{8}
$$

The receiving SNR γ can be expressed by the SNR per bit γ_b, given by $\gamma = \gamma_b \frac{R}{B}$, in which R is the transmission PHY rate, and B is the channel bandwidth, which is 20 MHz in the 802.11a/g OFDM PHY [14]. Fig. 1(a) shows the resulting HDD and SDD PER versus γ for each of the 802.11a/g modulation/coding schemes for a packet length of 1024 bytes. It can be seen that Viterbi SDD provides a lower PER than HDD for any given SNR.

2.0.2. Rayleigh channel

The hard-decision demodulation BER is calculated following the analysis in [15] for Nakagami-m fast-fading channels, i.e. each symbol experiences an independent fading realization. We note that a Rayleigh channel corresponds to a Nakagami channel with $m = 1$. The average hard-decision BER of BPSK in a Nakagami-m fast fading channel is given by

$$
\rho_{BPSK} = \frac{1}{2}\left(1 - \mu \sum_{k=0}^{m-1} \binom{2k}{k}\left(\frac{1-\mu^2}{4}\right)^k\right) \tag{9}
$$

where

$$
\mu = \sqrt{\frac{\overline{\gamma}_s}{m + \overline{\gamma}_s}} \tag{10}
$$

with $\overline{\gamma}_s$ being the average SNR per symbol.

The average hard-decision BER of M-QAM ($M \geq 4$) in Nakagami-m fast fading channels is given by

$$\rho_{M\text{-QAM}} = 4\left(\frac{\sqrt{M}-1}{\sqrt{M}}\right)\left(\frac{1}{\log_2 M}\right)$$

$$\times \sum_{i=1}^{\sqrt{M}/2} \frac{1}{2}\left(1 - \mu_i \sum_{k=0}^{m-1}\binom{2k}{k}\left(\frac{1-\mu_i^2}{4}\right)^k\right) \quad (11)$$

where

$$\mu_i = \sqrt{\frac{1.5(2i-1)^2\overline{\gamma}_s}{m(M-1) + 1.5(2i-1)^2\overline{\gamma}_s}}. \quad (12)$$

With the average demodulation BER, the HDD PER can be calculated in the same way as for an AWGN channel. The receiving SNR γ can be expressed by the average SNR per symbol $\overline{\gamma}_s$, given by $\gamma = \frac{\overline{\gamma}_s}{k} \cdot \frac{R}{B}$, in which k is the number of bits modulated per symbol. Fig. 1(b) shows the corresponding HDD PER versus γ for a Rayleigh channel.

The analysis for SDD PER for a Rayleigh channel is too complicated. In this paper we only consider Viterbi SDD for an AWGN channel. It is also worth noting that for a packet payload size of 992 bytes, as considered in the following numerical analysis, the upper bound on the HDD first-event error probability P_e is always tight for both the AWGN and Rayleigh channels [16].

3. Performance modelling

We consider an 802.11a/g single-hop downlink multicast network with one access point (AP) and M client stations. Without higher-layer coding, to achieve reliable multicast it is necessary for each client station to transmit higher-layer acknowledgement packets to inform the multicast sender of which packets were noise-corrupted and so need to be retransmitted. As the number M of client stations increases, the probability of a given packet being successfully received over a noisy channel by all stations often becomes small and we may quickly end up in a situation where almost every packet requires to be retransmitted at least once and perhaps multiple times. In contrast, fountain coding allows a block of N packets to be recovered, on average, from reception of *any* $N + \delta$ coded packets, where δ is the decoding overhead counted in terms of number of extra packets that need to be transmitted for decoding with high probability. Fountain coding therefore fundamentally changes the scaling behaviour of network performance with the number M of client stations. In this section we derive analytic expressions for the mean number of transmissions and acknowledgements required for all M clients to successfully receive a block of N multicast packets over a noisy channel with and without fountain coding.

3.1. Fountain encoding

We consider both systematic and non-systematic fountain codes and so the analysis encompasses both Raptor and LT codes. The non-systematic fountain code we use in our numerical results is the random linear fountain code [17]. Source packets are selected in terms of a Bernoulli (1/2) random vector and then summed, modulo 2, to construct a coded packet. The systematic fountain code we consider in the numerical results is proposed in [18], which uses uncoded source packets for the first N transmissions and the subsequent transmissions are coded packets constructed using an equiprobable random linear fountain code. It is shown in [18] that this class of fountain codes stochastically minimizes the number of received packets necessary for recovery of coded packets over a large class of fountain codes including Raptor and LT codes, and has low decoding complexity for small block sizes.

In our proposal fountain coding is carried out at the application layer. An application file is first divided into multiple equal-length blocks, and each block is then individually encoded. A block sequence number is included in the header of each coded packet to indicate which block it belongs to. In addition, the header contains the pseudo-random seed used to generate the Bernoulli (1/2) random vector associated with a coded packet, which is needed to reconstruct the generator matrix for decoding. We account for these overheads in our analysis, see Section 3.6, in addition to the coding overhead.

3.2. Higher-layer ACK/NACK modelling

For reliable multicast without fountain coding it is necessary for clients to use higher-layer ACK/NACK transmissions to inform the multicast sender of which packets have been successfully received. To derive the corresponding goodput expression we assume the use of the following signalling scheme. Namely, after transmission of a block of N packets, each client checks whether it has successfully received the whole block or not. If it has, an application layer ACK is transmitted to inform the multicast sender. Otherwise, an application layer NACK is transmitted which identifies the missing packets. Since our interest is in obtaining an upper bound on performance, we assume that higher-layer ACK/NACK transmissions are scheduled in such a way that they never collide and that their transmissions are error-free. At the next round, the multicast sender retransmits the union of any packets that were not received, and this is repeated until all clients have received the block of N packets.

When fountain coding is used, we assume that each client transmits an application layer ACK on successfully decoding a block of N packets. No further signalling is required, and again it is assumed that the higher-layer ACKs never collide and that their transmissions are error-free.

3.3. Mean #transmissions without higher-layer fountain coding

Let p denote the PER—given the channel SNR this can be obtained from Fig. 1 following the analysis presented in Section 2. We assume that every client has the same p at each PHY rate. The probability that client $j \in \{1, 2, \ldots, M\}$

receives packet $i \in \{1, 2, \ldots, N\}$ after $k \geq 1$ transmissions is

$$P(r_{i,j} = k) = p^{k-1}(1 - p). \tag{13}$$

That is, given a pair i and j, the number of transmissions $r_{i,j}$ is a Geometric$(1 - p)$ random variable. The number of transmissions required before the ith packet is received by all M clients is the random variable

$$t(i) = \max_{j \in \{1,2,\ldots,M\}} r_{i,j}. \tag{14}$$

The total number of transmissions to receive the full block of N packets is

$$T_{ucd} = \sum_{i=1}^{N} \max_{j \in \{1,2,\ldots,M\}} r_{i,j}. \tag{15}$$

The mean of the total number of transmissions is given by

$$
\begin{aligned}
E[T_{ucd}] &= \sum_{i=1}^{N} E[\max_{j \in \{1,2,\ldots,M\}} r_{i,j}] \\
&= N\left(1 + \sum_{t=1}^{\infty}\left(1 - \left(\sum_{x=1}^{t}(1-p)p^{x-1}\right)^{M}\right)\right) \\
&= N\left(1 + \sum_{t=1}^{\infty}\left(1 - (1 - p^t)^M\right)\right) \\
&= N\left(1 + \sum_{t=1}^{\infty}\sum_{i=1}^{M}\binom{M}{i}(-1)^{i+1}(p^t)^i\right) \\
&= N\left(1 + \sum_{i=1}^{M}\binom{M}{i}(-1)^{i+1}\sum_{t=1}^{\infty}(p^i)^t\right) \\
&= N\left(1 + \sum_{i=1}^{M}\binom{M}{i}(-1)^{i+1}\frac{p^i}{1 - p^i}\right). \tag{16}
\end{aligned}
$$

The number of higher-layer acknowledgement packets (including ACKs and NACKs) is

$$A_{ucd} = \sum_{j=1}^{M} \max_{i \in \{1,2,\ldots,N\}} r_{i,j} \tag{17}$$

and the mean of it is then given by

$$
\begin{aligned}
E[A_{ucd}] &= M\left(1 + \sum_{t=1}^{\infty}\left(1 - \left(\sum_{x=1}^{t}(1-p)p^{x-1}\right)^{N}\right)\right) \\
&= M\left(1 + \sum_{t=1}^{\infty}\left(1 - (1 - p^t)^N\right)\right) \\
&= M\left(1 + \sum_{i=1}^{N}\binom{N}{i}(-1)^{i+1}\frac{p^i}{1 - p^i}\right). \tag{18}
\end{aligned}
$$

3.4. Mean #transmissions with non-systematic fountain code

With fountain coding, to receive a block of N packets the requirement becomes that every client has on average to successfully receive $N + \delta$ packets, where δ is the coding overhead (usually a fixed sublinear function of N depending on the code). The probability that client j receives $N + \delta$ packets after $k \geq N + \delta$ transmissions is

$$P(Y_j = k) = \binom{k-1}{N + \delta - 1} p^{k-N-\delta}(1 - p)^{N+\delta}. \tag{19}$$

Thus Y_j is a Negative Binomial $(N + \delta, 1 - p)$ random variable. The number of transmissions required for all M clients to successfully decode a block of N packets is the random variable

$$T_{nonsys} = \max_{j \in \{1,2,\ldots,M\}} Y_j. \tag{20}$$

The mean number of transmissions is then given by

$$
\begin{aligned}
E[T_{nonsys}] &= N + \delta + \sum_{t=N+\delta}^{\infty}\left(1 - (1 - p)^{(N+\delta)M}\right. \\
&\qquad \left. \times \left(\sum_{x=N+\delta}^{t}\binom{x-1}{N+\delta-1}p^{x-N-\delta}\right)^{M}\right) \\
&= N + \delta + \sum_{t=N+\delta}^{t'}\left(1 - (1 - p)^{(N+\delta)M}\right. \\
&\qquad \left. \times \left(\sum_{x=N+\delta}^{t}\binom{x-1}{N+\delta-1}p^{x-N-\delta}\right)^{M}\right) \\
&\quad + \mathcal{O}(1) \tag{21}
\end{aligned}
$$

for some large t' and where $\mathcal{O}(1) \to 0$ as $t' \to \infty$. Unfortunately, we cannot analytically derive closed-form expressions for $E[T_{nonsys}]$ due to the series not being any standard known series. It is, however, easy to see that the series is summable so in this paper we use the summation of a finite number of terms to t' as an approximation, and t' is chosen to be 20 000. Unless otherwise stated, t' will be taken the same value for all our numerical results.

The corresponding number of higher-layer ACK transmissions is $A_{nonsys} = M$.

3.5. Mean #transmissions with systematic fountain code

When a systematic fountain code is used, the first N transmissions are uncoded packets while subsequent transmissions are coded. Let r_j denote the number of uncoded packets successfully received by client j, out of the N uncoded packets transmitted. If r_j is less than N, an additional $N + \delta - r_j$ coded packets must be received for the client to decode the block. Let Y_j denote the number of transmissions required for client j to successfully receive r_j uncoded packets and $N + \delta - r_j$ coded packets. We notice that Y_j is a Negative Binomial $(N + \delta, 1-p)$ random variable. It follows that the total number of transmissions required for all clients to recover a block of N packets is the random variable

$$T_{sys} = \begin{cases} N, & (1 - p)^{MN} \\ \max_{j \in \{1,\ldots,M\}} Y_j, & 1 - (1 - p)^{MN}. \end{cases} \tag{22}$$

The mean number of transmissions is thus given by

$$
E[T_{\text{sys}}] = (1 - p)^{MN} \cdot N + \left(1 - (1 - p)^{MN}\right)
$$
$$
\times E[\max_{j \in \{1, \dots, M\}} Y_j] \tag{23}
$$

where $E[\max_{j \in \{1, \dots, M\}} Y_j]$ is the mean number of transmissions when using the non-systematic fountain code, given by expression (21).

The number of higher-layer ACK transmissions is $A_{\text{sys}} = M$.

3.6. Goodput

Goodput here is defined as the mean number of application layer information bits correctly received by each client per second. For 802.11 WLANs goodput can be expressed as

$$
G = \frac{N \cdot L}{E[T] \cdot (D_{\text{data}} + D_{cw} + D_{\text{DIFS}}) + E[A] \cdot (D_{\text{ack}} + D_{\text{SIFS}})} \tag{24}
$$

where $E[T]$ is the mean number of transmissions required for all clients to receive a block of N packets, given by (16), (21) and (6); $E[A]$ is the corresponding mean number of higher-layer acknowledgement packets; L is the application layer payload size, which is taken to be 992 bytes in the following numerical analysis; D_{DIFS} and D_{SIFS} are respectively the durations of DIFS and SIFS, which are 34 and 16 μs in 802.11a. DIFS is the 802.11 distributed co-ordination function (DCF) interframe space, and SIFS is the short interframe space [14]. D_{cw} is the mean of the 802.11 contention window countdown duration. When there is a single multicast sender and no collisions between higher-layer acknowledgements and data packets, the random backoff number is uniformly selected between 0 and CW_{min}, where CW_{min} is the minimum 802.11 contention window, and $CW_{\text{min}} = 15$ in 802.11a/g. The value of D_{cw} is hence $\sigma CW_{\text{min}}/2$ where σ is 9 μs, the duration of a PHY idle slot. If desired, a more sophisticated calculation of D_{cw} could also be carried out using modelling approaches such as that in [19]. D_{data} is the transmission duration of a data packet. For a UDP packet over 802.11a/g this is given by

$$
D_{\text{data}} = D_{\text{PHYhdr}} + D_{\text{symbol}}
$$
$$
\cdot \left\lceil \frac{(L + \text{hdr}_{\text{APP}} + \text{hdr}_{\text{UDP}} + \text{hdr}_{\text{IP}} + \text{hdr}_{\text{MAC}}) \times 8 + 22}{\text{DBPS}} \right\rceil \tag{25}
$$

where hdr_{APP}, hdr_{UDP}, hdr_{IP} and hdr_{MAC} are respectively 3, 8, 20 and 24, the length in bytes of the APP, UDP, IP and MAC headers; DBPS is the number of data bits per symbol, which is determined by the transmission mode; D_{symbol} is the OFDM symbol duration, which is 4 μs in 802.11a/g; D_{PHYhdr} is the duration of the PHY header, including the PLCP preamble and PLCP header but excluding the 16 bit SERVICE field, which is 20 μs in 802.11a/g [14]. Similarly, D_{ack} is the transmission duration of a higher-layer ACK/NACK packet. In the following we assume that the acknowledgement packet has a 8 byte payload, and is transmitted at the basic PHY rate of 6 Mbps.

Fig. 2 shows typical behaviours of goodput vs. PER for a fixed 802.11a/g PHY rate of 6 Mbps and $M = 10$ clients.
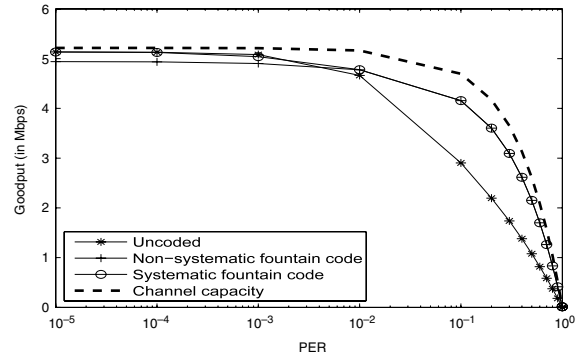

Fig. 2. Goodput vs. packet error rate (PER) for an 802.11a/g PHY rate of 6 Mbps. $M = 10$ client stations, block size $N = 50$ packets, coding overhead $\delta = 2$ packets.

We use a fountain code block size N of 50 packets and assume an overhead δ of 2 packets for both systematic and non-systematic fountain codes (this is a lower bound on the decoding overhead, see [18])—unless otherwise stated we also use these values of N and δ in the remainder of the paper. At low PERs, uncoded and systematic fountain coded traffic have similar goodputs and this is higher than the goodput with a non-systematic fountain code owing to the decoding overhead δ. At higher PERs, both the systematic and non-systematic fountain codes exhibit similar performance, with the achieved goodput substantially higher than the uncoded case.

4. Maximizing goodput: PHY modulation/rate choice

By combining the analysis in Section 3 (which gives goodput as a function of PER) with the channel models in Section 2 (which yield PER as a function of SNR and PHY modulation) we can obtain goodput as a function of SNR and PHY modulation. We can then determine (by numerical search) the choice of PHY modulation/rate that maximizes goodput for a range of channel SNRs and for both AWGN and Rayleigh channels. Fig. 3 shows a plot, obtained using this approach, of maximum goodput and the choice of PHY rates to achieve the maximum goodput versus channel SNR, $M = 10$ client stations and a block of $N = 50$ packets. A first observation is that fountain coding yields uniformly higher goodputs at all SNRs. At some SNRs the increase is significant, while at other SNRs it is minor. We observed this consistently in all of our results, even with Viterbi SDD. We also observed that at those SNRs where the difference in goodput is significant, the PER with fountain coding is around 10%–30% and higher than that without coding. At other SNRs however, the PER value is very low (close to 0), and therefore there is only a small difference in goodput.

Our primary interest here, however, is in the impact of fountain coding on the choice of PHY modulation/rate. It can be seen from Fig. 3 that use of fountain coding allows the PHY rate to be stepped up slightly sooner as SNR rises. However, the SNR interval over which fountain coding yields an advantage is rather narrow (less than 1 dB wide), and hence the performance loss of operating fountain coding with uncoded PHY rate selection is rather mi-
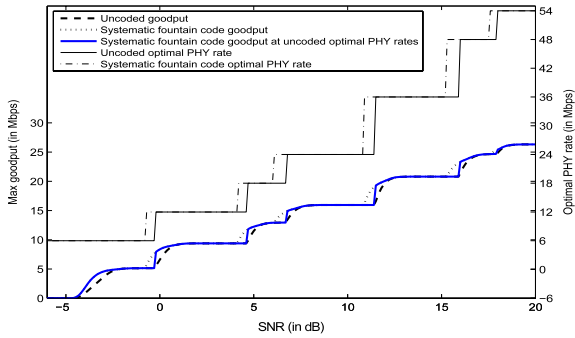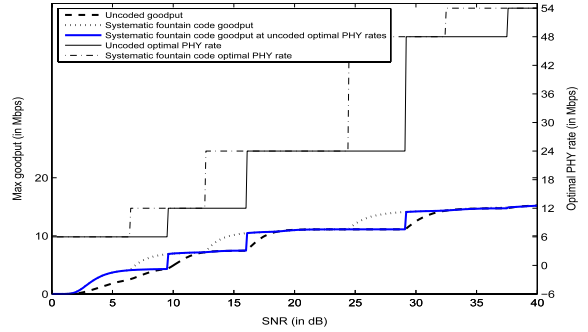
(a) AWGN, HDD.
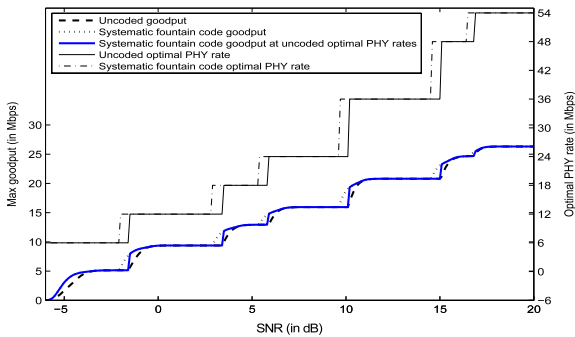
(b) AWGN, SDD.
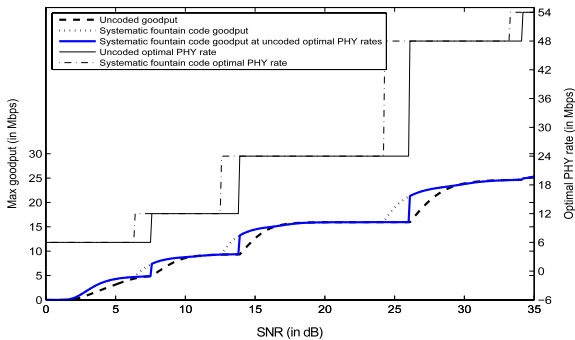
(c) Rayleigh channel, HDD.

Fig. 3. PHY rates maximizing goodput vs. SNR, $M = 10$ client stations, block size $N = 50$ packets. The three goodput lines at the bottom correspond to the left y-axis, and the two PHY rate lines at the top correspond to the right y-axis.

Fig. 4. PHY rates maximizing goodput vs. SNR, Rayleigh channel, HDD, $M = 100$ client stations, block size $N = 50$ packets. The three goodput lines at the bottom correspond to the left y-axis, and the two PHY rate lines at the top correspond to the right y-axis.
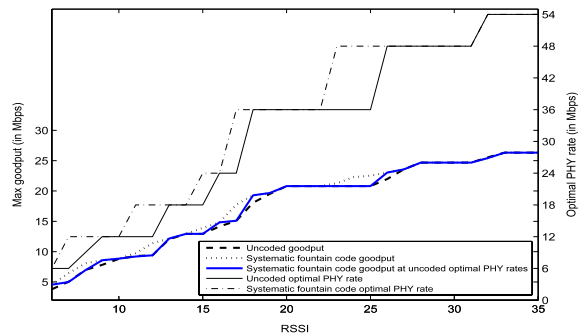
**Fig. 5.** PHY rates maximizing goodput vs. RSSI, with $M = 10$ clients, block size $N = 50$, outdoor measurement. The three goodput lines at the bottom correspond to the left y-axis, and the two PHY rate lines at the top correspond to the right y-axis.

nor. This is perhaps unsurprising with only 10 client stations, since there may be insufficient diversity. Fig. 4 shows the corresponding results for $M = 100$ client stations and a Rayleigh channel. This is, of course, a very large number of clients in the context of 802.11 WLANs. It can be seen that even with this larger number of clients and over a Rayleigh channel, the difference in goodput is still quite small, although the difference in the PHY rate choice between fountain coded and uncoded traffic is more pronounced with the SNR range where fountain coding allows a higher PHY rate used now around 4 dB wide. However, this is quite an extreme regime and for smaller numbers of client stations

(up to around 50 clients) the SNR interval continues to be only about 1 dB even over a Rayleigh channel.

We also evaluate the maximum goodput and the PHY rate choice using a set of experimental PER values for the set of 802.11a/g OFDM-based PHY rates. The measurements are taken in an outdoor open space. An Asus Eee PC 4G Surf equipped with an Atheros AR5BXB63 802.11b/g chipsets (AR2425, MAC 14.2, RF5424, PHY 7.0) was used as the transmitter, running FreeBSD 8.0 with the RELEASE kernel and using the standard FreeBSD ATH driver. A Fujitsu E series Lifebook equipped with a Netgear dual band 802.11a/b/g wireless PC card WAG511 using an Atheros AR5212 chipset was used as a receiver, running Ubuntu 11.04 with the RELEASE kernel and using a modified Linux Madwifi driver. The distance between them is altered to achieve different values of PER. Fig. 5 shows the maximum goodput and the choice of PHY rates versus RSSI (Received Signal Strength Indicator) with $M = 10$ clients and block size $N = 50$. It can be seen that the range of RSSI where the choices of PHY rates for coded and uncoded traffic are different becomes much wider, but again the difference in goodput is so small that the goodput loss of operating fountain coding at uncoded PHY rates is rather minor.

In summary, we find that fountain coding uniformly increases network goodput which is as expected. The choice of PHY rates that maximizes goodput with fountain

(a) Rayleigh, HDD.

(b) Outdoor measurement.
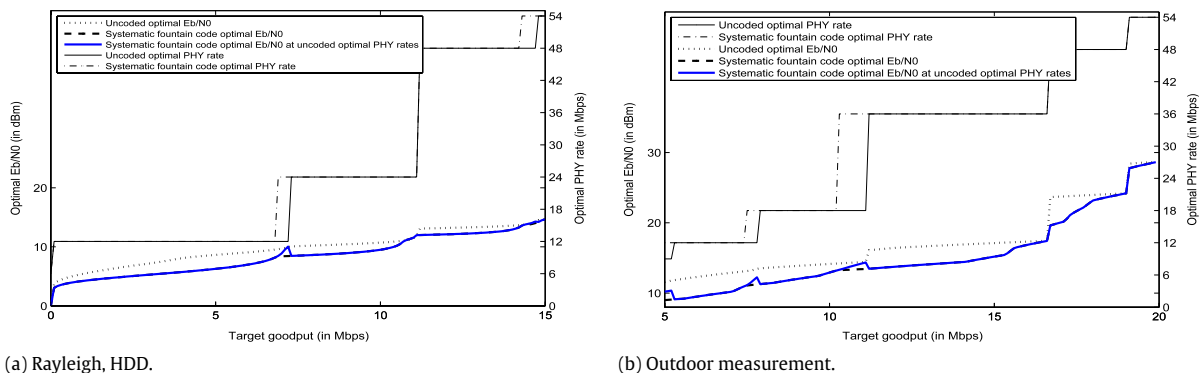
Fig. 6. PHY rates minimizing energy for a given target goodput, with $M = 100$ clients, block size $N = 50$. The three E_b/N_0 lines at the bottom correspond to the left y-axis, and the two PHY rate lines at the top correspond to the right y-axis.

coding generally differs from that for uncoded traffic, which is again as expected. Interestingly, however, the difference in goodput is extremely small over a wide range of network conditions, and hence the goodput loss associated with operating fountain coding at uncoded PHY rates is minor. Therefore the choice of PHY rates for uncoded traffic is close to that for fountain-coded traffic.

5. Minimizing energy: joint PHY power and modulation/rate choices

We now extend consideration to the choices of PHY power and rate to minimize energy expenditure for a given goodput. This is a key issue in the context of battery powered mobile devices, and in sensor networks.

Let E_b/N_0 denote the SNR per bit. We are interested in the choice of PHY rates which minimizes E_b/N_0 while achieving a target goodput. For a given choice of PHY modulation/rate, the analysis in Section 3 relates PER to goodput. We can invert this relationship to obtain the maximum permissible PER for a given PHY modulation/rate and target goodput, and then use the channel models described in Section 2 relating PER to SNR to determine the corresponding minimum SNR per bit E_b/N_0. Given a target goodput value, the minimum E_b/N_0 and the corresponding PHY rate at which the minimum E_b/N_0 is achieved can be obtained by numerical search amongst the PHY rate set supported by 802.11 a/g.

This yields Fig. 6(a). It can be seen that fountain coding uniformly reduces the energy required to achieve a target goodput, and the choice of PHY rates is almost identical for coded and uncoded traffic. Fig. 6(b) shows the corresponding results based on the experimental PER performance. Again, it can be seen that the choices of PHY rates for both fountain-coded and uncoded traffic are almost the same. Therefore, we find that fountain coding uniformly decreases energy requirements, and the PHY rates minimizing energy usage for coded/uncoded traffic are almost the same over a wide range of network conditions, so that the energy loss of operating fountain coding using uncoded PHY rate choices is extremely minor.

6. Discussion

The joint performance of a fountain code concatenated with a PHY layer code has previously been considered in [8,9] in the context of 3G cellular networks, and more recently in [5–7] in a general wireless context. In both cases it is concluded that when fountain codes are used at higher-layers then the overall system performance is improved if the PHY layer modulation/rate is selected to operate at much higher packet error rate (a 20%–30% loss rate in [9]) than would normally be used for uncoded traffic. That is, a cross-layer approach, whereby the PHY layer rate control exploits knowledge of the use of fountain coding at higher layers, can yield significant performance gains. We find in the context of the PHY modulation/rate choices available in 802.11a/g WLANs that a higher PHY rate associated with a higher PER (10%–30%) is selected at some conditions, i.e. at some SNR values, as can be seen in Figs. 3 and 4, or at some target goodput values, as can be seen in Fig. 6. However, we also find that it holds only at a small number of SNRs and target goodputs, and thus offers limited benefits. The main reason for this is that the PHY rate-set choice in 802.11a/g is not rich enough that one can always operate close to a 10%–30% packet error rate. Therefore over a wide range of network conditions the choice of PHY rate/modulation is similar for both fountain coded and uncoded traffic. Moreover this is true not only when the objective is to maximize goodput but also when we seek to minimize energy expenditure.

The paper [9] considers a very similar setup to that considered here, but focuses specifically on Raptor codes (a particular systematic fountain code) over a UMTS MBMS service that uses QPSK modulation and turbo-coded FEC concatenated with a physical layer FEC. The performance objective is to minimize energy rather than maximize goodput. The authors in [9] observe that an optimum level of PHY level FEC exists, however this is intimately related to the FEC options supported by the PHY layer. Additionally, they only consider scenarios with 500 client stations. This is far higher than the number of clients likely to be present in an 802.11 WAN. We have already noted in our own analysis that the impact of coding becomes more pronounced as the number of client stations increases. This
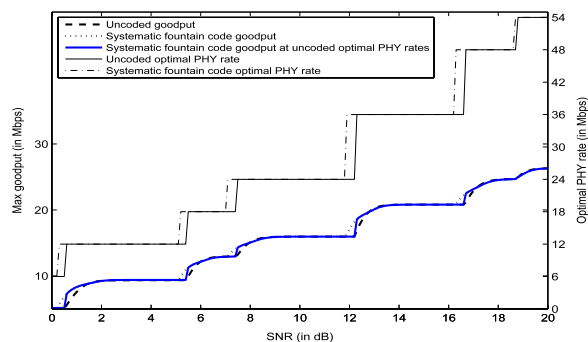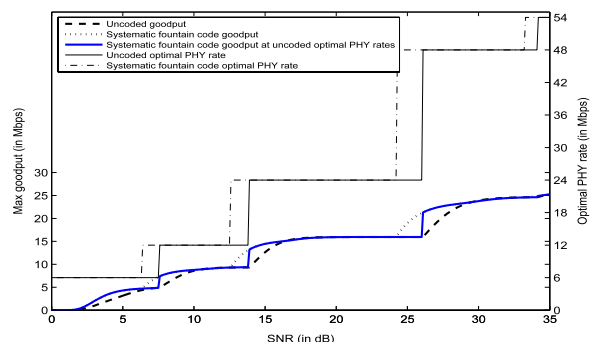
Fig. 7. PHY rates maximizing goodput vs. SNR of class 1 stations, with $M_1 = 5$ class 1 stations and $M_2 = 5$ class 2 stations, block size $N = 50$, AWGN channel. The three goodput lines at the bottom correspond to the left *y*-axis, and the two PHY rate lines at the top correspond to the right *y*-axis.

Fig. 8. PHY rates maximizing goodput vs. SNR, $M = 10$ client stations, block size $N = 50$ packets, Nakagami-6 fading channel. The three goodput lines at the bottom correspond to the left *y*-axis, and the two PHY rate lines at the top correspond to the right *y*-axis.

suggests that it might be interesting to repeat the analysis in [9] for smaller numbers of client stations to evaluate how this changes the cross-layer trade-off.

We have assumed that all client stations have the same PER at every PHY rate. This is not the case in a wireless network where clients are in different locations owing to different path loss, shadow fading and interference effects. However, this assumption has little effect on our conclusions. To illustrate this, consider a heterogeneous multicast network with two classes of client stations, with M_1 stations in class 1 and M_2 in class 2. All stations in the same class have the same SNR. Class 1 is subjected to noisy reception (PER $\geq$ 0); while class 2 experiences reliable reception at any of the available PHY rates (PER $=$ 0). Fig. 7 shows the maximum goodput and the choice of PHY rates versus SNR of class 1 for a network with $M_1 = 5$ class 1 stations and $M_2 = 5$ class 2 stations and an AWGN channel. It can be seen that even though the PER is different among client stations, the conclusion remains the same. The reason is that as our objective is for each individual client station to receive all N packets in each block, the number of transmissions is determined by the station(s) with the lowest PER. When the worst client(s) receive(s) the whole block, all other clients with higher SNRs have already received everything. The homogeneous setting is therefore the worst-case setting. Even though we may quantitatively get a different maximum goodput value, qualitatively the behaviour will be the same as in the heterogeneous case. If the objective changes, e.g. to ensure only 95% clients decode successfully, we just need to reduce the number of decoding overheads, hence the conclusion will still hold.

For the setting of this paper we have assumed a Rayleigh flat fast-fading channel. It is known that multipath-fading will lead to a diversity gain of the order of the number of distinguishable multipath components. Flat fading is the worst-case setting for our problem, in that it will yield the greatest difference between fountain coded and uncoded performances. For the purpose of illustration results for HDD Nakagami-6 fast-fading model are shown in Fig. 8. Similar behaviours to the AWGN and Rayleigh channels are observed. This indicates that our conclusion is insensitive to the channel model considered.

7. Conclusions

In this paper we consider the joint performance of fountain codes and 802.11a/g PHY modulation/coding. We consider optimality both in terms of maximizing goodput and minimizing energy, and results are presented for both theoretical and experimental channel models. To our knowledge the present paper is the first detailed study on the selection of PHY modulation/coding to use with and without fountain codes in standard 802.11a/g WLANs. In contrast to studies in cellular networks, we find that in 802.11a/g WLANs the cross-layer approach of a higher-layer fountain coding with a PHY layer modulation and FEC coding can yield very limited gains, and the PHY modulation/rate that maximizes uncoded goodput or minimizes uncoded energy is close to that for fountain-coded multicast traffic over a wide range of network conditions, principally because of the limited PHY rate-set choice available in 802.11a/g WLANs. This is potentially an important observation as it indicates that in 802.11a/g WLANs cross-layer design for multicast rate-control would bring few benefits, and PHY layer modulation/rate-control can be carried out without regard to the use of fountain coding at higher layers.

References

[1] M. Luby, LT codes, in: Proc. of 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002, pp. 271–280.
[2] A. Shokrollahi, Raptor codes, IEEE Transactions on Information Theory 52 (2006) 2551–2567.
[3] J.W. Byers, M. Luby, M. Mitzenmacher, A digital fountain approach to asynchronous reliable multicast, IEEE Journal on Selected Areas in Communications 20 (2002) 1528–1540.
[4] V. Badarla, V.G. Subramanian, D.J. Leith, Low-delay dynamic routing using fountain codes, IEEE Communications Letters 13 (2009) 552–554.
[5] C.R. Berger, S. Zhou, Y. Wen, K. Pattipati, P. Willett, Optimizing joint erasure- and error-correction coding for wireless packet transmissions, IEEE Transactions on Wireless Communications 7 (2008) 4586–4595.
[6] M. Samokhina, K. Moklyuk, S. Choi, J. Heo, Raptor code-based video multicast over IEEE 802.11 WLAN, in: Proc. of 5th IEEE VTS Asia Pacific Wireless Communications Symposium, 2008.
[7] T.A. Courtade, R.D. Wesel, A cross-layer perspective on rateless coding for wireless channels, in: Proc. of IEEE International Conference on Communications, 2009, pp. 1–6.

[8] M. Luby, M. Watson, T. Gasiba, T. Stockhammer, W. Xu, Raptor codes for reliable download delivery in wireless broadcast systems, in: Proc. of IEEE Consumer Communications and Networking Conference, 2006, pp. 192–197.

[9] M. Luby, M. Watson, T. Gasiba, T. Stockhammer, Mobile data broadcasting over MBMS: tradeoffs in forward error correction, in: Proc. of 5th International Conference on Mobile and Ubiquitous Multimedia, 2006.

[10] M.B. Pursley, D.J. Taipale, Error probabilities for spread-spectrum packet radio with convolutional codes and Viterbi decoding, IEEE Transactions on Communications 35 (1987) 1–12.

[11] A. Goldsmith, Wireless Communications, Cambridge University Press, 2005.

[12] A.J. Viterbi, Convolutional codes and their performance in communication systems, IEEE Transactions on Communication Technology 19 (1971) 751–772.

[13] R.C. Manso, Performance analysis of M-QAM with Viterbi soft-decision decoding, Master's Thesis, Naval Postgraduate School Monterey CA, 2003.

[14] IEEE 802.11: wireless LAN medium access control (MAC) and physical layer (PHY) specifications, IEEE, IEEE std 802.11-2007 edition, 2007.

[15] M.K. Simon, M.S. Alouini, Digital Communication Over Fading Channels, Wiley-IEEE Press, 2005.

[16] J. Lassing, T. Ottosson, E. Strom, On the union bound applied to convolutional codes, in: Proc. of 54th IEEE Vehicular Technology Conference, 2001, pp. 2429–2433.

[17] D.J.C. MacKay, Fountain codes, in: IEE Proceedings Communications, Vol. 152, 2005, pp. 1062–1068.

[18] V.G. Subramanian, D.J. Leith, On a class of optimal rateless codes, in: Proc. of 46th Annual Allerton Conference on Communication, Control and Computing, 2008, pp. 418–425.

[19] D. Malone, K. Duffy, D. Leith, Modeling the 802.11 distributed coordination function in nonsaturated heterogeneous conditions, IEEE/ACM Transactions on Networking 15 (2007) 159–172.

Xiaomin Chen received the M.Sc. degree from Beijing University of Posts and Telecommunications, China, in 2008, and the Ph.D. degree from the Hamilton Institute, NUI Maynooth, Ireland in 2012. She is currently a Post-doc Research Fellow at the Hamilton Institute, NUIM. Her research interests lie in coding, protocol design and resource allocation in wireless networks.

Vijay Subramanian received the B.Tech. degree from IIT Madras, in 1993, the M.S. degree from the IISc. Bangalore, in 1995, and the Ph.D. degree from the UIUC, Urbana, in 1999. From 1999 to 2006, he was with the Networks Business, Motorola. From 2006 to 2010 he was with the Hamilton Institute, NUIM, Ireland as a Research Fellow. From November 2010 to October 2011, he was a Senior Research Associate in the EECS Department at Northwestern University. He is at present a Research Assistant Faculty in the EECS Department at Northwestern University. His research interests include stochastic modeling, communications, information theory and applied mathematics.

Doug Leith graduated from the University of Glasgow in 1986 and was awarded his Ph.D., also from the University of Glasgow, in 1989. In 2001, Prof. Leith moved to the National University of Ireland, Maynooth to assume the position of SFI Principal Investigator and to establish the Hamilton Institute (www.hamilton.ie) of which he is Director. His current research interests include the analysis and design of network congestion control and resource allocation in wireless networks.