Universidad
Carlos III de Madrid

This is a postprint version of the following published document:

Durisi, Giuseppe; Koch, Tobias; Popovski, Petar. Towards Massive, Ultrareliable, and Low-Latency Wireless Communication With Short Packets. In: *Proceedings of the IEEE*, 104(9), pp. 1711-1726, Sept. 2016

DOI: <https://doi.org/10.1109/JPROC.2016.2537298>

© 2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Towards Massive, Ultra-Reliable, and Low-Latency Wireless Communication with Short Packets

Giuseppe Durisi, *Senior Member, IEEE*, Tobias Koch, *Member, IEEE*, Petar Popovski, *Fellow, IEEE*

Abstract—Most of the recent advances in the design of high-speed wireless systems are based on information-theoretic principles that demonstrate how to efficiently transmit long data packets. However, the upcoming wireless systems, notably the 5G system, will need to support novel traffic types that use short packets. For example, short packets represent the most common form of traffic generated by sensors and other devices involved in Machine-to-Machine (M2M) communications. Furthermore, there are emerging applications in which small packets are expected to carry critical information that should be received with low latency and ultra-high reliability.

Current wireless systems are not designed to support short-packet transmissions. For example, the design of current systems relies on the assumption that the metadata (control information) is of negligible size compared to the actual information payload. Hence, transmitting metadata using heuristic methods does not affect the overall system performance. However, when the packets are short, metadata may be of the same size as the payload, and the conventional methods to transmit it may be highly suboptimal.

In this article, we review recent advances in information theory, which provide the theoretical principles that govern the transmission of short packets. We then apply these principles to three exemplary scenarios (the two-way channel, the downlink broadcast channel, and the uplink random access channel), thereby illustrating how the transmission of control information can be optimized when the packets are short. The insights brought by these examples suggest that new principles are needed for the design of wireless protocols supporting short packets. These principles will have a direct impact on the system design.

I. INTRODUCTION

The vision of the Internet of Things promises to bring wireless connectivity to “...anything that may benefit from being connected...” [1], ranging from tiny static sensors to vehicles and drones. A successful implementation of this vision calls for a wireless communication system that is able to support a much larger number of connected devices, and that is able to fulfill

much more stringent requirements on latency and reliability than what current standards can guarantee. Among the various current research and standardization activities, the one aimed at the design of fifth generation (5G) wireless systems stands out as the largest globally orchestrated effort towards addressing these challenges.

So far, each new generation of cellular systems has been mainly designed with the objective to provide a substantial gain in data rate over the previous generation. 5G will depart from this scheme: its focus will not only be on enhanced broadband services and, hence, higher data rates. This is because the vast majority of wireless connections in 5G will most likely be originated by autonomous machines and devices rather than by the human-operated mobile terminals for which traditional broadband services are intended. 5G will address the specific needs of autonomous machines and devices by providing two novel wireless modes: ultra-reliable communication (URC) and massive machine-to-machine communications (MM2M) [2]–[4].

URC refers to communication services where data packets are exchanged at moderately low throughput (e.g., 50 Mbit/s) but with stringent requirements in terms of reliability (e.g., 99.999%) and latency (e.g., 4 ms). Example of URC include reliable cloud connectivity, critical connections for industrial automation, and reliable wireless coordination among vehicles [4]–[6].

With MM2M one refers to the scenario where a massive number of devices (e.g., 10 000) needs to be supported within a given area. This is relevant for large-scale distributed cyber-physical systems (e.g., smart grid) or industrial control. Also in this case, the data packets are short (and often contain correlated measurements) and reliability must be high to cope with critical events.

The central challenge with these two new wireless modes is the capability to support *short packet* transmission. Indeed, short packets are the typical form of traffic generated by sensors and exchanged in machine-type communications. This requires a fundamentally different design approach than the one used in current high-data-rate systems, such as 4G LTE and WiFi.

It is appropriate at this point to formally define what is meant by short/long packets. The transmission of a packet is a process in which the information payload (data bits) is mapped into a continuous-time signal, which is then transmitted over the wireless channel. A continuous-time signal with approximate duration T and approximate bandwidth B can be described by $n \approx BT$ complex parameters. It is then natural to refer to n as the packet length, i.e., the number of degrees of freedom (channel uses) that are required for the transmission of the information payload.

The work of G. Durisi has been in part supported by the Swedish Research Council under Grant 2012-4571. The work of T. Koch has been supported in part by the European Community's Seventh Framework Programme FP7/2007-2013 under Grant 333680, in part by the Ministerio de Economía y Competitividad of Spain under Grants TEC2013-41718-R, RYC-2014-16332, and TEC2015-69648-REDC, and in part by the Comunidad de Madrid under Grant S2013/ICE-2845. The work of P. Popovski has been in part supported by the European Research Council (ERC Consolidator Grant Nr. 648382 WILLOW) within the Horizon 2020 Program. The simulations were performed in part on resources at Chalmers Centre for Computational Science and Engineering (C3SE) provided by the Swedish National Infrastructure for Computing (SNIC).

G. Durisi is with the Department of Signals and Systems, Chalmers University of Technology, 41296, Gothenburg, Sweden (e-mail: durisi@chalmers.se).

T. Koch is with the Signal Theory and Communications Department, Universidad Carlos III de Madrid, 28911, Leganés, Spain and also with the Gregorio Marañón Health Research Institute, Madrid, Spain (e-mail: koch@tsc.uc3m.es).

P. Popovski is with the Department of Electronic Systems, Aalborg University, 9220 Aalborg, Denmark (e-mail: petarp@es.aau.dk).

The authors are listed in alphabetical order.

Fig. 1. An example of a packet structure with data and metadata. In many wireless systems, metadata consists of preamble (PA) and header (H). (a) Long data packet used in current wireless systems; (b) Short data packets needed to support novel 5G applications, such as URC and MM2M.

A *channel code* defines a map between the information payload and the signal transmitted over the n channel uses. The task of a wireless receiver is to recover the information payload from a distorted and noisy version of the transmitted signal. A fundamental result in information theory [7] tells us that when n is large (long packets), there exist channel codes for which the information payload can be reconstructed with high probability (in a sense we shall make precise in Section II). Intuitively, when n is large both the thermal noise and the distortions introduced by the propagation channel are averaged out due to the law of large numbers. However, when n is small (short packets) such averaging cannot occur.

Another defining element of long packets, besides the large number of channel uses, is the fact that the payload contained in a packet is much larger than the control information (metadata) associated with the packet. As a consequence, a highly suboptimal encoding of the metadata does not deteriorate the efficiency of the overall transmission, see Fig. 1(a). On the contrary, when the packets are short, the metadata is no longer negligible in size compared to the payload, see Fig. 1(b).

To summarize, in short-packet communications (i) classic information-theoretic results are not applicable because the law of large numbers cannot be put to work; (ii) the size of the metadata is comparable to the size of the payload and inefficient encoding of metadata significantly affects the overall efficiency of the transmission.

During the last few years, significant progress has been made within the information theory community to address the problem of transmitting short packets. Particularly for point-to-point scenarios, information theorists have gained some understanding of the theoretical principles governing short-packet transmission and possess metrics that allow them to assess their performance. In contrast, so far information theorists have mostly viewed the design of metadata as something outside their competence area. Consequently, the transmission of metadata has been largely left to heuristic approaches. In fact, practically all current protocols are based on a tacit assumption that the control information is perfectly reliable. A classic example is the proverbial “one-bit acknowledgement”, which is always assumed to be perfectly received.

In this article, we present a comprehensive review of the theoretical principles that govern the transmission of short packets

Fig. 2. Block diagram that illustrates how a packet is created.

and present metrics that allow us to assess their performance. We then highlight the challenges that need to be addressed to optimally design URC and MM2M applications by means of three examples that illustrate how the tradeoffs brought by short-packet transmission affect protocol design.

The paper is organized as follows. In Section II, we describe the structure of a packet and review two classic information-theoretic metrics that are relevant for long packets: the *ergodic capacity* and the *outage capacity*. In Section III, we introduce a performance metric, the *maximum coding rate* at finite packet length and finite packet error probability, that is more relevant for the case of short packets. By focusing on the case of additive white Gaussian noise (AWGN) channels and on the case of fading channels, we explain how to evaluate this quantity and discuss the engineering insights brought by it. In Section IV, we illustrate through three examples how to use the maximum coding rate performance metric to optimize the protocol design and the transmission of metadata in short-packet communications. Concluding remarks are offered in Section V.

II. ANATOMY OF A PACKET

Modern wireless systems transmit data in packets. Each transmitted packet over the air carries not only the information bits intended for the receiver but also additional bits that are needed for the correct functioning of the wireless protocols. Such bits, which will be referred throughout as control information or metadata—in contrast to the actual *data* to be transmitted—include packet initiation and termination, logical addresses, synchronization and security information, etc...

As illustrated in Fig. 2, a packet consists of k payload bits, which are made up of k_i information bits (information payload) and k_o additional bits, containing metadata from the media-access-control (MAC) layer and higher layers. The payload bits are typically encoded into a block of n_e data symbols (complex numbers) to increase reliability in packet transmission. Finally, n_o additional symbols are added to enable packet detection, efficient synchronization (in time and frequency), or estimation of channel state information (CSI), which is needed by the receiver to compensate for the distortion of the transmitted signal introduced by the wireless channel. The total packet length n is thus equal to $n_e + n_o$. With a slight abuse of notation, we shall refer to the additional k_o bits and n_o symbols as metadata.

The ratio $R = k_i/n$, i.e., the number of information bits per complex symbol (or, equivalently, the number of transmitted payload bits per second per unit bandwidth) represents the net transmission rate and is a measure of the spectral efficiency of a communication system. In some wireless standards (such as LTE) specific physical/logical channels are reserved to carry exclusively metadata (control channels). This lowers further the spectral efficiency.

In most current wireless systems, we have that $k_i \gg k_o$ and that $n_e \gg n_o$, so the net transmission rate R is roughly

k/n_e . Consequently, the performance of such systems essentially depends on the efficiency of the channel code. Furthermore, k_i (and hence also n_e) is typically large. It follows that information-theoretic metrics such as capacity [7] and outage capacity (also known as *capacity-versus outage*) [8] are accurate, in spite of being defined for asymptotically large packet sizes. In summary, encoding the data payload using a good channel code allows for reliable transmission at rates close to the capacity of the underlying channel.

In order to facilitate the review of the relevant information-theoretic metrics, we shall need a reference communication channel. A communication channel—the central part of a communication model—describes the relation between the input signal and the output signal over the available n channel uses. As mentioned in Section I, each channel use corresponds to the transmission of a complex symbol. Throughout most of the paper, we shall focus on the following channel model (and its multiple-antenna extension):

$$Y_k = H_k X_k + W_k, \quad k \in \mathbb{N}. \quad (1)$$

Here, X_k denotes the complex symbol transmitted over the k th channel use, Y_k is the corresponding channel output, H_k is the channel coefficient that represents fading and other propagation phenomena and W_k is the additive Gaussian noise, which we shall assume to be drawn from a stationary memoryless process.

If H_k is taken equal to a deterministic constant c independent of k and known to transmitter and receiver, i.e., $H_k = c$ for all k , then (1) describes an AWGN channel. The AWGN channel is an example of an ergodic channel, that is, it exhibits an ergodic behavior over the duration of each packet (recall that the noise process $\{W_k\}$ is assumed stationary and memoryless). For such ergodic channels, the relevant performance metric is the capacity C , defined as the largest rate k/n_e for which the packet error probability can be made arbitrarily small by choosing n_e sufficiently large. We shall treat the AWGN channel in more detail in Section III-B. Another example of an ergodic channel is the memoryless block fading channel, see Section III-C. In this model, H_k , which can be thought of as a multiplicative noise, is assumed not known a priori by the transmitter/receiver and to vary according to a block-memoryless process.

A completely different situation is the one in which the fading coefficient H_k is random but does not depend on k , i.e., $H_k = H$. Hence, the fading coefficient stays constant over the packet duration [9, p. 2631], [10, Sec. 5.4.1]. For this nonergodic channel, if H can take arbitrarily small values, the error probability cannot be made small by choosing n_e large. This is the case for most fading distributions, e.g., Rayleigh, Rician, and Nakagami. Indeed, when $|H|$ is small (deep fade), then the entire packet is lost, independently of its length. In such a nonergodic case, a relevant performance metric is the outage capacity C_e (also known as *capacity-versus-outage* or ϵ -capacity)—defined as the largest rate k/n_e for which a packet error probability less than a fixed $\epsilon > 0$ can be achieved by choosing n_e sufficiently large.

We note that both capacity and outage capacity require that the codeword length n_e (i.e., the packet size) and, hence, also the size of the data payload k be large. When the packets are short, the situation changes drastically. On the one hand, new information-theoretic performance metrics other than capacity

Fig. 3. Information-theoretic description of a communication system.

or outage capacity are needed to capture the tension between reliability and throughput, as well as the cost incurred in exploiting time-frequency and spatial resources (PHY overhead). On the other hand, when the packets are short, the MAC overhead is significant and needs to be designed optimally, perhaps together with the data. We shall address the former issue in Section III and the latter issue in Section IV.

III. RETHINKING PHY PERFORMANCE METRICS

A. Backing Off from the Infinite Blocklength Asymptotics

In this section, we discuss information-theoretic performance metrics for short-packet wireless communications. We account for the metadata symbols required for the estimation of CSI, but ignore other issues such as packet detection or synchronization. As is common in information theory, we view the blocks *channel encoder* and *PHY overhead* in Fig. 2 as one encoder block and consider the transmission of metadata symbols for channel estimation, such as *pilot symbols*, as a possible encoding strategy. We note that the use of pilot symbols to estimate the channel is a widely adopted heuristic strategy that may be strictly suboptimal in some cases.

Mathematically, the encoder is modeled as a function f_n that maps the k information bits $B_1, \dots, B_k$ to the sequence of symbols $X_1, \dots, X_n$ to be transmitted over the channel; see Fig. 3. We shall refer to the number of transmitted symbols n as the *packet length* or *blocklength* and to the sequence $X_1, \dots, X_n$ as a *codeword*. It is common to impose a power constraint ρ on the transmitted symbols to account for restrictions on the transmit power, e.g., due to the devices' limited battery life or regulatory constraints. An often-used power constraint is the *average power constraint*, under which the transmitted symbols must satisfy

$$\frac{1}{n} \sum_{k=1}^n |X_k|^2 \leq \rho. \quad (2)$$

The task of the decoder is to guess the information bits $B_1, \dots, B_k$ from the n channel outputs $Y_1, \dots, Y_n$. The decoding procedure is modeled as a function g_n that maps the channel outputs $Y_1, \dots, Y_n$ to the estimates $\hat{B}_1, \dots, \hat{B}_k$.

Let P_e denote the *packet error probability*, i.e., the probability that the decoder makes a wrong guess about the information bits $B_1, \dots, B_k$. Note that P_e does not only depend on the decoder g_n , but also on the encoder f_n .

The rate R of a communication system is defined as the fraction k/n of information bits to the number of transmitted symbols. Ideally, we would like to design communication systems for which R is as large as possible while, at the same time, the packet error probability P_e is as small as possible. We denote by $R^*(n, \epsilon)$ the *maximum coding rate* at finite packet length n and finite packet error probability ϵ , i.e., the largest rate k/n for which there exists an encoder/decoder pair (f_n, g_n) of packet length n whose packet error probability P_e does not exceed ϵ .

Traditional information-theoretic metrics, such as *capacity* [7] and *outage capacity* [8], can be directly obtained from $R^*(n, \epsilon)$ by taking appropriate limits. Specifically, the outage capacity C_ϵ is defined as the largest rate k/n such that, for every sufficiently large packet length n , there exists an encoder/decoder pair (f_n, g_n) whose packet error probability does not exceed ϵ . Thus, in contrast to $R^*(n, \epsilon)$, the definition of C_ϵ does not involve encoder/decoder pairs of a given fixed packet length n ; instead, we consider encoder/decoder pairs whose packet lengths are large enough for the error probability to fall below ϵ . It follows that C_ϵ can be obtained from $R^*(n, \epsilon)$ via

$$C_\epsilon = \lim_{n \rightarrow \infty} R^*(n, \epsilon). \quad (3)$$

The capacity C (in wireless communications also referred to as *ergodic capacity*) is defined as the largest rate k/n such that there exists an encoder/decoder pair (f_n, g_n) whose packet error probability can be made arbitrarily small by choosing the packet length sufficiently large. Thus, in contrast to the definition of the outage capacity that demands a packet error probability smaller than some ϵ , the definition of capacity is stronger in that it demands an arbitrarily small packet error probability. It follows that C can be obtained from C_ϵ by letting ϵ tend to 0:

$$C = \lim_{\epsilon \rightarrow 0} C_\epsilon = \lim_{\epsilon \rightarrow 0} \lim_{n \rightarrow \infty} R^*(n, \epsilon). \quad (4)$$

Intuitively, the capacity characterizes the largest transmission rate at which reliable communication is feasible when there are no restrictions on the packet length. Likewise, the outage capacity characterizes the largest transmission rate at which communication with packet error probability not exceeding ϵ is feasible, again provided that there are no restrictions on the packet length. It follows that both quantities are reasonable performance metrics for current wireless systems, where the packet size is typically large. However, assessing the performance of short packet communications requires a more refined analysis of $R^*(n, \epsilon)$. Unfortunately, the exact value of $R^*(n, \epsilon)$ is unknown even for channel models that are much simpler to analyze than the one encountered in wireless communications. Indeed, determining $R^*(n, \epsilon)$ is in general an NP-hard problem [11], and its complexity is conjectured to be doubly exponential in the packet length n .

Fortunately, during the last few years, significant progress has been made within the information theory community to address the problem of quantifying $R^*(n, \epsilon)$ and, hence, solve the long-standing problem of accounting for latency constraints in a satisfactory way. Building upon Dobrushin's and Strassen's previous asymptotic results, Polyanskiy, Poor, and Verdú [12] recently provided a unified approach to obtain tight bounds on $R^*(n, \epsilon)$. They showed that for various channels with positive capacity C , the maximal coding rate $R^*(n, \epsilon)$ can be expressed as

$$R^*(n, \epsilon) = C - \sqrt{\frac{V}{n}} Q^{-1}(\epsilon) + \mathcal{O}\left(\frac{\log n}{n}\right) \quad (5)$$

where $\mathcal{O}(\log n/n)$ comprises remainder terms of order $\log n/n$. Here, $Q^{-1}(\cdot)$ denotes the inverse of the Gaussian Q function and V is the so-called *channel dispersion* [12, Def. 1]. The approximation (5) implies that to sustain the desired error probability

Fig. 4. Upper bounds, lower bounds, and normal approximation on $R^*(n, \epsilon)$ for the AWGN channel with SNR $\rho = 0$ dB. The packet error probability ϵ is 10^{-3} . The upper bound is obtained using the metaconverse theorem [12, Th. 41]; the lower bound is the Shannon cone-packing bound [13], [12, Eq. (41)]. The normal approximation is indistinguishable from the lower bound.

ϵ for a given packet size n , one incurs a penalty on the rate (compared to the channel capacity) that is proportional to $1/\sqrt{n}$.

We next provide an interpretation for (5). The classic approach of approximating $R^*(n, \epsilon) \approx C$ for large packet sizes and small packet error rates according to (4) allows one to model a communication channel as a “bit pipe” that delivers reliably C bits per channel use. This holds under the assumption that good channel codes are used. The expansion provided in (5) suggests the following alternative model, which is more accurate when the packets are short: A communication channel can be thought of as a bit pipe of randomly varying size. Specifically, the size of the bit pipe behaves as a Gaussian random variable with mean C and variance V/n . Hence, V is a measure of the channel dispersion. In this interpretation, the packet error probability ϵ is the probability that $R^*(n, \epsilon)$ is larger than the size of the bit pipe.

B. AWGN Channel

Arguably, one of the best-understood channel models in the information theory literature is the average-power constrained AWGN channel. Its canonical form can be obtained from (1) by setting $H_k = 1$, which yields

$$Y_k = X_k + W_k. \quad (6)$$

Here, the inputs $\{X_k\}$ satisfy the average-power constraint (2). When the additive noise has unit variance, the power constraint ρ becomes equal to the signal-to-noise ratio (SNR).

For the AWGN channel, the capacity and the channel dispersion are given by [12, Th. 54]¹

$$C(\rho) = \log(1 + \rho) \quad (7)$$

$$V(\rho) = \rho \frac{(2 + \rho)}{(1 + \rho)^2} (\log e)^2. \quad (8)$$

¹The capacity of the *real-valued* AWGN channel has been obtained by Shannon [7]. The channel dispersion of the *real-valued* AWGN channel has been reported in [12, Eq. (293)]. One obtains (7) and (8) by noting that the transmission of a codeword of blocklength n over the *complex-valued* AWGN channel corresponds to the transmission of a codeword of blocklength $2n$ over the *real-valued* AWGN channel with the same SNR.

It has been observed that a good approximation for $R^*(n, \epsilon)$ can be obtained by replacing the remainder terms on the right-hand side of (5) by $(\log n)/(2n)$ [12], [14]. The resulting approximation, which is commonly referred to as *normal approximation*, is plotted in Fig. 4, together with nonasymptotic upper and lower bounds on $R^*(n, \epsilon)$ (see [12] for details).

As shown in the figure, the upper and lower bounds provide an accurate characterization of $R^*(n, \epsilon)$, which lies in the shaded region. According to the bounds, to operate at 70% of capacity with a packet error rate of 10^{-3} , i.e., at 0.7 bits/channel use, it is sufficient to use codes whose blocklength is between 110 and 138 channel uses. For the parameters considered in the figure, the normal approximation is indistinguishable from the achievability bound. We also see that capacity is an inaccurate performance metric for packet sizes that are as short as the ones considered in the figure.

C. Fading Channels

We shall next discuss how to extend the results reported in Section III-B for the AWGN case to multiple-antenna fading channels. Throughout, we shall focus on the *memoryless block-fading model* [15], depicted in Fig. 5, according to which the fading coefficient stays constant over n_c channel uses and then changes independently. In general, n_c can be interpreted as the number of “time-frequency slots” over which the channel does not change. We shall refer to each interval over which the fading coefficients do not change as a *coherence interval*.

The memoryless block-fading model is perhaps the simplest model to capture channel variations in wireless channels. Although inferior in accuracy to stationary channel models, where the channel varies continuously (see e.g., [16]), its simplicity enables analytical approaches that are currently out of reach for more sophisticated models.

For ease of notation, we shall write the symbols to be transmitted in each coherence interval in a $n_c \times m_t$ matrix whose entry at position (i, j) corresponds to the i th symbol transmitted from antenna j . Likewise, we write the received symbols in a $n_c \times m_r$ matrix. Within the k th coherence interval, the input-output relation of the block-fading channel with m_t transmit and m_r receive antennas is given by

$$\mathbf{Y}_k = \mathbf{X}_k \mathbf{H}_k + \mathbf{W}_k. \quad (9)$$

Here, $\mathbf{X}_k \in \mathbb{C}^{n_c \times m_t}$ and $\mathbf{Y}_k \in \mathbb{C}^{n_c \times m_r}$ are the transmitted and received matrices, respectively; $\mathbf{H}_k \in \mathbb{C}^{m_t \times m_r}$ denotes the fading matrix; $\mathbf{W}_k \in \mathbb{C}^{n_c \times m_r}$ denotes the additive noise, which is assumed to have independent and identically distributed (i.i.d.), zero-mean, unit-variance, complex Gaussian entries. For the sake of simplicity, we assume *Rayleigh fading*, i.e., we assume that the fading matrix $\mathbf{H}_k$ has i.i.d., zero-mean, unit-variance, complex Gaussian entries. However, this assumption is not essential. In fact, most results presented in this paper were either originally derived for more general fading distributions or can be generalized with some effort. For convenience, we shall assume that each codeword spans ℓ coherence intervals, i.e., $n = \ell n_c$.

We shall say that CSI is available at the transmitter, receiver, or both if the corresponding blocks have access to the realization of

Fig. 5. Block-fading model: the fading coefficient stays constant over n_c channel uses (coherence interval) and then changes to an independent realization. Coding is performed over ℓ coherence intervals (number of time-frequency diversity branches).

$\mathbf{H}_1, \dots, \mathbf{H}_n$. In practice, CSI at the transmitter allows for transmission strategies that make use of the actual fading realization, thereby using the available transmit power more efficiently; CSI at the receiver facilitates the decoding task. Note that CSI at the receiver can be acquired by transmitting training sequences (so-called *pilots*) that are used at the receiver to estimate the channel. CSI at the transmitter can, for example, be established by feeding channel estimates from the receiver back to the transmitter. However, the transmission of training sequences incurs a rate loss, sometimes referred to as *channel-estimation overhead*. Likewise, the creation of a feedback link is associated with additional costs or overheads. Analyses relying on the assumption that CSI is available at the transmitter, receiver or both simply ignore these overheads. In this spirit, analyses that are based on the assumption that no CSI is available at the receiver do not assume that the receiver does not perform a channel estimation. On the contrary, they account for the overhead associated with the acquisition of CSI. For example, the transmission of training sequences can be viewed as a specific form of coding. Thus, by analyzing the fading channel (9) under the assumption that no CSI is available at the receiver, the rate loss incurred by the transmission of pilot symbols is automatically accounted for.

1) *Capacity-versus-outage at finite blocklength*: We shall first discuss the case where the channel remains constant over the packet duration, i.e., $\ell = 1$. In this case, the fading channel is said to be *quasi-static*, to reflect that the fading matrix is random but stays constant during the packet transmission.² When communicating over quasi-static fading channels at a given rate R , the realization of the random fading matrix $\mathbf{H}_k$ may be very small, in which case the decoder will not be able to guess the transmitted information bits correctly, no matter how large we choose the packet length n . In this case, the channel is said to be in *outage*. For fading distributions for which the fading coefficient can be arbitrarily small (which is, for example, the case for Rayleigh fading), the probability of an outage is positive. Hence, the packet error probability is bounded away from zero for every positive rate $R > 0$ and the capacity, defined as the largest rate for which reliable communication is feasible, is zero [8], [9].

One may argue that the definition of capacity is too restrictive for quasi-static channels. Indeed, for sufficiently small (but positive) rates, the probability that the channel is in outage is typically small. Thus, while reliable communication cannot be guaranteed because there is always a chance that the channel is in outage, the probability that this happens is small. In other

²In the information theory literature, the quasi-static channel model belongs to the class of *composite channels* [9], [17], also known as *mixed channels* [18, Sec. 3.3].

words, most of the time the channel is not in outage and reliable communication can be achieved by choosing a sufficiently large packet length. Capacity, however, is determined by those rare events where the channel is in outage. For applications where a positive packet error probability is acceptable, the outage capacity C_ϵ is arguably a more relevant performance metric than capacity, because it allows for outage events as long as they happen with probability less than ϵ .

The outage capacity is often regarded as a performance metric for delay-constrained communication over slowly-varying fading channels (see, e.g., [19]). In fact, the assumption that the fading matrix stays constant during the packet transmission seems plausible only if the packet size is small. Nevertheless, the definition of outage capacity requires that the blocklength tends to infinity; cf. (3). For example, for a single-antenna system, the outage probability as a function of the rate R is given by [20], [19], [17]

$$P_{\text{out}}(R) = \mathbb{P}[\log(1 + |H|^2 \rho) < R] \quad (10)$$

and the outage capacity C_ϵ is the supremum of all rates R satisfying $P_{\text{out}}(R) \leq \epsilon$, namely

$$C_\epsilon = \sup\{R : P_{\text{out}}(R) \leq \epsilon\}. \quad (11)$$

The rationale behind this result is that, for every realization of the fading coefficient $H = h$, the quasi-static fading channel can be viewed as an AWGN channel with channel gain h , for which communication with arbitrarily small packet error probability is feasible if, and only if, $R < \log(1 + |h|^2 \rho)$, provided that the packet length is sufficiently large.³ However, it is *prima facie* unclear whether the quantity $\log(1 + |h|^2 \rho)$ is meaningful when the packet size is small.

To better understand the relevance of the outage capacity for delay-constrained communication, a more refined analysis of $R^*(n, \epsilon)$ was presented in [21]. It was shown that [21, Ths. 3 and 9]

$$R^*(n, \epsilon) = C_\epsilon + \mathcal{O}\left(\frac{\log n}{n}\right) \quad (12)$$

irrespective of the number of transmit and receive antennas, and irrespective of whether CSI is available to transmitter, receiver, or both. Comparing (12) with (5), we observe that for the quasi-static fading case the channel dispersion is zero, i.e., the $1/\sqrt{n}$ rate penalty is absent. This suggests that $R^*(n, \epsilon)$ converges quickly to C_ϵ as n tends to infinity, thereby indicating that the outage capacity is indeed a meaningful performance metric for delay-constrained communication over slowly-varying fading channels. Numerical examples that support this claim can be found in [21, Sec. VI]. Furthermore, a simple approximation for $R^*(n, \epsilon)$ is proposed in [21, Eqs. (59) and (95)]. For the single-antenna case, this approximation can be written in the following form [21], [22]

$$\epsilon \approx \mathbb{E} \left[Q \left(\frac{C(\rho |H|^2) + (\log n)/(2n) - R^*(n, \epsilon)}{\sqrt{V(\rho |H|^2)/n}} \right) \right]. \quad (13)$$

³Indeed, the capacity of the AWGN channel with channel gain h follows from (7) by changing the SNR from ρ to $|h|^2 \rho$.

Here, $C(\cdot)$ and $V(\cdot)$ are the functions defined in (7) and (8), respectively.

The asymptotic expansion (12) provides mathematical support to the observation reported by several researchers in the past that the outage probability describes accurately the performance over quasi-static fading channels of actual codes (see [19] and references therein). The intuition behind this result is that the dominant error event over quasi-static fading channels is that the channel is in a deep fade. Since the transmitted symbols experience all the same fading, it follows that coding is not helpful against deep fades in the quasi-static fading scenario, hence $R^*(n, \epsilon)$ is close to C_ϵ already for small blocklengths.

It has been observed that the outage capacity C_ϵ does not depend on whether CSI is available at the receiver [9, p. 2632], [21, Ths. 3 and 9]. Intuitively, this is true because the fading matrix stays constant during the whole transmission, so it can be accurately estimated at the receiver through the transmission of $\sqrt{n}$ pilot symbols with no rate penalty as the packet length n tends to infinity. This in turn implies that the outage capacity does not capture the channel-estimation overhead. Consequently, outage capacity is an inaccurate performance metric when the coherence interval n_c is small.

2) *Tradeoff between diversity, multiplexing, and channel estimation*: When communicating over multiple-input multiple-output fading channels, a crucial question is whether the spatial degrees of freedom offered by the antennas should be used to lower the packet error probability for a given data rate (through the exploitation of *spatial diversity*) or to increase the data rate for a given packet error probability (through the exploitation of *spatial multiplexing*). These two effects cannot be harvested concurrently, but there exists a fundamental tradeoff between diversity and multiplexing. This tradeoff admits a particularly simple characterization in the high-SNR regime [23].

Specifically, Zheng and Tse [23] defined the *diversity-multiplexing tradeoff* as follows. Assume that ℓ and n_c are fixed. Further assume that the packet error probability vanishes with increasing ρ as

$$\epsilon(\rho) = \rho^{-d\ell} \quad (14)$$

where $d \in \{1, \dots, m_t m_r\}$ is the so-called *spatial diversity gain*. The *multiplexing gain* $r(d)$ corresponding to the diversity gain d is defined as

$$r(d) = \lim_{\rho \rightarrow \infty} \frac{R^*(n, \epsilon(\rho))}{\log \rho}. \quad (15)$$

For the case where CSI is available at the receiver and $n_c \geq m_t$, one can show that $r(d)$ is the piecewise linear function connecting the points [23], [24]

$$r((m_t - k)(m_r - k)) = k, \quad k = 0, \dots, \min\{m_t, m_r\}. \quad (16)$$

Let $m^* = \min\{m_t, m_r, \lfloor n_c/2 \rfloor\}$, where $\lfloor a \rfloor$ denotes the largest integer that is not larger than a . For the case where no CSI is available at the receiver and $n_c \geq 2m^* + m_r + 1$, the multiplexing gain is given by [25], [26]

$$r((m_t - k)(m_r - k)) = \left(1 - \frac{m^*}{n_c}\right) k. \quad (17)$$

It is thus equal to (16) multiplied by $(1 - m^*/n_c)$. The expressions (16) and (17) describe elegantly and succinctly the tradeoff between diversity gain and multiplexing gain at high SNR.

Note that m^*/n_c is roughly the number of pilots per time-frequency slot needed to learn the channel at the receiver when m^* transmit antennas are used. A comparison of (17) with (16) thus illustrates how an analysis of the diversity-multiplexing tradeoff under the assumption of no CSI at the receiver captures the channel-estimation overhead.

It has been recently demonstrated that for data packets of 1000 channel uses or more and for moderately low packet-error probabilities (around 10^{-2}), one should typically operate at maximum multiplexing [27]. In this regime, which is relevant for current cellular systems, diversity-exploiting techniques are detrimental both for high- and for low-mobility users. For high-mobility users (where n_c is significantly smaller than the packet size n), abundant time and frequency selectivity is available, so diversity-exploiting techniques are superfluous. For low-mobility users (where n_c is large), the fading coefficients can be learnt at the transmitter and outage events can be avoided altogether by rate adaptation.

However, when the packet size becomes small and/or smaller packet-error probabilities are required, these conclusions may cease to be valid. For example, for packet lengths of, say, 100 channel uses (which is roughly equal to a LTE resource block) and packet error probability of 10^{-5} or lower, spatial diversity may be more beneficial than spatial multiplexing. Furthermore, when the coherence interval n_c is small, the cost of estimating the fading coefficients may be significant and must therefore be taken into consideration.

Studies based on capacity or outage capacity are inherently incapable of illuminating the entire *diversity-multiplexing-channel-estimation tradeoff*. Indeed, recall that the capacity is defined as the largest rate at which reliable communication is feasible as the packet length tends to infinity. Specialized to the block-fading channel, capacity is typically studied by letting the number of time-frequency diversity branches ℓ grow to infinity while holding the coherence interval n_c fixed. For example, when $n_c > 1$ and no CSI is available, the capacity is given by [28]

$$C(\rho) = m^* \left(1 - \frac{m^*}{n_c}\right) \log \rho + \mathcal{O}(1) \quad (18)$$

where $\mathcal{O}(1)$ comprises error terms that are bounded in the SNR. Observe that (18) reflects the cost of estimating the fading matrix through n_c , but it hides away the effects of spatial diversity, since by letting ℓ tend to infinity we achieve an infinite time-frequency diversity gain already through coding. Conversely, the definition of outage capacity is based on the assumption that the coherence interval n_c grows to infinity while the number of diversity branches ℓ is held fixed (cf. Section III-C1 where we chose $\ell = 1$). For example, in the absence of CSI, the outage capacity is given by [8]

$$C_\epsilon(\rho) = \sup \left\{ R : \inf_{\mathbf{Q}^\ell} P_{\text{out}}(R, \mathbf{Q}^\ell) \leq \epsilon \right\} \quad (19)$$

Fig. 6. Upper and lower bounds on the maximum coding rate $R^*(n, \epsilon)$ for a Rayleigh block-fading channel with $m_t = m_r = 2$, $n = 168$, $\epsilon = 10^{-5}$, $\rho = 6$ dB. The maximum coding rate lies in the shaded area between the upper and lower bound. Upper and lower bounds on the maximum coding rate achievable using an Alamouti inner code are also depicted to indicate the performance of a configuration in which transmit antennas are used to provide exclusively transmit diversity. The curve for the outage capacity has been obtained by numerically evaluating (19). The curve for the ergodic capacity follows by tightening (18); see [31] for more details. This figure appeared first in [31].

where $P_{\text{out}}(R, \mathbf{Q}^\ell)$ denotes the outage probability

$$P_{\text{out}}(R, \mathbf{Q}^\ell) = \mathbb{P} \left[\frac{1}{\ell} \sum_{k=1}^{\ell} \log \det (\mathbf{I} + \mathbb{H}_k^H \mathbf{Q}_k \mathbb{H}_k) \leq R \right] \quad (20)$$

and where the infimum in (19) is over all positive-definite $m_t \times m_t$ matrices $\{\mathbf{Q}_1, \dots, \mathbf{Q}_\ell\} = \mathbf{Q}^\ell$ whose traces satisfy $(1/\ell) \sum_{k=1}^{\ell} \text{tr}(\mathbf{Q}_k) \leq \rho$. In (20), the symbol $\mathbf{I}$ denotes the identity matrix, and $(\cdot)^H$ denotes Hermitian conjugation. For $\ell = 1$, the outage probability (20) specializes to (10). Observe that (19) captures the effects of spatial and time-frequency diversity through the dimension of $\mathbb{H}_k$ ($m_t \times m_r$) and the value of ℓ . However, as already mentioned at the end of Section III-C1, it hides away the cost of estimating the fading coefficient, since for an infinite coherence interval n_c the channel can be estimated perfectly without a rate penalty.

To investigate the entire diversity-multiplexing-channel-estimation tradeoff for small packet lengths, bounds on $R^*(n, \epsilon)$ were presented in [29]–[31]. Here, we provide an example, taken from [31], which illustrates the benefit of a nonasymptotic analysis of the diversity-multiplexing-channel-estimation tradeoff. Specifically, we consider a scenario based on the 3GPP LTE standard [27] where the packet size is $n = 168$ symbols, which corresponds to 14 OFDM symbols, each consisting of 12 tones. We set the SNR to 6 dB and the packet error rate to 10^{-5} , which corresponds to a URC scenario, and compute the bounds on the maximum coding rate obtained in [31] as a function of the coherence time n_c or, equivalently, the number of diversity branches ℓ (recall that $n = \ell n_c$) for a 2×2 MIMO system.

The upper and lower bounds on the maximum coding rate obtained in [31] for the above example are depicted in Fig. 6. We see from the figure that, given n and ϵ , the rate $R^*(n, \epsilon)$ is not monotonic in the coherence interval n_c , but there exists a value n_c^* (in this case 14) that maximizes the rate. This accentuates the fundamental tradeoff between time-frequency diversity (which

decreases with n_c) and the ability of estimating the fading coefficient (which increases with n_c).

We further observe that both outage capacity and capacity (computed for the scenario where CSI is not available at the receiver—see [32] for a recent review) fail to capture this tradeoff, although their intersection predicts surprisingly well the rate-maximizing coherence interval. Indeed, the outage capacity only captures the increase in time-frequency diversity, whereas capacity only captures the channel-estimation overhead. We also note that when the coherence interval is smaller than 8 channel uses, one of the two transmit antennas should be switched off, because the cost of estimating the fading coefficients overcomes the benefit of using two antennas at the transmitter.

In Fig. 6, we also depict bounds on the maximum coding rate obtainable using an Alamouti inner code [33], a configuration in which the transmit antennas are used to provide exclusively transmit diversity. Since the gap between the rate achievable using Alamouti and the maximum coding rate converse is small, we conclude that for the scenario considered in Fig. 6, the available transmit antennas should be used to provide diversity and not multiplexing.

D. Channel Dispersion versus Error Exponents

Traditionally, the tradeoff between reliability and throughput for small packet lengths has been studied by means of *error exponents*. In this section, we briefly discuss the relation between error exponents and asymptotic expansions of the maximum coding rate, such as (5), that express $R^*(n, \epsilon)$ as a function of channel capacity and channel dispersion.

Recall that the capacity C is the largest transmission rate for which the packet error probability P_e vanishes as the packet length n tends to infinity. It turns out that for every fixed transmission rate $R < C$, the packet error probability vanishes even exponentially in n [34]. It is therefore meaningful to expand P_e for every fixed $R < C$ as

$$P_e = e^{-n[E(R)+o(1)]} \quad (21)$$

where $o(1)$ comprises remainder terms that vanish as n tends to infinity. The exponent $E(R)$ in (21) is referred to as the *error exponent* corresponding to the rate R . For more details on error exponents, see [35] and references therein.

Intuitively, (21) characterizes the packet error probability P_e as a function of n and R . In contrast, (5) characterizes the transmission rate R as a function of n and P_e . It may therefore seem plausible to view the expansions (5) and (21) as two equivalent characterizations of the triple (R, n, P_e) . However, (5) and (21) contain remainder terms, specifically $\mathcal{O}(\log n/n)$ and $o(1)$, and are therefore only accurate if the packet length n is sufficiently large. Since P_e decays exponentially in n , it follows that for packet lengths for which (21) is a good approximation, P_e is very small. Likewise, $R^*(n, \epsilon)$ converges to the capacity C as n tends to infinity, so for packet lengths for which (5) is a good approximation, $R^*(n, \epsilon)$ is very close to C .

In summary, the error exponent $E(R)$ characterizes the triple (R, n, P_e) when the rate $R < C$ is held fixed and P_e is very small. In contrast, the channel dispersion V characterizes the triple (R, n, P_e) when $P_e \leq \epsilon$ is held fixed and R is very close

to capacity. For wireless communications, where a small but positive packet error probability can be tolerated, the asymptotic expansion of $R^*(n, \epsilon)$ provided in (5) seems more meaningful.

E. Further Works

The work by Polyanskiy, Poor, and Verdú [12] has triggered a renewed interest in the problem of finite-blocklength information theory. This is currently a very active research area. Here, we provide a (necessarily not exhaustive) list of related works dealing with wireless communications at finite blocklength.

When CSI is available at the receiver, the dispersion of fading channels was obtained in [37]–[39] for specific scenarios. Upper and lower bounds on the second-order coding rate of quasi-static multiple-input multiple-output (MIMO) Rayleigh-fading channels have been reported in [40] for the asymptotically ergodic setup when the number of antennas grows linearly with the blocklength. The channel dispersion of single-antenna, quasi-static fading channels with perfect CSI at both the transmitter and the receiver and a long-term power constraint has been given in [41], [42].

For discrete-memoryless channels, feedback combined with variable-length coding has been shown to dramatically improve the speed at which the maximum coding rate approaches capacity [43]. Such improvements can be achieved by letting the receiver feed back a stop signal to inform the transmitter that decoding has been successful (*stop feedback*, also known as *decision feedback*). One can relax the assumption that decoding is attempted after each symbol, with marginal performance losses [44].

Coding schemes approaching the performance predicted by finite-blocklength bounds have been also proposed. In [45], list decoding of polar codes is shown (through numerical simulations) to operate close to the maximum coding rate. The finite-blocklength gap to capacity exhibited by polar codes has been characterized up to second order (in terms of the so-called scaling exponent) in [46]–[48]. A comparison between the finite-blocklength performance of convolutional codes (both with Viterbi and with sequential decoding) and LDPC codes is provided in [49]. Bounds and exact characterizations on the error-vs-delay tradeoff for codes of very small cardinality have been recently provided in [50].

In Fig. 7, we provide an overview of the performance of codes for the binary-input AWGN channel from 1980 to present. The first eight codes in the legend of Fig. 7 are from [51]. The BCH (Koetter-Vardy) code is from [52, Fig. 2]; here, the decoder uses soft-decision list decoding. As shown in the figure, ordered-statistic decoding (OSD) [53] of BCH codes improves the performance further. OSD decoding of nonbinary LDPC codes turn out to yield similar performance as BCH-OSD. Indeed, this decoding technique seems to yield state-of-the-art performance for very short packets (between 100 and 200). For larger packet size, list decoding of polar codes combined with CRC [54] and multi-edge (ME) type LDPC codes [55] are a competitive benchmark.

Moving to coding schemes exploiting decision feedback, designs based on tail-biting convolutional codes combined with the reliability-output Viterbi algorithm have been proposed in [56]. Finally, second-order characterizations of the coding rates for

Fig. 7. Normalized rates (with respect to $R^*(n, \epsilon)$) over binary-input AWGN channel; $\epsilon = 10^{-4}$. An earlier version of this figure appeared first in [36].

some problems in network information theory have recently been obtained. A comprehensive review is provided in [57].

F. *spectre*: short-packet communication toolbox

To optimally design communication protocols for short-packet transmission, one needs to rely on accurate physical-layer performance metrics. *spectre*—*short-packet communication toolbox* [58] is a collection of numerical routines for the evaluation of upper and lower bounds on the maximum coding rate for popular channel models, including the AWGN channel, the quasi-static fading channel, and the Rayleigh block-fading channel. This toolbox can be freely accessed online and is under development. All the numerical simulations reported in this paper can be reproduced using *spectre* routines.

IV. COMMUNICATION PROTOCOLS FOR SHORT PACKETS

In simple terms, a communication protocol is a distributed algorithm that determines the actions of the actors involved in the communication process. Protocol information, also referred to as metadata or control information, can be understood as a *source code* [59] that ensures correct operation of the protocols and describes, e.g., the current protocol state, the packet length, or the addresses of the involved actors.

Only few results are available on the information-theoretic design of communication protocols, e.g., [60]–[62], and most of them deal with the (source coding) problem of how to encode the network/link state that needs to be communicated as a protocol information. The problem of how to transmit the protocol-related metadata has been largely left to heuristic approaches, such as the use of repetition coding. Broadly speaking, whereas information theorists busy themselves with developing capacity-approaching schemes for the reliable transmission of the information payload, they often see the design of metadata as something outside their competence area, or as stated in [43]: “...control information is not under the purview of the physical layer ...” Such a line of thinking is fully justifiable when the ratio between the data and metadata is the one depicted in Fig. 1(a), where the metadata occupies a small fraction of the overall packet length. However,

for applications where the data is comparable in size to the metadata—see Fig. 1(b)—this approach seems questionable.

In the following, we shall argue that a thorough understanding of how the maximum transmission rate $R^*(n, \epsilon)$ depends on the packet length n and on the packet error probability ϵ is also beneficial for protocol design. As mentioned above, only few results are available on the information-theoretic design of protocols, and there is even less work that considers protocol design for short-packet transmission, e.g., [63]–[66]. This section is therefore based on three simple examples that illustrate how the tradeoffs brought by short-packet transmissions affect protocol design. We believe that these examples unveil a number of interesting tradeoffs worth exploring and we hope that they may motivate the research community to pursue a better theoretical understanding of protocol design.

For simplicity, we assume throughout this section an AWGN channel with SNR $\rho = 10$, and we approximate $R^*(n, \epsilon)$ as

$$R^*(n, \epsilon) \approx C - \sqrt{\frac{V}{n}} Q^{-1}(\epsilon) + \frac{1}{2n} \log n \quad (22)$$

where C and V are given in (7) and (8), respectively.⁴ We expect that tradeoffs similar as the ones we shall illustrate for the AWGN case will occur also for the fading case (see [67] for an example that supports this claim). Solving (22) for ϵ yields the following approximation of the packet error probability as a function of the packet length n and the number of information bits $k = Rn$ which we shall use throughout this section:

$$\epsilon^*(k, n) \approx Q\left(\frac{nC - k + (\log n)/2}{\sqrt{nV}}\right). \quad (23)$$

A. *Reliable Communication Between Two Nodes*

Consider the two-way communication protocol illustrated in Fig. 8, where the nodes acknowledge the correct reception of a data packet by transmitting an ACK. The correct transmission of a data packet from, say, node 1 to node 2 would result in the following protocol exchange sequence:

⁴Recall that, as mentioned in Section III, replacing the remainder terms in (5) by $\frac{1}{2n} \log n$ yields a good approximation for $R^*(n, \epsilon)$.

Fig. 8. Scenario of a two-way communication with data from node 1 and acknowledgement from node 2.

- 1) The packet from node 1 is correctly received by node 2. We shall denote the probability of this event by $1 - \epsilon_1$;
- 2) Node 2 sends an ACK to node 1. We shall denote the probability that an ACK is received correctly by $1 - \epsilon_2$.

As noted in [68], if we communicate over a noisy channel and we are restricted to use a finite number of channel uses, then no protocol will be able to achieve perfectly reliable communication. Indeed, it is possible that either a packet is received incorrectly (an event which has probability ϵ_1) or that the ACK is received incorrectly (which happens with probability ϵ_2). By (23), decoding errors are particularly relevant if the packet size is small, in which case ϵ_1 and ϵ_2 are large. Thus, the often-made assumptions of perfect error detection or perfect ACK-transmission (so-called “1-bit feedback”) are particularly misleading if the considered packet length is small.

Let us consider the following example. Let each node have a 6-byte address and assume that node 1 has 12 data bytes to send. Assume that the packet sent by node 1 contains the source address, the destination address, one bit for flow control and the data bytes. Hence, node 1 transmits $k_{i,1} = 96$ data bits and $k_{o,1} = 97$ metadata bits, resulting in $k_1 = k_{i,1} + k_{o,1} = 193$ bits. The ACK packet sent by node 2 consists of the source address and the destination address and one ACK bit.⁵ For the ACK packet, this yields $k_{i,2} = 0$ data bits and $k_{o,2} = 97$ metadata bits, so $k_2 = k_{i,2} + k_{o,2} = 97$ bits. Let n be the total number of channel uses available to send the data and the ACK. To optimize the protocol, we may want to find the optimal number of channel uses n_1 by node 1 and $n_2 = n - n_1$ by node 2 such that the reliability of the transmission, given by $(1 - \epsilon^*(k_1, n_1))(1 - \epsilon^*(k_2, n_2))$, is maximized. These values can be found numerically using the approximation (23). For example, the minimum value of n that offers reliability of transmission

$$(1 - \epsilon^*(k_1, n_1))(1 - \epsilon^*(k_2, n_2)) > 0.999$$

is $n = 203$, out of which $n_1 = 132$ channel uses are for sending the data packet and $n_2 = 71$ channel uses are for sending the ACK.

As another example, fix $n = 250$ as the maximal allowed number of channel uses. The numerical optimization that yields the largest reliability $(1 - \epsilon^*(k_1, n_1))(1 - \epsilon^*(k_2, n_2))$ gives $n_1 = 158$ and $n_2 = 92$. The resulting reliability is almost 1 and the resulting throughput is $(1 - \epsilon^*(k_1, n_1))(1 - \epsilon^*(k_2, n_2))k_{i,1}/n = 0.384$ bits/channel use.

In many cases, it is not practical to have variable values for n_1 and n_2 , and a fixed time division duplex (TDD) structure in which $n_1 = n_2$ is preferred. In such a structure, there is no need

⁵Note that the source/destination addresses are necessary in order to uniquely identify the link to which the ACK belongs.

Fig. 9. Example of a scenario with downlink communication from a Base Station over a broadcast channel to three nodes.

of explicit ACK packets, since the acknowledgement is typically piggybacked on a data packet. In order to align this scenario with the last example, we assume that $n_1 = n_2 = 125$, such that the acknowledgment for the packet arrives within $n = 250$ channel uses from the start of the data transmission. A packet sent by nodes 1 and 2 contains 194 bits, of which 96 are data bits, 96 are bits for addresses, 1 bit is for flow control, and 1 bit for the acknowledgment. Evaluating (23) for these parameters gives $\epsilon^*(k_1, n_1) = \epsilon^*(k_2, n_2) = 0.0118$. Observe that the reliability is markedly decreased, although the throughput is almost doubled to 0.759 bits/channel use.

These simple examples show that adjusting the packet length and the coding rate has the potential to yield high reliability. Note, however, that flexibility in the packet length necessarily implies that the receiver needs to acquire information about it. This means that the protocol needs to reserve some bits within each packet for the metadata that describes the packet length. Our simple calculations have not accounted for this overhead.

The use of a predefined slot length yields a robust system design, since no additional error is caused by the exchange of length-related metadata. This indicates that, in designing protocols that support ultra-high reliability, a holistic approach is required that includes all elements of the protocol/metadata that are commonly assumed to be perfectly received.

B. Downlink Multi-User Communication

We now turn to an example in which a base station (BS) transmits in the downlink to M devices; see Fig. 9. The BS needs to unicast D bits to each device. Hence, it sends in total MD bits. As a reference, we consider a protocol where the BS serves the users in a time division multiple access (TDMA) manner: each device receives its D bits in a dedicated time slot that consists of n channel uses. Thus, the TDMA frame consists of M slots with a total of Mn channel uses. In order to avoid transmission of metadata, we assume that the system operates in a *circuit-switched* TDMA manner: (a) all devices and the BS are perfectly synchronized to a common clock; (b) each device knows the slot in which it will receive its data. The performance of this idealized scheme can be considered as an upper bound on the performance of practical systems, such as GSM, as it assumes that there is a genie that helps the devices remain synchronized.

The approximation on $\epsilon^*(k, n)$ in (23) suggests that, for short packet sizes, it may be more efficient to encode a larger amount

Fig. 10. Example of a scenario with uplink communication from a set of three nodes over a multiple access channel to a common Base Station.

of data than the one intended to each device. Thus, instead of using TDMA, the BS may *concatenate* all the data packets for the individual devices. In this way, the BS constructs a single data packet of MD bits that should be *broadcasted* by using Mn channel uses. Each receiving device then decodes the whole data packet and extracts the bits it is interested in from the decoded MD bits.

As a concrete example, assume that the BS wishes to transmit $D = 192$ bits to each device and that there are $M = 10$ devices. Furthermore, assume that $n = 125$. We consider for simplicity one-shot communication. Accounting for retransmissions would require a more elaborate discussion.

In the reference scheme, the probability of error experienced by each device is 0.007. If concatenation is used, however, the probability of error drops to about 10^{-12} , which puts the transmission scheme in a different reliability class, while preserving the same overall delay. The price paid is the fact that each device needs to decode more data than in the reference scheme.

Note that if one ignored the dependency of the packet error probability ϵ^* on the packet size n , one would conclude that the circuit-switched TDMA protocol is the most efficient, since all channel uses can be devoted to the transmission of payload bits. In contrast, by taking the dependence of ϵ^* on n into account, we see that an unconventional protocol that concatenates the data intended to different devices outperforms the traditional TDMA protocol by orders of magnitude in terms of reliability.

C. Uplink Multi-User Communication

Our last example is related to the scenario depicted in Fig. 10 in which M devices run a random access protocol in order to transmit to a common receiver BS. Specifically, there are M users, each sending D bits to the BS. Each packet should be delivered within a time that corresponds to n channel uses. These n channel uses are divided into K equally-sized slots of $n_K = n/K$ channel uses. The devices apply a simple framed ALOHA protocol: each device picks randomly one of the K slots in the frame and sends its packet. If two or more users pick the same slot, then a collision occurs and none of the packets is received correctly (see [67] for a more elaborate example). If only one device picks a particular slot (singleton slot), then the error probability is calculated using (23) for D bits and n_K channel uses.

We are interested in the following question: given M , D , and n , how should we choose the slot size n_K in order to

maximize the packet transmission reliability experienced by each individual device? This problem entails a tradeoff between the probability of collision and the number of channel uses available for each packet, which by (23) affects the achievable packet error probability in a singleton slot. Indeed, if K increases, then the probability of a collision decreases, while the packet error probability for a singleton slot increases. Conversely, if K decreases, then the probability of collision increases, while the packet error probability for a singleton slot decreases. The probability of successful transmission is given by

$$P_S = \frac{M}{K} \left(1 - \frac{1}{K}\right)^{M-1} \cdot (1 - \epsilon^*(D, n_K)). \quad (24)$$

Here, $(M/K)(1 - 1/K)^{M-1}$ is the probability of not experiencing collision, and $\epsilon^*(D, n_K)$ is the probability of error for a packet of D bits sent over n_K channel uses, which can be approximated by (23).

As a concrete example, let us consider the setup where $D = 192$ bits, $M = 10$ devices, and $n = 800$ channel uses. The number of slots that maximizes (24) is $K = 6$. In contrast, the classic framed-ALOHA analysis, which assumes that packets are decoded correctly if no collisions occur (i.e., $\epsilon^* = 0$ in (24)), yields $K = M = 10$. In fact, the same is true for any positive error probability ϵ^* that does not depend on n_K .

V. CONCLUSIONS

Motivated by the advent of novel wireless applications such as massive machine-to-machine and ultra-reliable communications, we have provided a review of recent advances in the theory of short-packet communications and demonstrated through three examples how this theory can help designing novel efficient communication protocols that are suited to short-packet transmissions. The key insight is that—when short packets are transmitted—it is crucial to take into account the communication resources that are invested in the transmission of metadata. This unveils tradeoffs that are not well understood yet and that deserve further research, both on the theoretical and on the applied side.

ACKNOWLEDGMENT

We would like to thank Yuri Polyanskiy for letting us reproduce Fig. 7, Gianluigi Liva for bringing our attention to BCH codes combined with OSD (see Fig. 7), and Erik G. Ström for fruitful discussions. We further would like to thank the Managing Editor, Vaishali Damle, as well as the anonymous referees for their valuable comments.

REFERENCES

- [1] E. Dahlman, G. Mildh, S. Parkvall, J. Peisa, J. Sachs, and Y. Selén, “5G radio access,” *Ericsson Review*, no. 6, Jun. 2014.
- [2] F. Boccardi, R. Heath, A. Lozano, T. Marzetta, and P. Popovski, “Five disruptive technology directions for 5G,” *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 74–80, Feb. 2014.
- [3] A. Osseiran, F. Boccardi, V. Braun, K. Kusume, P. Marsch, M. Maternia, O. Queseth, M. Schellmann, H. Schotten, H. Taoka, H. Tullberg, M. Uusitalo, B. Timus, and M. Fallgren, “Scenarios for 5G mobile and wireless communications: the vision of the METIS project,” *IEEE Commun. Mag.*, vol. 52, no. 5, pp. 26–35, May 2014.
- [4] P. Popovski, “Ultra-reliable communication in 5G wireless systems,” in *IEEE Int. Conf. 5G for Ubiquitous Connectivity*, Levi, Finland, Nov. 2014.

- [5] N. A. Johansson, Y.-P. E. Wang, E. Eriksson, and M. Hessler, "Radio access for ultra-reliable and low-latency 5G communications," in *Proc. IEEE Int. Conf. Commun. (ICC)*, London, U.K., Jun. 2015.
- [6] O. N. C. Yilmaz, Y.-P. E. Wang, N. A. Johansson, N. Barhmi, S. A. Ashraf, and J. Sachs, "Analysis of ultra-reliable and low-latency 5G communication for a factory automation use case," in *Proc. IEEE Int. Conf. Commun. (ICC)*, London, U.K., Jun. 2015.
- [7] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379–423 and 623–656, July and October 1948.
- [8] L. H. Ozarow, S. Shamai (Shitz), and A. D. Wyner, "Information theoretic considerations for cellular mobile radio," *IEEE Trans. Veh. Technol.*, vol. 43, no. 2, pp. 359–378, May 1994.
- [9] E. Biglieri, J. G. Proakis, and S. Shamai (Shitz), "Fading channels: Information-theoretic and communications aspects," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2619–2692, Oct. 1998.
- [10] D. N. C. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [11] R. Costa, M. Langberg, and J. Barros, "One-shot capacity of discrete channels," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Austin, TX, Jun. 2010, pp. 211–215.
- [12] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Channel coding rate in the finite blocklength regime," *IEEE Trans. Inf. Theory*, vol. 56, no. 5, pp. 2307–2359, May 2010.
- [13] C. E. Shannon, "Probability of error for optimal codes in a Gaussian channel," *Bell Syst. Tech. J.*, vol. 38, pp. 611–656, 1959.
- [14] V. Y. F. Tan and M. Tomamichel, "The third-order term in the normal approximation for the AWGN channel," *IEEE Trans. Inf. Theory*, vol. 61, no. 5, pp. 2430–2438, May 2015.
- [15] T. L. Marzetta and B. M. Hochwald, "Capacity of a mobile multiple-antenna communication link in Rayleigh flat fading," *IEEE Trans. Inf. Theory*, vol. 45, no. 1, pp. 139–157, Jan. 1999.
- [16] G. Matz and F. Hlawatsch, "Fundamentals of time-varying communication channels," in *Wireless Communications over Rapidly Time-Varying Channels*, F. Hlawatsch and G. Matz, Eds. Academic Press, Mar. 2011, ch. 1, pp. 1–63.
- [17] M. Effros, A. Goldsmith, and Y. Liang, "Generalizing capacity: New definitions and capacity theorems for composite channels," *IEEE Trans. Inf. Theory*, vol. 56, no. 7, pp. 3069–3087, Jul. 2010.
- [18] T. S. Han, *Information-Spectrum Methods in Information Theory*. Berlin, Germany: Springer-Verlag, 2003.
- [19] G. Caire, G. Taricco, and E. Biglieri, "Optimum power control over fading channels," *IEEE Trans. Inf. Theory*, vol. 45, no. 5, pp. 1468–1489, Jul. 1999.
- [20] I. E. Telatar, "Capacity of multi-antenna Gaussian channels," *Eur. Trans. Telecommun.*, vol. 10, pp. 585–595, Nov. 1999.
- [21] W. Yang, G. Durisi, T. Koch, and Y. Polyanskiy, "Quasi-static multiple-antenna fading channels at finite blocklength," *IEEE Trans. Inf. Theory*, vol. 60, no. 7, pp. 4232–4265, Jul. 2014.
- [22] E. MolavianJazi and J. N. Laneman, "On the second-order coding rate of non-ergodic fading channels," in *Proc. Allerton Conf. Commun., Contr., Comput.*, Monticello, IL, USA, Oct. 2013.
- [23] L. Zheng and D. Tse, "Diversity and multiplexing: a fundamental tradeoff in multiple-antenna channels," *IEEE Trans. Inf. Theory*, vol. 49, no. 5, pp. 1073–1096, May 2003.
- [24] P. Elias, S. A. Pawar, P. V. Kumar, and H.-F. Lu, "Explicit space-time codes achieving the diversity-multiplexing gain tradeoff," *IEEE Trans. Inf. Theory*, vol. 52, no. 9, pp. 3869–3884, Sep. 2006.
- [25] L. Zheng and D. Tse, "The diversity-multiplexing tradeoff for non-coherent multiple antenna channels," in *Proc. Allerton Conf. Commun., Contr., Comput.*, Monticello, IL, U.S.A., Oct. 2002, pp. 1011–1020.
- [26] L. Zheng, "Diversity-multiplexing tradeoff: a comprehensive view of multiple antenna systems," Ph.D. dissertation, University of California at Berkeley, Berkeley, CA, Nov. 2002.
- [27] A. Lozano and N. Jindal, "Transmit diversity vs. spatial multiplexing in modern MIMO systems," *IEEE Trans. Wireless Commun.*, vol. 9, no. 1, pp. 186–197, Sep. 2010.
- [28] L. Zheng and D. N. C. Tse, "Communication on the Grassmann manifold: A geometric approach to the noncoherent multiple-antenna channel," *IEEE Trans. Inf. Theory*, vol. 48, no. 2, pp. 359–383, Feb. 2002.
- [29] W. Yang, G. Durisi, T. Koch, and Y. Polyanskiy, "Diversity versus channel knowledge at finite block-length," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Lausanne, Switzerland, Sep. 2012, pp. 572–576.
- [30] J. Östman, W. Yang, G. Durisi, and T. Koch, "Diversity versus multiplexing at finite blocklength," in *Proc. IEEE Int. Symp. Wirel. Comm. Syst. (ISWCS)*, Barcelona, Spain, Aug. 2014.
- [31] G. Durisi, T. Koch, J. Östman, Y. Polyanskiy, and W. Yang, "Short-packet communications over multiple-antenna Rayleigh-fading channels," *IEEE Trans. Commun.*, vol. 64, no. 2, pp. 618–629, Feb. 2016.
- [32] R. Devassy, G. Durisi, J. Östman, W. Yang, T. Eftimov, and Z. Utkovski, "Finite-SNR bounds on the sum-rate capacity of Rayleigh block-fading multiple-access channels with no priori CSI," *IEEE Trans. Commun.*, vol. 63, no. 10, pp. 3621–3632, Oct. 2015.
- [33] S. Alamouti, "A simple transmit diversity technique for wireless communications," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 8, pp. 1451–1458, Oct. 1998.
- [34] A. Feinstein, "Error bounds in noisy channels without memory," *IRE Trans. Inform. Theory*, vol. 1, no. 2, pp. 13–14, Sep. 1955.
- [35] R. G. Gallager, *Information Theory and Reliable Communication*. New York, NY, U.S.A.: Wiley, 1968.
- [36] Y. Polyanskiy, "Channel coding: non-asymptotic fundamental limits," Ph.D. dissertation, Princeton University, Princeton, NJ, Nov. 2010.
- [37] Y. Polyanskiy and S. Verdú, "Scalar coherent fading channel: Dispersion analysis," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Aug. 2011, pp. 2959–2963.
- [38] S. Vituri and M. Feder, "Dispersion of infinite constellations in MIMO fading channels," in *Proc. Conf. of Electrical & Electronics Eng. Israel (IEEEI)*, Eilat, Israel, Nov. 2012.
- [39] A. Collins and Y. Polyanskiy, "Orthogonal designs optimize achievable dispersion for coherent MISO channels," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Honolulu, HI, Jul. 2014.
- [40] J. Hoydis, R. Couillet, and P. Piantanida, "The second-order coding rate of the MIMO Rayleigh block-fading channel," *IEEE Trans. Inf. Theory*, vol. 61, no. 12, pp. 6591–6622, Dec. 2015.
- [41] W. Yang, G. Caire, G. Durisi, and Y. Polyanskiy, "Optimum power control at finite blocklength," *IEEE Trans. Inf. Theory*, vol. 61, no. 9, pp. 4598–4615, Sep. 2015.
- [42] —, "Finite-blocklength channel coding rate under a long-term power constraint," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Honolulu, HI, USA, Jul. 2014.
- [43] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Feedback in the non-asymptotic regime," *IEEE Trans. Inf. Theory*, vol. 57, no. 8, pp. 4903–4925, Aug. 2011.
- [44] T.-Y. Chen, A. R. Williamson, N. Seshadri, and R. D. Wesel, "Feedback communication systems with limitations on incremental redundancy," Sep. 2013. [Online]. Available: <http://arxiv.org/abs/1309.0707>
- [45] I. Tal and A. Vardy, "List decoding of polar codes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Saint Petersburg, Russia, Aug. 2011.
- [46] S. Hassani, K. Alishahi, and R. Urbanke, "Finite-length scaling for polar codes," *IEEE Trans. Inf. Theory*, vol. 60, no. 10, pp. 5875–5898, Oct. 2014.
- [47] M. Mondelli, S. H. Hassani, and R. Urbanke, "Unified scaling of polar codes: Error exponent, scaling exponent, moderate deviations, and error floors," Jan. 2015. [Online]. Available: <http://arxiv.org/abs/1501.02444>
- [48] D. Goldin and D. Burshtein, "Improved bounds on the finite length scaling of polar codes," *IEEE Trans. Inf. Theory*, vol. 60, no. 11, pp. 6966–6978, Nov. 2014.
- [49] S. Maiya, J. Costello, D. J., and T. Fuja, "Low latency coding: Convolutional codes vs. LDPC codes," *IEEE Trans. Commun.*, vol. 60, no. 5, pp. 1215–1225, May 2012.
- [50] P.-N. Chen, H.-Y. Lin, and S. Moser, "Optimal ultrasmall block-codes for binary discrete memoryless channels," *IEEE Trans. Inf. Theory*, vol. 59, no. 11, pp. 7346–7378, Nov. 2013.
- [51] S. Dolinar, D. Divsalar, and F. Pollara, "Code performance as a function of the block size," Jet Propulsion Laboratory, Pasadena, CA, JPL TDA Progress Report 133, 1998.
- [52] R. Koetter and A. Vardy, "Algebraic soft-decision decoding of Reed-Solomon codes," U.S. Patent WO2001035536 A1, 2001.
- [53] M. Fossorier and S. Lin, "Soft-decision decoding of linear block codes based on ordered statistics," *IEEE Trans. Inf. Theory*, vol. 41, no. 5, pp. 1379–1396, Sep. 1995.
- [54] I. Tal and A. Vardy, "List decoding of polar codes," *IEEE Trans. Inf. Theory*, vol. 61, no. 5, pp. 2213–2226, May 2015.
- [55] T. Richardson and R. Urbanke, "Multi-edge type LDPC codes," 2004. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.106.7310>
- [56] A. R. Williamson, T.-Y. Chen, and R. D. Wesel, "Variable-length convolutional coding for short blocklengths with decision feedback," *IEEE Trans. Commun.*, vol. 63, no. 7, pp. 2389–2403, Jul. 2015.
- [57] V. Y. F. Tan, *Asymptotic estimates in information theory with non-vanishing error probabilities*. Foundations and Trends Commun. Inf. Theory, 2014, vol. 11, no. 1–2.

- [58] A. Collins, G. Durisi, V. Kostina, J. Östman, Y. Polyanskiy, I. Tal, and W. Yang, *SPECTRE: short-packet communication toolbox*, v0.2, Dec. 2015. [Online]. Available: <https://github.com/yp-mit/spectre>
- [59] R. G. Gallager, "Basic limits on protocol information in data communication networks," *IEEE Trans. Inf. Theory*, vol. 22, no. 4, pp. 385–398, Jul. 1976.
- [60] A. Ephremides and B. Hajek, "Information theory and communication networks: an unconsummated union," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2416–2434, Oct. 1998.
- [61] B. P. Dunn and J. N. Laneman, "Basic limits on protocol information in slotted communication networks," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Toronto, ON, Canada, Jul. 2008, pp. 2302–2306.
- [62] D. Wang and A. Abouzeid, "On the cost of knowledge of mobility in dynamic networks: An information-theoretic approach," *IEEE Trans. Mobile Comput.*, vol. 11, no. 6, pp. 995–1006, Jun. 2012.
- [63] B. Makki, T. Svensson, and M. Zorzi, "Finite block-length analysis of spectrum sharing networks," in *Proc. IEEE Int. Conf. Commun. (ICC)*, London, UK, Jun. 2015, pp. 7665–7670.
- [64] —, "Finite block-length analysis of spectrum sharing networks: Interference-constrained scenario," *IEEE Wireless Commun. Lett.*, vol. 4, no. 4, pp. 433–436, Aug. 2015.
- [65] N. Farsarakis-Hilliard, P. N. Alevizos, and A. Bletsas, "Coherent detection and channel coding for bistatic scatter radio sensor networking," *IEEE Trans. Commun.*, vol. 63, no. 5, pp. 433–436, May 2015.
- [66] P. N. Alevizos and A. Bletsas, "Noncoherent composite hypothesis testing receivers for extended range bistatic scatter radio WSNs," in *Proc. IEEE Int. Conf. Commun. (ICC)*, London, UK, Jun. 2015, pp. 4448–4453.
- [67] R. Devassy, G. Durisi, P. Popovski, and E. G. Ström, "Finite-blocklength analysis of the ARQ-protocol throughput over the Gaussian collision channel," in *Int. Symp. Commun., Cont., Signal Process. (ISCCSP)*, Athens, Greece, May 2014, invited paper.
- [68] D. P. Bertsekas and R. G. Gallager, *Data Networks*, 2nd ed. Englewood Cliffs, NJ: Prentice Hall, 1992.

Tobias Koch (S'02–M'09) is a Visiting Professor and a Ramón y Cajal Research Fellow at the Department of Signal Theory and Communications, Universidad Carlos III de Madrid (UC3M), Leganés, Spain. He received the M.Sc. degree (with distinction) in electrical engineering in 2004 and the Ph.D. degree in electrical engineering in 2009, both from ETH Zurich, Switzerland. From June 2010 to May 2012, he was a Marie Curie Intra-European Research Fellow with the University of Cambridge, U.K. He was also a Research Intern at Bell Labs, Murray Hill, NJ, USA, in 2004, and the Universitat Pompeu Fabra (UPF), Barcelona, Spain, in 2007. He joined UC3M in 2012. His research interests include digital communication theory and information theory.

Dr. Koch is serving as the Vice Chair of the Spain Chapter of the IEEE Information Theory Society in 2013–2016. He was the recipient of a Ramón y Cajal Research Fellowship, a Marie Curie Intra-European Fellowship, a Marie Curie Career Integration Grant, and a Fellowship for Prospective Researchers from the Swiss National Science Foundation.

Giuseppe Durisi (S'02–M'06–SM'12) received the Laurea degree summa cum laude and the Doctor degree both from Politecnico di Torino, Italy, in 2001 and 2006, respectively. From 2006 to 2010 he was a postdoctoral researcher at ETH Zurich, Zurich, Switzerland. In 2010, he joined Chalmers University of Technology, Gothenburg, Sweden, where he is now associate professor. He is also guest researcher at Ericsson, Sweden.

Dr. Durisi is a senior member of the IEEE. He is the recipient of the 2013 IEEE ComSoc Best Young Researcher Award for the Europe, Middle East, and Africa Region, and is co-author of a paper that won a "student paper award" at the 2012 International Symposium on Information Theory, and of a paper that won the 2013 IEEE Sweden VT-COM-IT joint chapter best student conference paper award. In 2015, he joined the editorial board of the IEEE Transactions on Communications as associate editor. From 2011 to 2014, he served as publications editor for the IEEE Transactions on Information Theory. His research interests are in the areas of communication and information theory.

Petar Popovski (S97–A98–M04–SM10–F'16) is a Professor in Wireless Communications at Aalborg University, Denmark. He received Dipl.-Ing. in electrical engineering (1997) and Magister Ing. in communication engineering (2000) from Sts. Cyril and Methodius University, Skopje, Republic of Macedonia, and Ph.D. from Aalborg University, in 2004. He has more than 240 publications in journals, conference proceedings and books and has more than 25 patents and patent applications. He is a Fellow of IEEE, a holder of a Consolidator Grant from the European Research Council and recipient of the Danish award for Elite Researcher in 2016. In the past he served as an Editor in IEEE Trans. Wireless Communications, IEEE JSAC Cognitive Radio Series and Senior Editor in IEEE Communications Letters. He is currently an Editor for IEEE Transactions on Communications. From 2012 to 2014 he served as the Chair of IEEE ComSoc Emerging Technology Committee on Smart Grid Communications. He is a Steering Committee member for IEEE Internet of Things Journal, as well as Steering committee member of IEEE SmartGridComm. His research interests are in the broad area of wireless communication/networking and communication theory.