





Signal-to-Noise Ratio, Dynamic Range, and Power Dissipation

*Paying attention to their interrelation
can greatly benefit analog circuit design*

Yannis Tsividis

Low power dissipation is paramount in an increasing number of applications. These notably include the Internet of Things and wireless sensor networks, in which a battery must last a very long time (even ten years in some cases) or in which energy harvesting is used. Battery life is also a key consideration in more traditional applications such as hearing aids and communication devices. All of these applications involve analog signals and require careful micropower circuit design. Such design can be more efficient if guided by an understanding of the fundamental limits [1]–[4],

Digital Object Identifier 10.1109/MSSC.2018.2867246
Date of publication: 16 November 2018



FOOTAGE FIRM, INC.

beyond which nature does not allow us to go. In this article, we review the relation between power dissipation and signal-to-noise ratio (SNR) for analog circuits. We make a crucial distinction between SNR and usable dynamic range (UDR) and review ways to allow the SNR to vary as needed, thus extending the UDR and resulting in adaptive power dissipation.

To be able to provide some depth, we focus on only conventional, analog continuous-time, single-ended circuits using active elements, resistors, and capacitors. References to discussions of other circuits are provided. The circuits are assumed to be in or near thermal equilibrium. All signals are supposed to be held below a maximum value, beyond which unacceptable distortion occurs. We consider first-order effects only, which allows us to make several simplifying assumptions.

Supply Excess Factors and Power Dissipation

A key to efficient micropower design is to be able to use supply voltage and current as efficiently as possible. For simplicity, signals will be assumed to be sinusoidal. Later in this article, we will use a transconductor as an example, but the definitions apply to other circuits as well. The ideal supply voltage, V_{DD} , would be just equal to the peak-to-peak value of the signal, V_{PP} ; then no voltage would be wasted. The ideal average supply current, $\overline{I_{DD}}$, would be that for class B operation, which is I_P/π , with I_P being the peak output current [5]. To quantify how far we are from those ideal values, we define two excess factors, one for sup-

In this article, we review the relation between power dissipation and signal-to-noise ratio for analog circuits.

ply voltage and the other for supply current, as in Figure 1. We want both excess factors to be as close as possible to 1.

We can now express the supply power dissipation as follows:

$$P = V_{DD} \overline{I_{DD}} = (E_V V_{PP}) \left(E_I \frac{I_P}{\pi} \right). \quad (1)$$

Noise

We use the well-known noise models shown in Figure 2 [5]–[8], where T is the absolute temperature and k is the Boltzmann constant (1.38×10^{-23} J/K). We assume that we are not interested in extremely high frequencies or extremely low temperatures, wherein quantum corrections become necessary. Δf denotes a very small measurement bandwidth, and the mean square noise voltage and current with spectral

components in that bandwidth are denoted by $\Delta \overline{v_n^2}$ and $\Delta \overline{i_n^2}$, respectively. The ratios in Figure 2 approach the corresponding one-sided power spectral densities (PSDs) [9] if Δf is allowed to approach zero. E_N is the so-called excess noise factor because of the extra noise added by the various devices in the transconductor (e.g., 1 to 5).

To find the contribution of a noise source with a known PSD to the output mean square noise, $\overline{v_{on}^2}$, of a circuit, we make use of the transfer function from the noise source to the output, $H_n(f)$. The resulting output mean square noise, denoted by N , is given by [9]

$$N = \overline{v_{on}^2} = \int_0^\infty \text{PSD} \times |H_n(f)|^2 df. \quad (2)$$

N is often called the *noise power* although its units are V^2 , not W. N




FIGURE 1: An active element with supply voltage and current excess factors defined.




FIGURE 2: White noise models for low- and medium-frequency work. (a) and (b) Equivalent models for a resistor and (c) and (d) equivalent models for a transconductor.

Each 3-dB improvement in SNR requires doubling the capacitor area.

includes all noise; for a frequency-selective circuit, noise in both the pass-band and the stopband is included.

Signal-to-Noise Ratio

The so-called signal power, S , is the mean-square value of the output signal voltage v_{os}

$$S = \overline{v_{os}^2}. \quad (3)$$

The SNR is defined as

$$\text{SNR} = \frac{S}{N}. \quad (4)$$




FIGURE 3: A first-order RC low-pass filter.




FIGURE 4: A first-order Gm-C low-pass filter.

Note that this quantity is defined as a ratio of mean square quantities. [The reader is cautioned that in some treatments, root mean square (rms) quantities are used instead.]

The corresponding ratio in decibels is

$$\text{SNR}_{\text{dB}} = 10 \log_{10}(\text{SNR}) \quad (\text{dB}). \quad (5)$$

We stress that, in the above definitions, SNR is a power ratio. Thus, a 3-dB increase in SNR_{dB} means doubling the SNR.

In this article, we will assume sinusoidal signals for simplicity, so that

$$S = \overline{v_{os}^2} = V_{os,\text{rms}}^2 = \left(\frac{V_{pp}/2}{\sqrt{2}} \right)^2 = \frac{V_{pp}^2}{8}. \quad (6)$$

The maximum value of S , denoted by S_{max} , is that beyond which distortion becomes unacceptable.

First-Order Resistance-Capacitance Circuit Example

We now calculate the mean square voltage noise at the output of the resistance-capacitance (RC) low-pass filter in Figure 3. With the input signal set to zero, we use (2) with the PSD from Figure 2(a) and $H_n(f) = (1 + j2\pi fRC)^{-1}$; this gives the value in Figure 3.

Another way to obtain this result is to use the concept of equivalent noise bandwidth [5], which for the above circuit is $1/(4RC)$; when this is multiplied by the PSD of $4kTR$, we again obtain the previously given result. This makes clear why the noise generated by the resistor is independent of the resistance value: increasing the resistance increases the noise PSD but decreases the bandwidth in proportion.

Yet another approach is to equate the capacitor's average energy, $(1/2)C\overline{v_{on}^2}$, to the thermal energy expected for this system in thermal equilibrium from statistical physics, $kT/2$ [10], [11]. This gives again the same result.

The capacitance value required in this circuit to achieve a given SNR can be determined from (4), (6), and the equation for N in Figure 3; the result is shown by the third equation in the same figure. The smaller the signal and the larger the required SNR, the larger the capacitance (and area) needed to keep noise down. Note that each 3-dB improvement in SNR requires doubling the capacitor area (recall that SNR is defined as a power ratio).

First-Order Gm-C Circuit Example

Consider now the first-order low-pass Gm-C circuit in Figure 4 [4]. From the output's point of view, this corresponds to a resistor $1/G_m$ driving C but with noise power multiplied by the excess noise factor. Working as before, we find the noise power given by the second equation in the figure; the equation after that gives the required capacitance for a given SNR. We see that these results are similar to those in Figure 3, only multiplied by the excess noise factor, E_N .

We can determine the power dissipation from the supply as follows. The peak capacitor current, I_p , is given by $(V_{pp}/2) \times (2\pi fC)$. Using in this C from the third equation in Figure 4 and the result in (1), we obtain the power dissipation as given by the last equation in the same figure, where we have assumed that the signal frequency is equal to the cutoff frequency of the filter, f_c . Note that the first three factors in the right-hand side depend on specifications and design prowess; the rest of the factors, though, are there due to the laws of nature. To double the SNR, we need to use twice the capacitance and thus twice the transconductance to leave the corner frequency $G_m/(2\pi C)$ intact. This halving of impedance levels requires twice the current for the same swing, thus doubling the power. Each 3-dB improvement in SNR (meaning, doubling this power ratio) requires doubling the power dissipation and capacitor area. In this respect, life is easier with digital circuits. For them, adding just one more bit improves the




FIGURE 5: The power-to-frequency ratio for an ideal first-order Gm-C circuit [4].

signal-to-quantization-error ratio by 6 dB, yet it does not increase power dissipation by a large factor in multi-bit systems.

Assume now an ideal world in which $E_V = 1$, $E_I = 1$, and $E_N = 1$. Then, from the power equation in Figure 4, we obtain the lowest possible power dissipation per unit of frequency:

$$\frac{P_{\text{lower bound}}}{f} = 8kT(\text{SNR}). \quad (7)$$

This is plotted in Figure 5 [4]. This result, although it cannot be achieved in practice, can serve as a sanity check; if the specifications call for power lower than this bound, we know that this is not possible (under our stated assumptions). The lower bound also allows us to see how much room there is, theoretically, for improvement.

According to the previous equation, the minimum required power can be expected to scale linearly with SNR. Compare this to digital [1]–[4], where the power is logarithmic with SNR [however, analog-to-digital converters (ADCs) and digital-to-analog converters do not allow taking full advantage of this because they are governed by analog limitations].

Highly Selective Bandpass Filter Example

Figure 6 shows a bandpass filter circuit [12], in which we assume that the quality factor, Q , is much larger than 1. Proceeding in the same way,




FIGURE 6: A bandpass Gm-C filter. A quality factor Q much larger than 1 is assumed.

we obtain the results shown next to the circuit.

Although these results are derived analytically [4], [7], [13], [14], intuition can be gained by comparing them to the results in Figure 4. We see that the noise in the second-order circuit is larger by a factor of $2Q$. This noise is dominated by the two large transconductors, for which large internal gains exist from their equivalent input noise voltage to the output. The power dissipation is dominated by the two large transconductors, too. Thus, it is not surprising that the 8 in the power equation in Figure 4 is multiplied by $2 \times 2Q$, resulting in the equation shown [4]. Distortion components can be thought of as disturbances at internal points, which are also amplified by high gains. This

makes small-amplitude signals necessary and, thus, requires smaller amounts of noise, further increasing the required capacitance; power dissipation increases as well because of the resulting large E_V and/or E_I . Similar equations hold for resistive noise in the corresponding active RC filter in Figure 7.

Other Circuits

Detailed calculations give similar results as those already given for a variety of circuits [4], in the general form

$$P = [E_V E_I E_N M] \times kTf(\text{SNR}), \quad (8)$$

where M is a circuit-specific quantity [4], and f is a characteristic frequency depending on the circuit, e.g., a




FIGURE 7: A bandpass RC filter. A quality factor Q much larger than 1 is assumed.

Does this mean that a high dynamic range with low power dissipation is an impossible dream?

band-edge or center frequency for a filter or the noise bandwidth for current mirrors and amplifiers. Results for oscillators are highly dependent on type and assumptions [4], [15]–[17]. A qualitative increase with frequency and SNR is still observed. Switched-capacitor circuits, when well designed, obey similar laws [2], [18].

The product in brackets in (8) depends strongly on the circuit type, specifications, and a number of assumptions [4], of which there are too many to list here for each circuit. The voltage excess factor E_V is large if V_{PP} must be limited to avoid distortion. Similarly, the current excess factor E_I is large for practical

class A circuits; M has a minimum value of 8 (as in Figure 4) for known circuits, and it is much larger for highly selective circuits (as in Figure 6). Thus, the product in brackets can be very large for practical circuits, depending on the details; values of several thousand are not uncommon [4].

Numerical Example

To get an order-of-magnitude idea of what we are talking about, let us consider two Gm-C circuits: the first-order one in Figure 4 and the second-order one in Figure 6. We will assume $V_{DD} = 1$ V, $V_{PP} = 0.5$ V, $E_I = 2\pi$, $E_N = 3$, $T = 300$ K, and $f = 10$ MHz.

Type of Circuit	Q	SNR _{dB}	Total Capacitance	Power Dissipation
First-Order Gm-C	–	30 dB	400 aF	12 nW
		60 dB	400 fF	12 μ W
		90 dB	400 pF	12 mW
Second-Order Gm-C	20	30 dB	32 fF	1 μ W
	20	60 dB	32 pF	1 mW
	20	90 dB	32 nF	1 W

FIGURE 8: A numerical example.




FIGURE 9: SNR in (a) a conventional circuit with constant noise and (b) an adaptive circuit in which the noise is allowed to vary with input level. The desired signal is assumed to be dominant at the input of the filter.

For the second-order circuit, we assume $Q = 20$. Using the equations in Figures 4 and 6, we obtain the results shown in Figure 8. These results exemplify the fact that high SNR, highly selective active filters are impractical.

If inductors can be used, the picture is very different. For a parallel resistor-inductor-capacitor circuit, the noise is still kT/C , but the power dissipation is much lower [4]. If inductors with limited Q are used with active loss cancellation, savings are still possible.

Technology-Dependent Issues

There are several technology-related issues not addressed in the above presented first-order estimates. They include the following:

- **Parasitic capacitances:** These can limit speed; thus, extra currents are needed to drive them.
- **Flicker noise:** Flicker noise can be lowered by using larger gate areas, but this results in larger parasitics, as discussed above. Flicker noise can be addressed through autozeroing/correlated double sampling/chopper stabilization [19], [20] (but with the overhead of a clock).
- **Mismatches:** These can be lowered by using larger dimensions, thus larger parasitics and larger power [21], as discussed above. Mismatches can be addressed through digital calibration, with an ensuing overhead that is sometimes quite significant.
- **Leakage currents:** To be able to ignore leakage currents, bias currents must be kept well above them.
- **Voltage headroom:** Headroom issues become more severe with lower supply voltages [22], [23].

The large power dissipation needed even for a modest SNR, exemplified in Figure 8, appears to be bad news. Does this mean that a high dynamic range with low power dissipation is an impossible dream? The answer is fortunately not, as discussed next.

Maximum Signal-to-Noise Ratio and Usable Dynamic Range

To understand the reason for the large amount of power required for




FIGURE 10: Fixed power versus adaptive power.

conventional circuits with high SNR, we use the following definitions:

$$\text{SNR}_{\max} = \frac{S_{\max}}{N}, \quad (9)$$

$$\text{UDR} = \frac{S_{\max}}{S_{\min}}. \quad (10)$$

These two quantities are not the same. $\text{SNR}_{\max}$ is often called dynamic range, but this can cause confusion and lead to wrong conclusions. We use the above given definition of UDR [24] to avoid the possibility of such confusion. As will be seen, UDR can be very different from $\text{SNR}_{\max}$.

Power Waste in Conventional Circuits

To understand how conventional circuits can waste large amounts of power, consider a simplified example of a filter, in which the desired signal at its input is dominant; the filter is supposed to filter out small amounts of interference and noise in its stopband. Noise and the specification for SNR are assumed to be independent of signal level. Then, the situation is as shown in Figure 9(a). To satisfy the specification, SNR_{SPEC} , for SNR at the minimum signal level, we end up greatly exceeding that specification at the maximum signal level. This can result in a large waste of power dissipation and area, as exemplified by the equations in Figures 4 and 6. The problem can be traced to the fixed nature (time-invariance) of conventional circuits, which is discussed next.

Adaptive Power Dissipation

If the attributes of a circuit are allowed to vary, so that when the signal is large the noise is also allowed to be large, as in Figure 9(b), SNR never needs to be much larger than SNR_{SPEC} , and the power dissipation needed can be drastically lower than that of a fixed circuit; yet the UDR remains the same as in Figure 9(a). (We stress that, for now, we are assuming that the desired signal at the input of the filter is dominant; if there are strong interferers, the situation is more complicated, as will be discussed shortly.) With proper design, one can use the minimum power needed at each point in time, as illustrated in Figure 10. This approach is widely used today, e.g., in smartphones.

Adaptive power dissipation can be accomplished by allowing noise and signal-handling capability to vary depending on signal and interference strengths (for the case of strong interferers, see the “Handling Interferers” section). This can be done by making gain, bias, and impedance levels controllable [24]–[34], as indicated in Figure 11.

Gain adaptation is a classical technique that was first used in telephony and is widely employed in wireless receivers. We will combine it with two power adaptation techniques: bias-current scaling and bias-impedance scaling. Bias and impedance can be varied by adjusting a current, and/or activating/deactivating a number of

unit-cell capacitors, resistors, operational transconductance amplifiers (OTAs), transconductors, and so on, in parallel, as illustrated in Figure 12 [8] (internal nodes may also have to be strapped together). Note that we differentiate between OTAs and transconductors: the former are just operational amplifiers with high output impedance, used closed-loop, with their input driven to a virtual short; the latter can be used open loop and need to remain linear when large signals are applied to their input.

We now present several examples. These examples are idealized; only first-order effects are discussed.

Bias-Current Scaling

A simple example of bias-current scaling is shown in Figure 13 for a class A circuit. Rather than keeping the bias current constant, as in Figure 13(a), we allow it to decrease when the signal is




FIGURE 11: The three adaptive power strategies: adaptive gain, adaptive bias, and adaptive impedance.




FIGURE 12: A transconductor composed of unit transconductors in parallel, with total transconductance G_m . Internal nodes may also have to be strapped together. A smaller transconductance can be implemented by disconnecting some of the units and shutting them down.




FIGURE 13: (a) A current bias and signal in a conventional class A circuit; (b) the corresponding quantities in an adaptively biased circuit; and (c) the log-log plot of current signal, current noise, and current SNR versus the bias current for the case in (b).

small, as in Figure 13(b). Then not only does the power go down, but current noise also decreases. This is shown in the log-log plot in Figure 13(c), where the subscript i indicates current-based quantities.

Another example is shown in Figure 14. The left side of the figure shows an OTA, composed by a number of unit OTA cells in parallel, as was done with transconductors in Figure 12. When the situation allows it, a number of those cells are deactivated [30], [31], resulting in the situation shown on the right side of the figure. The power dissipation decreases n times but so does the maximum current swing. Assuming the transconductance also decreases n times, the mean square noise current at the OTA output also decreases by the same factor [see Figure 2(d)]; this means that the mean square equivalent input noise increases n times [see Figure 2(c)].

Bias-Impedance Scaling

Figure 15 shows a first-order Gm-C filter and its bias-impedance-scaled version. We divide all biases, transconductances, and capacitances by n , e.g., using the approach in Figure 12. Then, all voltage transfer functions

remain unchanged. Voltage swings remain unchanged, and linearity is not affected. The output mean square voltage noise increases n times, and SNR decreases n times. The power dissipation decreases n times.

The corresponding strategy for an active RC circuit using an OTA is shown in Figure 16. The OTA bias and transconductance are divided by n ; all impedances are multiplied by n . Then, all voltage transfer functions remain unchanged. Voltage swings also remain unchanged, and linearity is not affected. The output mean square voltage noise increases n times, and the power dissipation goes down n times. The effect on the SNR is shown at the bottom of the figure (if the scaling is done in discrete steps, then the line for the noise would instead look like a staircase).

Handling Interferers

To put these techniques together, consider a simple wireless receiver architecture as shown in Figure 17. Signal-strength indicators provide information about the size of the desired signal and the out-of-band interferers. The control signals adapt gain (which for simplicity is shown in front of the filter, rather than being distributed),




FIGURE 14: Bias-current scaling for an OTA.




FIGURE 15: Bias-impedance scaling for a Gm-C circuit.

biasing, and impedance levels as needed. Depending on the strength of the in-band signals and those that are out of band (interferers or blockers), different strategies can be adopted, as is now illustrated using the baseband filter as an example. (Similar techniques can be used for other blocks, e.g., mixers and oscillators.)

First, consider the worst case of minimum desired signal and maximum interference at the input of the filter, shown in Figure 18. The horizontal dimension is frequency; the line labeled noise is the noise power in a given, small bandwidth. Thus, this figure is an idealized version of what would be observed with a spectrum analyzer.

The circuit must be designed for maximum signal-handling capability

Adaptive power dissipation can be accomplished by allowing noise and signal-handling capability to vary.

and minimum noise, thus a large SNR. This results in maximum power dissipation. But for any other case, the power dissipation can be decreased. Several cases are now considered [28].

- Small desired signal, small interference (Figure 19). In this case, we can amplify the input and allow the baseband filter noise floor to increase, as shown, via bias-impedance scaling. SNR remains adequate, but it does not need to be made large. Thus, the achievable power dissipation is small. As the signal is varied

over the desired UDR, adaptation makes sure that the SNR remains small, thus keeping the power dissipation small, too. Another possibility is shown in Figure 20; bias scaling is used to reduce the maximum signal handling capability. The power again decreases, at the expense of decreased linearity and decreased (but adequate) SNR.

- Large desired signal, small interference (Figure 21). Here bias-impedance scaling is used to allow the noise to increase relative to the worst-case scenario. The linearity remains the same, and the power goes down. If the signal is now decreased over its UDR, the noise floor is correspondingly decreased to keep the SNR within the specification, thus requiring large power dissipation only when needed.
- Large desired signal, large interference (Figure 22). Again bias-impedance scaling is used to allow the noise to increase. The linearity remains the same, and the power goes down. Again, if the signal decreases over its UDR, the noise floor is decreased to keep the SNR within the specification, thus increasing the power dissipation only when needed.
- Medium desired signal, large interference (Figure 23). This case is trickier and must be handled with care. Here bias scaling is used to reduce the signal-handling capability, and




FIGURE 16: Bias-impedance scaling for an active RC circuit.




FIGURE 17: A wireless receiver with adaptive power control. SSI: signal-strength indicator; LNA: low-noise amplifier.




FIGURE 18: The worst-case scenario: minimum desired signal with maximum interference.




FIGURE 19: Power adaptation in the baseband filter in the case of a small desired signal and small interference.




FIGURE 20: Alternative power adaptation in the case of a small desired signal and small interference.




FIGURE 21: Power adaptation in the case of a large desired signal and small interference.




FIGURE 22: Power adaptation in the case of a large desired signal and large interference.

bias-impedance scaling is used to increase the noise floor. The power goes down, but the linearity decreases, which may result in signal contamination, e.g., by intermodulation products resulting from the blockers. However, as long as the signal is still well above those products, it can be decoded.

This discussion is limited to first-order effects. During design, one must, of course, consider second-order effects including the presence of parasitic capacitances. Still, practical designs are possible and have been well tested. Some are widely used. For example, a careful use of the techniques presented can decrease the average power dissipation of a cell phone receiver by a significant factor.

Avoiding Output Disturbances During Scaling

Gain, bias, and impedance adaptation can be applied during intervals in which no signal service is needed, to avoid disturbances during the transitions, e.g., in between frames in certain wireless applications. However, techniques for applying the discussed techniques even in the presence of signals, without causing disturbances at the output, have been developed [24], [29]–[34].

Conclusions

We have reviewed ways to estimate the power dissipation expected of simple circuits from fundamental considerations. For a given circuit, a lower bound can be determined for power dissipation, which is proportional to frequency and SNR (as a power ratio). Various practical limitations contribute to exceeding this lower bound; this has been quantified by defining three excess factors: one for supply voltage, one for supply current, and one for noise. High-SNR circuits can be prohibitive in terms of power dissipation, especially if they are highly selective. However, in most cases, high SNR is used just to provide a large UDR. In these cases, as we have shown, we can design for a large UDR, while the SNR is only as large as the specifications dictate. Gain, bias,




FIGURE 23: Power adaptation in the case of a medium desired signal and large interference.

and impedance scaling techniques have been reviewed for making high-UDR circuits with low power dissipation possible.

Acknowledgments

The author would like to thank Jacob Baker, Peter Kinget, and the anonymous reviewers for useful comments. The material in this article was developed based on a presentation by the author at the Forum on Energy Efficient Analog Design, IEEE Solid-State Circuits Conference, 2018. The author thanks Prof. Man-Kay Law for his invitation to the forum and his comments on the presentation draft.

References

- [1] B. J. Hosticka, "Performance comparison of analog and digital circuits," *Proc. IEEE*, vol. 73, no. 1, pp. 25–29, 1985.
- [2] R. Castello and P. Gray, "Performance limitations in switched-capacitor filters," *IEEE Trans. Circuits Syst.*, vol. 32, no. 9, pp. 865–876, 1985.
- [3] E. A. Vittoz, "Future of analog in the VLSI environment," in *Proc. IEEE 1990 Int. Symp. Circuits and Systems*, 1990, pp. 1372–1375.
- [4] E. A. Vittoz and Y. P. Tsividis, "Frequency-dynamic range-power," in *Trade-Offs in Analog Circuit Design*, C. Toumazou, G. Moschytz, and B. Gilbert, Eds. New York: Springer, 2002, pp. 283–313.
- [5] P. R. Gray, P. J. Hurst, R. G. Meyer, and S. H. Lewis, *Analysis and Design of Analog Integrated Circuits*. Hoboken, NJ: Wiley, 2008.
- [6] H. Nyquist, "Thermal agitation of electric charge in conductors," *Phys. Rev.*, vol. 32, no. 1, p. 110, 1928.
- [7] D. Blom and J. Voorman, "Noise and dissipation of electronic gyrators," *Philips Res. Rep.*, vol. 26, pp. 103–113, 1971.
- [8] E. A. Klumperink and B. Nauta, "Systematic comparison of HF CMOS transconductors," *IEEE Trans. Circuits Syst. II: Analog Digit. Signal Process.*, vol. 50, no. 10, pp. 728–741, 2003.
- [9] A. Papoulis and S. U. Pillai, *Probability, Random Variables, and Stochastic Processes*. New York: Tata McGraw-Hill Education, 2002.
- [10] N. G. Van Kampen, *Stochastic Processes in Physics and Chemistry*, vol. 1. New York: Elsevier, 1992.
- [11] R. Sarpeshkar, T. Delbruck, and C. A. Mead, "White noise in MOS transistors and resistors," *IEEE Circuits Devices Mag.*, vol. 9, no. 6, pp. 23–29, 1993.
- [12] R. Schaumann, H. Xiao, and V. V. Mac, *Design of Analog Filters*, 2nd ed. New York: Oxford Univ. Press, 2009.
- [13] A. A. Abidi, "Noise in active resonators and the available dynamic range," *IEEE Trans. Circuits Syst. I: Fundam. Theory Appl.*, vol. 39, no. 4, pp. 296–299, 1992.
- [14] G. Groenewold, "The design of high dynamic range continuous-time integratable band-pass filters," *IEEE Trans. Circuits Syst.*, vol. 38, no. 8, pp. 838–852, 1991.
- [15] A. Hajimiri, S. Limotyrakis, and T. H. Lee, "Jitter and phase noise in ring oscillators," *IEEE J. Solid-State Circuits*, vol. 34, no. 6, pp. 790–804, 1999.
- [16] P. Kinget, "Integrated GHz voltage controlled oscillators," in *Analog Circuit Design*. New York: Springer, 1999, pp. 353–381.
- [17] T. H. Lee and A. Hajimiri, "Oscillator phase noise: A tutorial," *IEEE J. Solid-State Circuits*, vol. 35, no. 3, pp. 326–336, 2000.
- [18] B.-S. Song and P. R. Gray, "Switched-capacitor high-Q bandpass filters for IF applications," *IEEE J. Solid-State Circuits*, vol. 21, no. 6, pp. 924–933, 1986.
- [19] C. C. Enz and G. C. Temes, "Circuit techniques for reducing the effects of op-amp imperfections: autozeroing, correlated double sampling, and chopper stabilization," *Proc. IEEE*, vol. 84, no. 11, pp. 1584–1614, 1996.
- [20] G. C. Temes, "Autozeroing and correlated double sampling techniques," in *Analog Circuit Design*. New York: Springer, 1996, pp. 45–64.
- [21] P. R. Kinget, "Device mismatch and tradeoffs in the design of analog circuits," *IEEE J. Solid-State Circuits*, vol. 40, no. 6, pp. 1212–1224, 2005.
- [22] K. Bult, "The effect of technology scaling on power dissipation in analog circuits," in *Analog Circuit Design*. New York: Springer, 2006, pp. 251–294.
- [23] P. R. Kinget, "Scaling analog circuits into deep nanoscale CMOS: Obstacles and ways to overcome them," in *Proc. IEEE 2015 Custom Integrated Circuits Conf.*, 2015, pp. 1–8.
- [24] Y. Tsividis, N. Krishnapura, Y. Palaskas, and L. Toth, "Internally varying analog

circuits minimize power dissipation," *IEEE Circuits Devices Mag.*, vol. 19, no. 1, pp. 63–72, 2003.

- [25] B. B. North, "Apparatus and method for low power operation with high sensitivity in a communications receiver," U.S. Patent 6,081,558, June 27, 2000.
- [26] F. Behbahani, A. Karimi-Sanjaani, W.-G. Tan, A. Roithmeier, J. C. Leete, K. Hoshino, and A. A. Abidi, "Adaptive analog IF signal processor for a wide-band CMOS wireless receiver," *IEEE J. Solid-State Circuits*, vol. 36, no. 8, pp. 1205–1217, 2001.
- [27] S. C. Ciccarelli, R. E. Kaufman, and P. E. Peterzell, "Amplifier with adjustable bias current," U.S. Patent 6,175,279, Jan. 16, 2001.
- [28] Y. Tsividis, "Power dissipation reduction in wireless transceivers," U.S. Patent 7,010,330, Mar. 7, 2006.
- [29] N. Krishnapura and Y. Tsividis, "Dynamically biased 1 MHz low-pass filter with 61 dB peak SNR and 112 dB input range," in *IEEE 2001 Solid-State Circuits Conf. Dig. Tech. Papers*, 2001, pp. 360–361.
- [30] M. Ozgun, Y. Tsividis, and G. Burra, "Dynamically power-optimized channel-select filter for zero-IF GSM," in *Proc. IEEE 2005 Solid-State Circuits Conf., Dig. Tech. Papers*, 2005, pp. 504–613.
- [31] M. T. Ozgun and Y. P. Tsividis, "System and method for dynamic power-optimization of analog active filters," U.S. Patent 7,436,251, Oct. 14, 2008.
- [32] I. Akita, K. Wada, and Y. Tadokoro, "A 0.6-V dynamic biasing filter with 89-dB dynamic range in 0.18- μ m CMOS," *IEEE J. Solid-State Circuits*, vol. 44, no. 10, pp. 2790–2799, 2009.
- [33] Y. Palaskas, Y. Tsividis, V. Prodanov, and V. Boccuzzi, "A 'divide and conquer' technique for implementing wide dynamic range continuous-time filters," *IEEE J. Solid-State Circuits*, vol. 39, no. 2, pp. 297–307, 2004.
- [34] A. Yoshizawa and Y. Tsividis, "A blocker-vigilant channel-select filter with adaptive IIP3 and power dissipation," in *IEEE 2006 Int. Solid-State Circuits Conf. Dig. Technical Papers*, pp. 1850–1859.

About the Author

Yannis Tsividis (tsividis@columbia.edu) received the B.S. degree from the University of Minnesota, Minneapolis, and the M.S. and Ph.D. degrees from the University of California, Berkeley, in 1972, 1973, and 1976, respectively. He is an Edwin Howard Armstrong professor of electrical engineering at Columbia University, New York. He received the 1984 IEEE W.R.G. Baker Award for the best IEEE publication, the 2005 IEEE Undergraduate Teaching Award, and the IEEE Gustav Robert Kirchhoff Award in 2007, and is a corecipient of the Lewis L. Winner Best Paper Award from the IEEE International Solid-State Circuits Conference in 2003. He is a Life Fellow of the IEEE.

SSC