

Spectral Versus Energy Efficiency in 6G: Impact of the Receiver Front-End

Abstract—This article puts the spotlight on the receiver front-end (RFE), an integral part of any wireless device that information theory typically idealizes into a mere addition of noise. While this idealization was sound in the past, as operating frequencies, bandwidths, and antenna counts rise, a soaring amount of power is required for the RFE to behave accordingly. Containing this surge in power expenditure exposes a harsher behavior on the part of the RFE (more noise, nonlinearities, and coarse quantization), setting up a tradeoff between the spectral efficiency under such nonidealities and the efficiency in the use of energy by the RFE. With the urge for radically better power consumptions and energy efficiencies in 6G, this emerges as an issue on which information theory can cast light at a fundamental level. More broadly, this article advocates the interest of having information theory embrace the device power consumption in its analyses. In turn, this calls for new models and abstractions such as the ones herein put together for the RFE, and for a more holistic perspective.

Motivation

Energy efficiency was an important driver in the design of 1G and 2G wireless standards where, for instance, more efficient amplifiers were enabled by the adoption of (analog and digital, respectively) signaling formats tolerant of nonlinear amplification. By the time 3G came to be, however, the perception of spectrum scarcity in the low-frequency bands had brought about a shift in priorities, and energy efficiency has since taken a back seat. Indeed, a

chief thrust in 3G, 4G, and 5G, was to improve the spectral efficiency. In the post-5G era, the pendulum is swinging back—not so much because the demand for spectral efficiency is relenting, but because energy efficiency is becoming an imperative. The ICT sector is anticipated to devour a staggering 20% of the global electricity production by 2030 and, homing in on wireless networks, the radio access is the most energy-hungry portion [1]. Meanwhile, at the device end, autonomy and battery life are expected to acquire paramount importance. This mounting pressure on the use of energy brings renewed interest in the trade-off between spectral and energy efficiency, for years skewed all the way toward the former.

At the same time that they regain importance, energy efficiency assessments need to become more holistic. Classically, the only power component that information theory has concerned itself with is the transmit power, but with the progression toward ever higher carrier frequencies, much broader bandwidths, and multiplied antenna counts, the power consumed by the circuitry is bound to become comparable to the transmit power at the infrastructure end [1, Fig. 1] and might outright dwarf it at many devices.

By far the biggest contributor to the circuit power consumption is certain to be the receiver front-end (RFE), whose consumption is swelling already in 5G. For 6G, carrier frequencies could reach 300 GHz, with multigigahertz bandwidths and on the order of 64 antennas at mobile units [1], posing a major challenge. Even for lower frequencies, bandwidths, and antenna numbers, the RFE consumption becomes an issue if high autonomies and/or miniature batteries are desired. As illustrated in Figure 1, in contrast with existing paradigms of high bit rates with a high transmit power or low bit rates with a low transmit power, many 6G devices may be in a new class of their own: high bit rates over short distances with a low transmit power. In such devices, the RFE power consumption might overshadow the transmit power, and lowering the former would enable a powerful regimes of use cases for applications such as smart




Figure 1

Power and rate regions of operation of existing wireless technologies, alongside the new low power, high bit rate regime of potential interest for 6G. An example in this regime is pinpointed based on the models in this article combined with current RF device power versus performance characteristics (see Appendix for details).

wearables, untethered cameras, virtual reality goggles, connectivity modules, or compact short-range access points devoid of cooling systems.

Crucially, any effort to tame the power expended by an RFE pushes it away from ideality, aggravating its imperfections: higher noise floor, nonlinear behavior, and coarser quantization [2]. Characterizing the impact of the RFE on the fundamental performance of a communication system requires models that abstract these essential aspects of noisiness, nonlinearity, and quantization coarseness in a manner that is useful from an information-theoretic vantage, as well as a physically based model for the concomitant power consumption. Coarse quantization has been dealt with extensively already, and important results have emerged [3]. This article expands the modeling scope to also subsume the other aspects.

Armed with a model for the RFE, both the spectral efficiency that it enables (in bits/second/Hz) and its energy efficiency (in bits/second/Watt or bits/Joule) can be characterized, with the latter also expressible through its reciprocal, the energy per bit (in Joules/bit). The spectral efficiency and the RFE energy per bit inform, respectively, of the utilization of the bandwidth and of the energy stored by the device's battery. These two performance measures are informative on their own, yet they are even more revealing when pitted against each other. This amounts to expressing the spectral efficiency, not as a function of the energy expended by the RFE per unit of time, but per unit of information. Indeed, pushing the RFE toward ideality increases its power consumption, but that may




Figure 2

Transmission chain with transmitter, channel, and receiver. From the information-theoretic vantage of encoding and decoding, the RFE can be subsumed into the channel.

be worthwhile if sufficiently more bits can then be sent through. Conversely, relaxing the RFE specifications lowers its consumption, but that need not be fruitful from an energy per bit standpoint, depending on how much the bit rate drops. A spectral-versus-energy efficiency assessment is the appropriate framework to make these determinations, examine the interplay of the RFE key knobs, and glean design guidelines that explicitly account for the energy cost of concealing the RFE's imperfections.

These RFE imperfections, which motivate the analysis in the first place, also complicate it, as they give rise to settings that deviate from the familiar linear channel with additive white Gaussian noise (AWGN). These broader settings are described, relevant works are surveyed, and open issues are identified.

Altogether, this article contemplates the impact of the RFE in potential 6G designs operating at higher carrier frequencies, much wider bandwidths, increased antenna counts, and/or with far lower power consumption limits and energy budgets. Looking beyond, the formulation could be augmented with other expanding contributions to the power consumption (say the channel decoder), for an ever more comprehensive assessment of the cost of operating a receiver. This could then be conceivably blended with its transmitter counterpart. Indeed, in terms of efficiency what matters is the total expended energy [4], [5]; holistic assessments align with 3GPP's energy consumption and efficiency metrics, which account for the total energy irrespective of where it is consumed. For proofs of some of the results and further considerations, readers are referred to [6].

RFE Model

For starters, let us consider the setting in Figure 2, with a single-antenna receiver and bandwidth B . The discrete-time complex baseband transmit signal is x while the noiseless received signal is $r = hx$ for a given complex




Figure 3

RFE modeled as a cascade of four operations: additive noise, saturation, gain, and quantization.

channel gain h . The energies per symbol are $\mathcal{E}_x = \mathbb{E}[|x|^2]$ and $\mathcal{E}_r = \mathbb{E}[|r|^2] = \mathbb{E}[|h|^2]\mathcal{E}_x$.

Linear RFE Model

A wireless receiver consists of two stages: the RFE, which effects the downconversion, filtering, and digitalization, and the baseband processor, which demodulates and decodes. The standard model for the RFE is

$$y = r + z \quad (1)$$

where z is complex Gaussian noise with variance kTF ; here, kT is the minimum theoretical value (-174 dBm/Hz at room temperature) while $F > 1$ is the noise figure quantifying the increase in noise due to RFE nonideality. With that, the signal-to-noise ratio is $\text{SNR} = \text{SNR}_{\text{ideal}}/F$ given

$$\text{SNR}_{\text{ideal}} = \frac{\mathcal{E}_r}{kT}. \quad (2)$$

This AWGN setting has spawned much of the wisdom on the fundamental limits of reliable communication. In particular, the highest achievable bit rate is

$$R = B \log_2(1 + \text{SNR}) \quad (3)$$

attained when x is complex Gaussian. The corresponding spectral efficiency is $C = R/B$.

Having the RFE approach the behavior in (1) is however very costly in terms of power consumption and, as advanced, as one attempts to lower that cost, the behavior becomes less benign. This less benign functioning can be captured by the broader model $y = \Phi(r, z)$, where $\Phi(\cdot)$ is a memoryless, generally nonlinear function.

RFE Operations

The specific form for $\Phi(\cdot)$ propounded here is, as illustrated in Figure 3,

$$\Phi(r, z) = \mathcal{Q}_b(\sqrt{\alpha}S(r + z)) \quad (4)$$

which expresses a cascade of components, each a critical feature of the RFE as detailed next.

➤ **Additive noise:** As in the linear model, z is the noise, with variance kTF .

➤ **Saturation:** The transformation $S(\cdot)$ brings into the model the finite linear range via

$$S(\xi) = \sqrt{\mathcal{E}_{\max}} \phi\left(\frac{|\xi|}{\sqrt{\mathcal{E}_{\max}}}\right) \frac{\xi}{|\xi|} \quad (5)$$

where $\phi(\cdot) \in (0, 1]$ is a function satisfying $\phi(0) = 0$ and $\phi'(0) = 1$. For the sake of analysis, two concrete such functions can be $\phi(\cdot) = \tanh(\cdot)$, which is smooth, and $\phi(\cdot) = \min(\cdot, 1)$, an outright clipping. The transform behaves linearly for small inputs, but it saturates at a signal power of $P_{\max} = \mathcal{E}_{\max} B$.

➤ **Quantization:** The function $\mathcal{Q}_b(\cdot)$ models the analog-to-digital converter (ADC), with b bits per each of the in-phase and quadrature dimensions.

➤ **Gain control:** Critical to proper operation of the ADC is a scaling of its input to prevent overflow or underflow. This is represented by α .

Altogether, the RFE is characterized by the knobs $(F, P_{\max}, b)$. In addition, the performance is influenced by the strategy for setting the gain control.

Power Consumption

There is considerable range in the RFE power consumption for a set of performance attributes, with deviations due to the process technology, form factor, and other considerations. For the purpose here, what is sought is not a model to make precise consumption predictions, but to—based on the scaling laws of circuits and devices—abstract the dependencies on the RFE knobs and on the system parameters, chiefly frequency and bandwidth. The noise figure is determined primarily by the low-noise amplifier (LNA), at whose input the signal is weakest. The most widely used scaling for the LNA power consumption is with $f_c/(F - 1)$, where f_c is the carrier frequency [7]; this is the scaling adopted here, even if [8] suggests that, at very low powers, the scaling may be more aggressive, namely with $f_c^2/(F - 1)^3$. As of the saturation, it occurs in both the LNA and mixer, and the incurred power is reasonably modeled as proportional to $P_{\max}$ [9], [10], [11], [12]. With no single accepted rule for how the noise figure and the saturation power consumptions should be combined [10], [12], this text espouses their addition; this conservative choice is sure to be valid for moderate perturbations around nominal values of F and $P_{\max}$. Then, the two ADCs required to process a complex signal consume a power that scales with Bk^b , where $\kappa \approx 4$ for signal-to-quantization ratios in excess of 40–50 dB while $\kappa \approx 2$ for the lower ratios at which wireless systems operate [13]; $\kappa = 2$ is thus considered in the sequel. Although, beyond a few hundred megahertz, the ADC consumption

Table 1. Figures of Merit Based on Recent Device and Circuit Surveys

Figure of Merit	Units	Value	Remark
γ_{ADC}	fJ/qt	165	Based on ADC data [13]
γ_{NF}	fJ	140	Based on LNA data [8]
γ_{max}	-	5000	Based on mixer data [2], [14], [15]

would become quadratic in the bandwidth, the linear scaling can always be retained through parallelization and ADC pipelining (multiple low-resolution stages cascaded to obtain a higher resolution). Only if it were truly necessary to quantize a single band exceeding hundreds of megahertz would a quadratic behavior be experienced.

The gain control, finally, is typically performed at base-band and consumes negligible power relative to the rest. All in all, the proposed model for the RFE power consumption is

$$P_{\text{RFE}} = \underbrace{\gamma_{\text{NF}} \frac{f_c}{F-1}}_{\text{noise figure}} + \underbrace{\gamma_{\text{max}} P_{\text{max}}}_{\text{saturation}} + \underbrace{\gamma_{\text{ADC}} B 2^b}_{\text{ADC}} \quad (6)$$

with indicative values for the figures of merit γ_{ADC} , γ_{NF} , and γ_{max} , presented in Table 1. In particular, γ_{ADC} improves by about 1.5 dB per year, with a fundamental limit imposed by physics anticipated at $\gamma_{\text{ADC}} \approx 0.1$ fJ/qt [13]. See [6], for more details on the figures of merit.

Interestingly, the power consumption term induced by each knob involves a distinct system parameter:

➤ The power that must be burned to attain a certain noise figure is determined by the carrier frequency, but neither the bandwidth nor the received signal strength. And, rewriting it as

$$\frac{\gamma_{\text{NF}} f_c}{F-1} = \frac{\gamma_{\text{NF}} f_c}{\frac{\text{SNR}_{\text{ideal}}}{\text{SNR}}} \quad (7)$$

it cleanly connects the consumed power, the carrier frequency, and the SNR degradation.

➤ The power that needs to be expended to stretch the range of unsaturated signals depends on the received signal strength, and hence it relates to $\text{SNR}_{\text{ideal}}$.

➤ The power consumed to operate at a certain ADC resolution is contingent on the bandwidth, but not on the carrier frequency or the received signal strength.

Spectral and Energy Efficiency

When a nonlinear transformation sits between encoder and decoder, the noise is compounded by distortion and quantifying the information-theoretic performance limits, as well as strategies that can approach them, becomes notoriously difficult. In such a context, a very useful result can be formulated on the basis of the signal-to-noise-and-distortion ratio

$$\text{SNDR} = \frac{\rho_{xy}^2}{1 - \rho_{xy}^2} \quad (8)$$

where noise and distortion are blended via the squared correlation coefficient

$$\rho_{xy}^2 = \frac{|\mathbb{E}[x^* y]|^2}{\mathcal{E}_x \mathcal{E}_y}. \quad (9)$$

Based on the SNDR, the spectral efficiency with complex Gaussian signalingsatisfies [16]

$$C \geq \log_2(1 + \text{SNDR}). \quad (10)$$

This bound can be obtained from the Bussgang-Rowe decomposition, whereby the output of a nonlinear transformation is split into a scaled version of the input plus uncorrelated distortion (in general neither Gaussian nor independent of the input); the bound arises by then replacing that distortion with Gaussian noise of the same variance. When a unitary transformation is applied between the RFE and the decoder, for instance the time-frequency transformation in multicarrier signaling, the distortion is thrown around and asymptotically (in the dimensions of the unitary transformation, say the number of subcarriers) it is rendered Gaussian; then, the bound gives the actual achievable spectral efficiency [16]. Without a post-RFE transformation, the distortion remains signal-dependent and non-Gaussian; these attributes could conceivably be taken advantage of by the decoder, rendering the lower bound somewhat conservative. Customarily though, the decoder treats the distortion as additional Gaussian noise, whereby its effect (on Gaussian signals) is indeed that of additional Gaussian noise [17], [18]. Motivated by this argument, the lower bound in (10) is regarded as the achievable spectral efficiency in the sequel.

As of the RFE's energy efficiency, it equals R/P_{RFE} . At the infrastructure end, this is often measured in megabits per kilowatthour, reflecting that electricity is billed in

kilowatthour, while at mobile devices bits per joule is a more fitting unit. More germane to information theory is actually its reciprocal, the RFE energy per bit

$$\mathcal{E}_b^{\text{RFE}} = \frac{P_{\text{RFE}}}{R} = \frac{1}{C} \left(\frac{\gamma_{\text{NF}}}{F-1} \frac{f_c}{B} + \gamma_{\text{max}} \mathcal{E}_{\text{max}} + \gamma_{\text{ADC}} 2^b \right) \quad (11)$$

in Joules per bit; this is the measure adopted henceforth. Note that we have used the fact that $P_{\text{max}} = \mathcal{E}_{\text{max}} B$. Importantly, $\mathcal{E}_b^{\text{RFE}}$ relates to the RFE power consumption should not be confused with the transmit and received energy per bit, both are associated with the radiated power [19, Sec. 4.2].

An important implication of (11) is that the dependence of the RFE energy per bit on the bandwidth and carrier frequency is only through the ratio, B/f_c .

A Closer Look at the SNDR

Without loss of generality, the ADC can be designed for a unit-energy input, whereby the required gain control is

$$\alpha = \frac{1}{\mathbb{E}[|S(r+z)|^2]}. \quad (12)$$

With this gain, and given how x , r , and z are related through $\text{SNR}_{\text{ideal}}$, (8) can be manipulated (see [6]) into

$$\text{SNDR} = \frac{\text{SNR}_{\text{ideal}} \rho^2(b, v)}{F + \text{SNR}_{\text{ideal}} (1 - \rho^2(b, v))} \quad (13)$$

where

$$\rho^2(b, v) = \frac{|\mathbb{E}[(r+z)^* y]|^2}{(\mathcal{E}_r + N_0) \mathcal{E}_y} \quad (14)$$

implicitly depends on the function $\phi(\cdot)$ in (5) and must generally be computed numerically. Interestingly, $\rho(b, v)$ depends only the ADC resolution and on the saturation backoff, the latter given by

$$v = \frac{\mathcal{E}_{\text{max}}}{\mathcal{E}_r + N_0}. \quad (15)$$

No matter how high $\text{SNR}_{\text{ideal}}$, the SNDR is curbed, namely

$$\text{SNDR} \leq \frac{\rho^2(b, v)}{1 - \rho^2(b, v)} \quad (16)$$

whose right-hand side is the noiseless signal-to-distortion ratio (SDR). Figure 4 shows this SDR for $\phi(\cdot) = \tanh(\cdot)$, illustrating the effects of the resolution and the saturation backoff.

Only Noise and Quantization

When the backoff is large, the input is always below the saturation level and the distortion is caused only by the ADC. Then, the gain control becomes $\alpha = 1/(\mathcal{E}_r + N_0)$ and (13)




Figure 4

SDR as a function of the ADC resolution and the saturation backoff.

continues to hold, only with $\rho(\cdot)$ a function of solely the resolution, namely

$$\rho^2(b) = \frac{|\mathbb{E}[u^* \mathcal{Q}_b(u)]|^2}{\mathbb{E}[|\mathcal{Q}_b(u)|^2]} \quad (17)$$

where $u \sim \mathcal{N}_{\mathbb{C}}(0, 1)$ represents the gain-controlled signal being fed to the ADC. Since $\rho(b) \in [0, 1]$

$$0 \leq \text{SNDR} \leq \frac{\text{SNR}_{\text{ideal}}}{F} \quad (18)$$

as one would expect, with the highest value corresponding to infinite resolution. At low $\text{SNR}_{\text{ideal}}$

$$\text{SNDR} \approx \text{SNR}_{\text{ideal}} \frac{\rho^2(b)}{F} \quad (19)$$

indicating that the noise figure and the ADC affect the SNDR in a multiplicative fashion. At high $\text{SNR}_{\text{ideal}}$, in turn, the SNDR is bounded by the SDR as per (16).

Optimum Quantization

For a complex Gaussian signal, the quantization distortion is minimized, in the mean-square sense, by a vector quantizer operating over asymptotically long blocks and having itself a complex Gaussian codebook [20]. Then, the distortion is itself Gaussian and independent of the quantized signal, and

$$\mathcal{Q}(u) = \sqrt{1 - 2^{-2b}} \left(\sqrt{1 - 2^{-2b}} u + d \right) \quad (20)$$

with $d \sim \mathcal{N}_{\mathbb{C}}(0, 2^{-2b})$; this relationship evinces a loss in signal energy and the appearance of the quantization distortion d ,




Figure 5

Noiseless SDR as a function of the ADC resolution: Optimum vector quantization versus scalar uniform quantization and its asymptotic approximation in (23).

which in this case does amount to additional Gaussian noise. Note that, with vector quantization over asymptotically long blocks, b embodies the average number of bits per symbol, which need not be integer. The number of bits Mb representing a block of M symbols is to be an integer, but, for $M \rightarrow \infty$, that allows b to take any rational value. From (17) and (20),

$$\rho^2(b) = 1 - 2^{-2b} \quad (21)$$

which plugged into (13) gives the SNDR. Moreover, with the distortion introduced by the optimum vector quantizer being complex Gaussian and independent of the signal, the bound in (10) is then the actual spectral efficiency achievable with complex Gaussian signaling, even without a post-RFE unitary transformation. However, the situation is not akin to an AWGN setting because, for a given resolution, the SNDR is curbed and thus caution must be exercised when applying known results, in particular notions such as the degrees-of-freedom that are inherently asymptotic.

Interestingly, the optimum vector quantizer turns out to be a remarkably faithful representation of the scalar uniform quantizers that are preferred from an implementation standpoint, with b integer-valued. Although, in general, the ensuing $\rho^2(b)$ has to be determined numerically, for growing b [21]

$$\rho^2(b) \approx 1 - cb2^{-2b} \quad (22)$$

for some constant c . Consequently,

$$\text{SNDR} \approx \frac{(1 - cb2^{-2b}) \text{SNR}_{\text{ideal}}}{F + cb2^{-2b} \text{SNR}_{\text{ideal}}} \quad (23)$$




Figure 6

Minimum power consumption over all RFE configurations as a function of the dB-difference between $\text{SNR}_{\text{ideal}}$ and SNDR for $f_c = 28$ GHz, $B = 400$ MHz, and scalar uniform quantization.

Presented in Figure 5 is a comparison between the noiseless SDR with optimum vector quantization and with scalar uniform quantization; also shown is how (23) matches the latter for $b \geq 3$, requiring only the calibration of c . When the backoff is sufficient and the distortion is introduced solely by the ADC, information-theoretic analyses are thereby enabled with only a correction for $b \leq 2$.

Design Guidelines: A Case Study

To make the analysis concrete, consider an exemplary mmWave system with $f_c = 28$ GHz and $B = 400$ MHz, corresponding to $B/f_c = 1/70$. The behavior of the SNDR in this case is presented in Figure 6, which shows, for every value (referenced to $\text{SNR}_{\text{ideal}}$), the minimum power consumption over all combinations of (F, P_{max}, b) . The picture vividly illustrates the soaring power required to bring the SNDR ever closer to $\text{SNR}_{\text{ideal}}$, even in the best possible configuration of the RFE. Also noteworthy is how the power consumption shifts up with $\text{SNR}_{\text{ideal}}$, as better—hence more costly in terms of power—noise figures, saturation levels, and resolutions can be capitalized on. This is a manifestation of the intricate relationship between operating point and RFE knobs that is explored in this section, as an illustration of how the RFE model and the performance measures derived from it can be put to use.

Figure 7 breaks the RFE power consumption down into its three components for the optimum combination of (F, P_{max}, b) at every SNDR. The values increase in discrete steps, reflecting how (F, P_{max}, b) have been discretized to search for the optimum combinations. Save at the highest




Figure 7

Minimum power consumption as a function of the SNDR, parameterized by $\text{SNR}_{\text{ideal}}$, with $f_c = 28$ GHz, $B = 400$ MHz, and scalar uniform quantization.

$\text{SNR}_{\text{ideal}}$, the saturation power is negligible relative to the other components. Some caution is in order concerning this, as more work is needed to finely model the behavior of mixers at very low power levels, and also because the present analysis considers a single link. Additional linear range might have to be budgeted to handle interference, a point that we return to later in the article. Nevertheless, under the premise that the saturation power can indeed be neglected, (11) simplifies to

$$\mathcal{E}_b^{\text{RFE}} = \frac{1}{\log_2(1 + \text{SNDR})} \left(\frac{\gamma_{\text{NF}}}{F-1} \frac{f_c}{B} + \gamma_{\text{ADC}} 2^b \right) \quad (24)$$

which is used in the sequel.

Noise Figure

Also evidenced by Figure 7 is that, when the RFE is optimally configured, the strongest component of the power consumption is virtually always the one associated with the noise figure; physically, this power is invested in

improving the quality of the LNA. Under this premise, it follows from an inspection of (24) that, while relaxing the noise figure worsens the spectral efficiency, it lowers the power consumption even faster, improving the energy per bit. This tradeoff holds up to the point, that depends on the resolution, where the noise-figure power ceases to be dominant; further relaxing the noise figure becomes pointless. This is illustrated in Figure 8, for $f_c/B = 70$ (corresponding for instance to 400 MHz at 28 GHz, or to 2 GHz at 140 GHz), with the SNDR evaluated numerically for 6-bit scalar uniform quantization. The RFE energy per bit is minimized by $F = 2.6$, 3.5, and 6.5 dB, respectively, for $\text{SNR}_{\text{ideal}} = 0$, 10, and 30 dB. Noise figures above these values yield a simultaneously lower spectral efficiency and energy per bit; below these values, a tradeoff unfolds.

Also interesting in Figure 8 is that a higher $\text{SNR}_{\text{ideal}}$ improves both the spectral efficiency and the RFE energy per bit. In contrast, the radiated energy per bit worsens with a higher $\text{SNR}_{\text{ideal}}$ because of the concavity of the bit rate as a function




Figure 8

Spectral efficiency versus RFE energy per bit for $\text{SNR}_{\text{ideal}} = 0, 10, \text{ and } 30 \text{ dB}$, with 6-bit scalar uniform quantization and $f_c/B = 70$. For each $\text{SNR}_{\text{ideal}}$, the noise figure is varied and the value minimizing $\mathcal{E}_b$ is indicated.

thereof [19, Sec. 4.1]. Although beyond the scope of this article, we note that a tension would arise in any holistic optimization of the operating point that involved both transmitter and receiver.

ADC Resolution

Homing in on the ADC, consider first a high $\text{SNR}_{\text{ideal}}$. While increasing the resolution improves the spectral efficiency, the behavior of the RFE energy per bit is more nuanced, and three regimes arise:

- 1) For small and moderate resolutions, the power consumption is dominated by the noise figure term, hence the energy per bit actually shrinks as the resolution grows and more bits are pushed through. In this regime, it is pointless not to increase the resolution.
- 2) At some point, given its exponential dependence on the resolution, the ADC consumption becomes predominant and the energy per bit starts moving north, setting up a tradeoff with the spectral efficiency.
- 3) Eventually, the spectral efficiency ceases to improve as the performance becomes limited by noise rather than quantization. Past this resolution, energy is squandered for no significant improvement in spectral efficiency.

Operation in the first and last regimes is ill-advised, and a well-designed system should target the (rather narrow) intermediate one. This insight would be missed if only the ADC power consumption were accounted for, rather than that of the entire RFE, as the conclusion would then be that the energy per bit is minimized as the resolution is minimized. A more refined analysis reveals that intermediate resolutions are actually preferable. In particular, the resolution that minimizes the energy per bit




Figure 9

Spectral efficiency versus RFE energy per bit for $\text{SNR}_{\text{ideal}} = 0, 10, \text{ and } 30 \text{ dB}$ with $F = 5 \text{ dB}$ and $f_c/B = 70$. For each $\text{SNR}_{\text{ideal}}$, the resolution of the scalar uniform quantizer is varied starting at $b = 1$ and the value that minimizes $\mathcal{E}_b$ is indicated.

is obtainable by solving the transcendental equation that emerges from the condition $(\partial \mathcal{E}_b)/(\partial b) = 0$. For $\text{SNR}_{\text{ideal}} = 30 \text{ dB}$, $f_c/B = 70$, and $F = 5 \text{ dB}$, this returns $b = 4$. Shown in Figure 9 is how operation below this resolution is indeed unwise while operation above $b = 7$ is rather pointless; for $4 \leq b \leq 7$, roughly a four-fold factor in energy per bit can be traded for a roughly 40% increase in spectral efficiency.

At lower $\text{SNR}_{\text{ideal}}$, the same insight holds, only at lower resolutions. For the example in Figure 9 translated to $\text{SNR}_{\text{ideal}} = 10 \text{ dB}$, a 3-bit resolution minimizes the energy per bit. This suggests that an SNR-adaptive resolution would be desirable, and such adaptation is a challenge worth posing to device engineers. Short of that, the resolution must be based on the interval that a system is meant to operate on.

Altogether, the interest in signaling strategies and receiver architectures that are tailored to coarse quantization is well justified, not only to prevent excessive power consumption at the receiver, but for the sake of efficiency. In particular, attention should be paid to medium-resolution converters.

Multiantenna Receivers

Continuing with the case study in the previous section, let us now graduate to multiantenna receivers, progressing from the optimum solution at low SNR (beamforming) to the optimum solution at high SNR (multiple-input multiple-output, or MIMO).

Digital Beamforming

Soaring frequencies and bandwidths sink $\text{SNR}_{\text{ideal}}$, because the omnidirectional pathloss scales with f_c^2 and the noise

power scales with B . Beamforming is the antidote that can bring things back to the operating range of interest without the need to dramatically shrink the communication distance and/or increase the transmit power. Digital beamforming at baseband requires an RFE per antenna, followed by a maximal-ratio combiner. The exact expression for the ensuing SNDR is somewhat involved [22], but, reasonably regarding the distortion at different antennas as uncorrelated, it equals the single-antenna SNDR multiplied by the number of antennas N . With that, (16) becomes

$$\text{SNDR} \leq \frac{N\rho^2(b, v)}{1 - \rho^2(b, v)}. \quad (25)$$

Likewise, the power consumption would in principle be multiplied by N . As mentioned though, the real appeal of beamforming is that, as the SNDR is boosted, proportionally lower $\text{SNR}_{\text{ideal}}$ can be supported. Hence, each constituent RFE can operate at a lower resolution, with a lower power consumption. Weighing the plunging $\text{SNR}_{\text{ideal}}$ against the increasing beamforming gain, and the growing number of RFEs against a lower per-RFE power consumption, looms as a most interesting exercise.

Analog Beamforming

To skirt the need for a full RFE per antenna, some of the beamforming operations can be conducted in the analog domain, and the resulting architectures are indeed widely used in 5G millimeter-wave systems. The signal received by each antenna is passed through an LNA and a phase shifter. The resulting cophased signals are then added and downconverted, and the ensuing single stream is digitized. Proceeding as in the single-antenna case, only with $\mathcal{E}_r$ and $\text{SNR}_{\text{ideal}}$ multiplied by N

$$\text{SNDR} = \frac{N \text{SNR}_{\text{ideal}} \rho^2(b, v)}{F + N \text{SNR}_{\text{ideal}} (1 - \rho^2(b, v))} \quad (26)$$

where, rather than (15), the saturation backoff is now

$$v = \frac{\mathcal{E}_{\text{max}}}{N\mathcal{E}_r + N_0}. \quad (27)$$

Note that

$$\text{SNDR} \leq \frac{\rho^2(b, v)}{1 - \rho^2(b, v)} \quad (28)$$

which, unlike (25), is unaffected by N . It follows that a higher resolution and saturation level are needed for the same noiseless SDR as with digital beamforming. Indeed, as the ADC and mixer are shared by all N channels, a higher resolution and saturation power are to be expected.




Figure 10

SNDR after beamforming as a function of the power consumption with $N = 16$ receive antennas at $f_c = 28$ GHz, $B = 400$ MHz, and $\text{SNR}_{\text{ideal}} = 0$ dB: digital versus analog beamforming.

As of the RFE power consumption, it abides by

$$P_{\text{RFE}} = \gamma_{\text{NF}} \frac{Nf_c G}{F - 1} + \gamma_{\text{max}} P_{\text{max}} + \gamma_{\text{ADC}} B 2^b \quad (29)$$

where G is an additional LNA gain required overcome the insertion loss of the phase shifter. At millimeter-wave frequencies, a typical value is $G = 10$ dB [23]. As only the LNA is replicated N times, only the noise-figure term is affected by the number of antennas, yet this term is now further multiplied by the insertion loss being corrected. And, as mentioned, a higher resolution is needed, meaning that the ADC term is also indirectly enlarged.

A comparison of the power consumption, digital versus analog, is presented in Figure 10, for $N = 16$ antennas and with the configuration of (F, P_{max}, b) optimized at every point. In both cases, having SNDR approach $N \text{SNR}_{\text{ideal}}$ (in this case 12 dB) entails an escalating power consumption that reaches levels unaffordable for many mobile devices; this exposes the challenge of reaping the full potential of beamforming with 16 antennas, let alone with the even higher numbers envisioned for 6G.

An important additional conclusion of this analysis is that the savings in ADC power in the analog structure is more than offset by the increase in LNA power, such that the digital option ends up consuming substantially less power at any given SNDR. This observation is consistent with studies such as [23], yet, as advanced, some caution is in order. As pointed out earlier, the analysis herein considers that the saturation power component can be made arbitrarily low, which would require driving the mixer at very low powers. While designs such as [2] have used very low power mixers

to sacrifice the saturation for power consumption, further work is needed to confirm that such highly aggressive power savings can be realized.

Digital beamforming is also preferable from a performance point of view as it enables communication in multiple directions simultaneously; this enables spatial multiplexing as well as dramatically faster beam search. Also, digital beamforming can mitigate frequency selectivity in the channel, hence its appeal lies beyond the pure spectrum-energy tension. Our derivations suggest that, looking beyond current commercial designs that employ analog beamforming to save power, the benefits of digital beamforming could eventually come with a power benefit.

MIMO

Finally turning to MIMO, which, like digital beamforming, necessitates of a full RFE per antenna, it can be abstracted by the combination of a multiplexing gain and a beamforming gain. There are some caveats concerning the dependence on the channel-state information available at the transmitter, and the analysis could branch accordingly, but the multiplexing gain is robust in that respect. More nuanced is the dependence on the channel matrix itself, whose singular value spread determines the balance between multiplexing gain and beamforming gain. Having already studied the latter, it is sensible to now focus on situations where the former is maximized; this befits either rich multipath channels at lower frequencies or near-field channels at higher frequencies [24]. In these situations, at a fixed $\text{SNR}_{\text{ideal}}$ both the spectral efficiency and the consumed power scale with the number of antennas, and thus the energy per bit is held constant. Alternatively, with a diminishing $\text{SNR}_{\text{ideal}}$, the spectral efficiency can be held constant while the energy per bit shrinks. Both possibilities are decidedly attractive.

1-Bit Communication

Although, as seen earlier, very low resolutions are not the best choice in terms of energy per bit, and are the worst choice in terms of spectral efficiency, the 1-bit case stands apart for two reasons:

- It brings the RFE power consumption down to its minimum, which may be enticing for devices limited, besides battery storage, by sheer power consumption as a proxy for cost, form factor, and heat dissipation.
- This extreme approach, in conjunction with 1-bit resolution also at the transmitter, is implementationally attractive: nonlinear power amplifiers can be tolerated at both ends—with the caveat of spectral regrowth and adjacent channel interference—as only the signs of the in-phase and quadrature components matter.

Note that dropping the quadrature component could be seen as an even further reduction in resolution, but in actuality it would sacrifice half the bandwidth, hence it is information-theoretically sound to preserve both complex dimensions.

With 1-bit resolutions at the two ends of the link, the single-antenna spectral efficiency is [25]

$$C = 2 \left[1 - \mathcal{H}_b(Q(-\sqrt{\text{SNR}})) \right] \quad (30)$$

where $\mathcal{H}_b(\cdot)$ is the binary entropy function and $Q(\cdot)$ is the Gaussian Q-function. Meanwhile,

$$\mathcal{E}_b^{\text{RFE}} = \frac{1}{C} \left(\frac{\gamma_{\text{NF}}}{F-1} \frac{f_c}{B} + \gamma_{\text{ADC}} 2^b \right) \quad (31)$$

where since the spectral efficiency ceases to improve beyond $\text{SNR} \approx 7$ dB, the noise figure can be relaxed. The price of the superior energy per bit is, of course, that $C \leq 2$ b/s/Hz. This is compounded by the difficulties in engineering a link with such strong nonlinearity early in the processing [26]; in that sense, 1-bit communication is a rather uncharted territory, with its own set of challenges.

Beamforming does not remove the 2-b/s/Hz ceiling, although it does lessen the SNR required to attain a certain spectral efficiency. Analog beamforming, in particular, can be straightforwardly subsumed by using N SNR in lieu of SNR. Digital beamforming, alternatively, requires a generalization of (30), depending on the manner in which the signals are combined; from an information-theoretic perspective, the optimum manner is the one that maximizes the mutual information [27].

MIMO, if feasible, is the clearest way to higher spectral efficiencies in a 1-bit architecture, and the RFE count is no different from that of digital beamforming: on full RFE per receive antenna. Again, to the extent that the spectral efficiency scales with the number of antennas, the energy per bit does not worsen, making 1-bit MIMO an attractive combination.

Alternatively to MIMO, there are intriguing information-theoretic results indicating that oversampling can make up for shortages in resolution. This is the case if the continuous-time signal is first hard-limited, discriminating it to two levels, and then sampled. The hard-limited signal is not bandlimited, hence sampling it at the Nyquist rate is suboptimal [28]. For a properly crafted transmit signal distribution, the 2-b/s/Hz ceiling can be raised logarithmically in the oversampling factor [29]. Likewise, if the signal is sampled and subsequently 1-bit quantized, oversampling can be exploited provided the receiver low-pass filter is not set to the signal bandwidth, but expanded by the oversampling factor. Although that increases the noise power, it also ensures that the various samples corresponding to each

symbol are contaminated by independent noise and, again for carefully designed signal distributions, a ceiling above 2-b/s/Hz ceiling can then be attained [30]. Altogether, oversampling can be an interesting alternative to 1-bit communication whenever MIMO is not an option.

No oversampling has been posited throughout this article, and the bandwidth has been identified with the sampling rate. If that is not the case, then the sampling rate should be plugged in lieu of the bandwidth in the RFE power consumption formulas.

Conclusion

As strides are taken toward 6G, energy efficiency is poised to regain its importance alongside spectral efficiency. To provide information-theoretic cover to this development, classical formulations should be augmented with the power consumed by the circuitry. At mobile devices in particular, that consumption is to be dominated by the RFE as we move to higher frequencies, bandwidths, and antenna counts.

Instrumental to any analysis involving the RFE is a model for its power consumption, and this article has set forth such a model. Then, following in the footsteps of other works, the spectral efficiency has been expressed with explicit account of the added noise and distortion introduced by the RFE. The tradeoff between such spectral efficiency and the energy per bit sheds light on the effect of the operating point and the FRE configuration. As a case study, receivers with escalating frequencies and bandwidths have been contemplated, and the gleaned insights include the optimality of intermediate (SNR-dependent) resolutions. The 12-bit or even 14-bit resolutions employed by 4G and 5G become decidedly inadequate, and suitably adapting the resolution to the SNR appears as desirable should the hardware be able to accommodate it. The machinery put forth in the article could be put to use for other situations of interest to 6G. Particularly, multiuser channels with strong near-far conditions as well as adversarial systems would demand a much higher saturation backoff, bringing into the picture additional distortion and/or power consumption, and possibly altering the contrast of digital and analog beamforming. The corresponding analysis would be a welcome development.

In closing, some final comments are in order:

➤ In the face of the non-Gaussian distortion generated by scalar uniform quantizers, and possibly by signal saturation, Gaussian signaling ceases to be optimum. In particular, as far as the quantization distortion is concerned, discrete signal distributions have been shown to be optimum [3]. With optimized discrete constellations, therefore, somewhat better spectral efficiency and energy per bit than the ones considered

in this article are possible, yet that need not be the case with generic constellations. Ongoing research on this matter could augment the framework herein presented.

➤ Interference would exhibit a more favorable behavior than thermal noise in that its spectral density would abate as the bandwidth of the interfering transmitters grows; ultimately, the large-bandwidth regime is inherently not limited by interference, but by noise and distortion.

➤ Fading and the acquisition of channel-state information are problems on their own right, and they would affect the spectral-energy tradeoff as the bandwidth grows unboundedly—provided the scattering were endlessly rich. However, multipath propagation becomes decidedly sparser as the frequency grows into the realm where very large bandwidths are possible, eroding the inherent channel uncertainty and its impact.

A host of research avenues open up as energy efficiency, device power consumption, coarse quantization, distortion, oversampling, and related aspects are injected into information-theoretic formulations. Opportunities to continue probing the boundaries and guiding the evolution of wireless systems, and in particular to keep closing the knowledge gap in terms of how much energy is fundamentally needed to reliably convey one bit of information across a wireless channel.

Appendix A Low-Power, High Bit Rate Use Case

An important implication of the analysis herein is the potential for a high bit rate, short range, but very low power wireless system. As an example, consider the parameters in Table 2, with a relatively low transmit power of 1mW (similar to Bluetooth low energy). The frequency and bandwidth are consistent with private 5G networks in the CBRS band. With these parameters, $\text{SNR}_{\text{ideal}} = 10.3$ dB. Numerically computing $\rho^2(b, v)$ and applying (13), we obtain $\text{SNDR} \approx 6.0$ dB; this is about 4.3 dB below $\text{SNR}_{\text{ideal}}$. To enable a low power digital implementation, assume that only 80% of capacity at 6 dB is attained. Then, $R \approx 160$ Mbps.

What is remarkable is that the consumption can be small. Combining Table 1 with (6), the total RFE power consumption is below 1 mW; see breakdown in Table 2. Of course, this only accounts for the RFE power, to which the digital processing required for filtering, equalization, and decoding, must be added. Nevertheless, this simple calculation supports the possibility of high bit rate applications at very low powers.

Table 2. Theoretical Bit Rate and Power Consumption for a Short Range, Low Power, High Rate Application

Parameter	Value	Remarks
Transmit power	1 mW	
f_c	3.5 GHz	Common in private 5G
B	200 MHz	
Path loss model	3GPP InH	Indoor office
Distance	20 m	
$\text{SNR}_{\text{ideal}}$	10.3 dB	
F	4 dB	
b	4 bits/dim	
ν	30 dB	
Spectral efficiency	$0.8 \log_2(1 + 0.25 \text{ SNDR})$ [bps/Hz]	
Calculated Values for the Bit Rate		
$\text{SNR}_{\text{ideal}}$	10.3 dB	
SNDR	6.0 dB	
R	160 Mbps	
Calculated Values for the RFE Power		
Saturation power	0.003 mW	$\gamma_{\text{max}} P_{\text{max}}$
ADC power	0.52 mW	$\gamma_{\text{ADC}} B 2^b$
NF power	0.32 mW	$\gamma_{\text{NF}} \frac{f_c}{F-1}$
P_{RFE}	0.84 mW	

Acknowledgments

The feedback provided by the guest editorial team and the anonymous reviewers is gratefully acknowledged. The work of Angel Lozano was supported by the Fractus-UPF Chair on Tech Transfer and 6G, and by the ICREA Academia Program. The work of Sundeep Rangan was supported in part by the National Science Foundation under Grant 1952180, Grant 2133662, Grant 2236097, Grant 2148293, and Grant 1925079, along with the industrial affiliates of NYU Wireless.

References

- [1] S. Wesemann, J. Du, and H. Viswanathan, "Energy efficient extreme MIMO: Design goals and directions," *IEEE Commun. Mag.*, early access, 2023, doi: [10.1109/MCOM.004.2200958](https://doi.org/10.1109/MCOM.004.2200958).
- [2] P. Skrimponis et al., "Towards energy efficient mobile wireless receivers above 100 GHz," *IEEE Access*, vol. 9, pp. 20704–20716, 2021.
- [3] J. Singh, O. Dabeer, and U. Madhow, "On the limits of communication with low-precision analog-to-digital conversion at the receiver," *IEEE Trans. Commun.*, vol. 57, no. 12, pp. 3629–3639, Dec. 2009.
- [4] J. N. Murdock and T. S. Rappaport, "Consumption factor and power-efficiency factor: A theory for evaluating the energy efficiency of cascaded communication systems," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 2, pp. 221–236, Feb. 2014.
- [5] O. Kanhere et al., "A power efficiency metric for comparing energy consumption in future wireless networks in the millimeter-wave and terahertz bands," *IEEE Wireless Commun.*, vol. 29, no. 6, pp. 56–63, Dec. 2022.
- [6] A. Lozano and S. Rangan, "Spectral vs. energy efficiency in 6G: Impact of the receiver front-end," *arXiv:2310.02622*, 2023.
- [7] I. Song et al., "A simple figure of merit of RF MOSFET for low-noise amplifier design," *IEEE Electron Device Lett.*, vol. 29, no. 12, pp. 1380–1382, Dec. 2008.
- [8] L. Belostotski and E. A. Klumperink, "Figures of merit for CMOS low-noise amplifiers and estimates for their theoretical limits," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 69, no. 3, pp. 734–738, Mar. 2022.
- [9] R. Brederlow, W. Weber, J. Sauerer, S. Donnay, P. Wambacq, and M. Vertregt, "A mixed-signal design roadmap," *IEEE Des. Test Comput.*, vol. 18, no. 6, pp. 34–46, Nov./Dec. 2001.
- [10] R. Ramzan, F. Zafar, S. Arshad, and Q. Wahab, "Figure of merit for narrowband, wideband, and multiband LNAs," *Int. J. Electron.*, vol. 99, no. 11, pp. 1603–1610, 2012.
- [11] H.-K. Chen, J. Sha, D.-C. Chang, Y.-Z. Juang, and C.-F. Chiu, "Design trade-offs for low-power and high figure-of-merit LNA," in *Proc. Int. Symp. VLSI Design, Autom. Test*, 2006, pp. 1–4.

- [12] A. Verma, P. K. Yadav, S. Ambulker, M. Goswami, and P. K. Misra, "A 36.7 mW, 28GHz receiver frontend using 40 nm RFCMOS technology with improved figure of merit," *Analog Integr. Circuits Signal Process.*, vol. 107, pp. 135–144, 2021.
- [13] B. Murmann, "The race for the extra decibel: A brief review of current ADC performance trajectories," *IEEE Solid-State Circuits Mag.*, vol. 7, no. 3, pp. 58–66, Summer 2015.
- [14] C.-Y. Wang and J.-H. Tsai, "A 51 to 65 GHz low-power bulk-driven mixer using 0.13 μm CMOS technology," *IEEE Microw. Wireless Compon. Lett.*, vol. 19, no. 8, pp. 521–523, Aug. 2009.
- [15] B. Guo, H. Wang, and G. Yang, "A wideband merged CMOS active mixer exploiting noise cancellation and linearity enhancement," *IEEE Trans. Microw. Theory Techn.*, vol. 62, no. 9, pp. 2084–2091, Sep. 2014.
- [16] S. Dutta, A. Khalili, E. Erkip, and S. Rangan, "Capacity bounds for communication systems with quantization and spectral constraints," in *Proc. IEEE Int. Symp. Inform. Theory*, 2020, pp. 2038–2043.
- [17] A. Lapidoth, "Nearest neighbor decoding for additive non-Gaussian noise channels," *IEEE Trans. Inform. Theory*, vol. 42, no. 5, pp. 1520–1529, Sep. 1996.
- [18] A. Lapidoth and P. Narayan, "Reliable communication under channel uncertainty," *IEEE Trans. Inform. Theory*, vol. 44, no. 6, pp. 2148–2177, Oct. 1998.
- [19] R. W. Heath Jr and A. Lozano, *Foundations of MIMO Communication*. Cambridge U.K.: Cambridge Univ. Press, 2018.
- [20] A. Gersho and R. Gray, *Vector Quantization and Signal Compression*, vol. 159. Berlin, Germany: Springer, 2012.
- [21] D. Hui and D. L. Neuhoff, "Asymptotic analysis of optimal fixed-rate uniform scalar quantization," *IEEE Trans. Inform. Theory*, vol. 47, no. 3, pp. 957–977, Mar. 2001.
- [22] A. Khalili, E. Erkip, and S. Rangan, "Quantized MIMO: Channel capacity and spectrospectral power distribution," in *Proc. IEEE Int. Symp. Inform. Theory*, 2022, pp. 2303–2308.
- [23] S. Dutta, C. N. Barati, D. Ramirez, A. Dhananjay, J. F. Buckwalter, and S. Rangan, "A case for digital beamforming at mmWave," *IEEE Trans. Wireless Commun.*, vol. 19, no. 2, pp. 756–770, Feb. 2019.
- [24] H. Do, S. Cho, J. Park, H.-J. Song, N. Lee, and A. Lozano, "Terahertz line-of-sight MIMO communication: Theory and practical challenges," *IEEE Commun. Mag.*, vol. 59, no. 3, pp. 104–109, Mar. 2021.
- [25] A. J. Viterbi and J. K. Omura, *Principles of Digital Communication and Coding*. New York, NY, USA: McGraw-Hill, 1979.
- [26] J. Singh, S. Ponnuru, and U. Madhow, "Multi-gigabit communication: The ADC bottleneck," in *Proc. IEEE Int. Conf. Ubiquitous Wireless Broadband*, 2009, pp. 22–27.

- [27] K. Gao, J. N. Laneman, and B. Hochwald, "Beamforming with multiple one-bit wireless transceivers," in *Proc. Inform. Theory Appl. Workshop*, 2018, pp. 1–9.
- [28] E. N. Gilbert, "Increased information rate by oversampling," *IEEE Trans. Inform. Theory*, vol. 39, no. 6, pp. 1973–1976, Nov. 1993.
- [29] S. Shamai, "Information rates by oversampling the sign of a bandlimited process," *IEEE Trans. Inform. Theory*, vol. 40, no. 4, pp. 1230–1236, Jul. 1994.
- [30] S. Krone and G. Fettweis, "Capacity of communications channels with 1-bit quantization and oversampling at the receiver," in *Proc. 35th IEEE Sarnoff Symp.*, 2012, pp. 1–7.



Angel Lozano (Fellow, IEEE) received the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 1998.

He is a Professor at Universitat Pompeu Fabra, Barcelona. In 1999, he joined Bell Labs (Lucent Technologies, now Nokia), where he was a member of the Wireless Communications Research Department until 2008. He has been

an Area Editor for IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS, an Associate Editor for IEEE TRANSACTIONS ON INFORMATION THEORY, and IEEE TRANSACTIONS ON COMMUNICATIONS. His papers have received several awards, including the 2016 Communications Society & Information Theory Society joint paper award.



Sundeep Rangan (Fellow, IEEE) received the B.A.Sc. degree from the University of Waterloo, ON, Canada, and the M.Sc. and Ph.D. degrees from the University of California, Berkeley, CA, USA, in 1992, 1995, and 1997, respectively, all in electrical engineering.

In 2000, he co-founded (with four others) Flarion Technologies, a spin-off of

Bell Labs, that developed Flash OFDM, a precursor to 4G systems including LTE and WiMAX. In 2006, Flarion was acquired by Qualcomm Technologies where Dr. Rangan was a Senior Director of Engineering. He joined NYU in 2010, where he is now an Associate Director of NYU Wireless, an academic-industry research center.