Open Column

By Hanspeter Schmid

Aaargh! I Just Loooove Flicker Noise

My love-hate relationship with flicker noise started with my attempt to explain it to my brother Martin, who is a limnologist[1]. As soon as I got to the "power spectrum inversely proportional to frequency" bit, he said "ah, you mean this, yes, we see the same effect in lakes."

What, in *lakes*? Yes, it is something the limnologists have observed long ago. Kolmogorow explained how it happens: there are turbulences of different scales of size in a lake, and larger-scale turbulences will transfer their energy to smaller-scale turbulences, at the same energy transfer rate all the way from spatially huge turbulences down to the dissipation of the energy into heat on molecular level, or, as physicist Lewis Fry Richardson put it:

> Big whirls have little whirls
> That feed on their velocity;
> And little whirls have lesser whirls,
> And so on to viscosity.

The measured spectrum of water molecule movement then has an energy that is proportional to $1/f^{5/3}$. The exponent is 5/3 because the energy transfer rate is scale invariant; it is independent of the absolute value of that rate. The lower cut-off frequency of that noise, for which some among us have been desperately searching for a long time, is well known in lakes: it is the lowest frequency at which the lake is periodically excited: the Coriolis frequency, the period of which is 12 hours at the pole and approaches infinity at the equator.

Is a MOSFET a Very Small Lake?

Isn't that the material scientific success stories are made of? Two brothers who, by chance, discover that there is something common in their respective research fields and find a way to transfer knowledge [1]? Only that in this case it is not a scientific success story; it just took us two years to find out that, although Yannis Tsividis likes to explain the MOSFET using a hydrodynamic analogy [2], the MOSFET is *still* not a lake. There is, however, a lake of carriers both on the gate and in the channel, and although there are really no turbulences, these lakes of carriers are the cause of flicker noise: they flicker because carriers fall into so-called interface traps, which are oxygen vacancies in the SiO_2 layer [3]. And there are indeed a few things we can learn from the real lakes.

The History of Electronic Flicker Noise

But first we have to understand a bit more about flicker noise. It generally has a power spectral density proportional to $1/f^x$ with $x \approx 1$ and is called flicker noise because if a lamp had this distribution in its light intensity, we would perceive it as flickering. Sometimes this noise is also called pink noise, because light with such a spectrum looks pink to our eyes.

In our domain it was first observed in vacuum tube circuits, for which Bernamont modelled it in 1937. The first models for flicker noise in semiconductors were made by McWhorter already in 1957. The mathematician Benoit Mandelbrot, famous through the beautiful Mandelbrot fractals, made a few observations on the fractal (or self-similar) nature of flicker noise in several publications between 1967 and 1969.

Lagging behind this theory, flicker noise became a really big topic for IC designers only with the advent of

For more information, contact the editors:

Prof. Joseph Chang
(ejschang@ntu.edu.sg)

Prof. Robert Newcomb
(newcomb@glue.umd.edu)

Prof. Stephen Wong
(stephen_wong@hms.harvard.edu)

[1]Limnology is a discipline that concerns the study of in-land waters (both saline and fresh), specifically lakes, ponds and rivers (both natural and man-made), including their biological, physical, chemical, and hydrological aspects.

 1531-636X/07/$25.00©2007 IEEE

CMOS amplifiers, but almost from the start, people asked about the lower cut-off frequency of this noise, as it was mathematically meaningless to assume that there is no such lower cut-off. Caloyannides actually managed to measure MOSFET flicker noise down to $10^{-6.3}$ s (which is one cycle in three weeks) in 1974, but could not find an end to it. By the time Keshner wrote his famous, partly philosophical overview paper in 1982 [4], there were already abundant occurrances of flicker noise in very different systems: the voltage and current of vacuum tubes, diodes and MOSFETs; the resistance of carbon microphones, semiconductors, thin films and ionic solutions; the frequency of quartz oscillators; the average seasonal temperature; the annual amount of rain fall; the rate of traffic flow; the voltages across nerve cell membranes; the rate of insulin uptake by diabetics; economic data; the loudness and pitch of music; and, of course, water molecule movement in lakes. In fact, flicker noise does not seem to be the exception, but rather the rule.

Flicker noise generally has a power spectral density proportional to $1/f^x$ with $x \approx 1$ and is called flicker noise because if a lamp had this distribution in its light intensity, we would perceive it as flickering.

The Paradox Nature of Flicker Noise

Thinking about flicker is sometimes very disturbing. Take the lower cut-off frequency I mentioned above: if we take the integral of $1/f$ noise from a certain frequency down to zero, we end up with an infinite power. An infinite power over a finite frequency band? Not nice. There are two related problems we actually face here. One is that, as much as we would like it to be otherwise, the duality of time and frequency is only a mathematical concept, but when we get down to it, time is real and frequency is not [5]. We love to think about band-limited signals; our whole communication world bases on them, but when we look at it closely, we find no such signals. My mobile phone did not exist five years ago, and will probably not be around anymore in ten years, so whatever I do with it is strictly limited in time and, therefore, by Fourier theory, *cannot* have any band limited signals. The signals are sufficiently band limited for practical purposes, though. So while this does not normally have to worry us when we deal with stationary processes, it leads us to the infinite-power problem when we look at $1/f$ noise, because this noise is not stationary [4]. It is therefore much more productive, intellectually, to look at the time domain.

The time domain already calms down our nerves regarding the apparent paradox of infinite power over a finite frequency band: after all it takes an eternity for something to happen at zero frequency, and the possibility of having infinite power within an eternity should not surprise us. There is more, though. The inverse Fourier transform of a power spectral density (PSD) is the autocorrelation function (ACF), and in the case of a PSD proportional to $1/f$, the ACF is constant [4]. This actually means: what happened one microsecond ago has precisely the same influence on the present state as what happened one week ago. Flicker noise comes from a process with time-independent memory. And, even more importantly, any process with time-independent memory or long-time memory flickers!

World-Wide Web: A Cooool Example for Flicker Noise

I can illustrate this quite well with another system you know that also has time-independent memory: the world-wide web. In many cases you will not be interested whether something was written yesterday or one year ago, particularly when it is a trivial matter or a joke you are reading about. One example is the level of annoyedness as expressed by the number of letters ‘a’ in the word “argh”. You may write “aaargh”, and then I feel that I am more annoyed than you and write “aaaaaargh”, which again may inspire someone else to use even more ‘a’s. As with flicker noise, this process is not time-dependent: I just express that I am more annoyed than the person whose text I read last, irrespective of when that text was written. The process is also not stationary, because the web is getting larger all the time. A simple Perl script can use Yahoo[2] to search for the word “argh” with a certain number of ‘a’s, and while we’re at it, also for the words “love” and “cool” with a specific number of ‘o’s. Figure 1 shows the frequency (i.e., the number of pages containing that word) on the x-axis and the power of the word (i.e., the number of vowels) on the y-axis. The red line is “love”, the green line “argh”, the blue line “cool”, and the dotted magenta line is a plain $1/f^x$ behavior. While the magnitude is not the same for “argh”, “love” and “cool”, the exponent is the same, an observation that extends to several four-letter words with ‘i’s and ‘u’s that I will not discuss here. As with lakes, the exponent of flicker noise depends on the fundamental operation principle of a system, and not on the actual content. It is the same as for MOSFETs, where the flicker noise magnitude depends very much on how clean the SiO_2 interface can be made,

[2]The same can be done with Google, of course, but there all curves show a very curious anomaly, which must be due to a problem in the search algorithm used by Google. Additional information is available from the author on request.

but the exponent is independent of this. Whether we can now conclude from Figure 1 that the web is quite cool and that people express love more often than annoyedness is material for another column.

The Incredibly Long Memory of a SiO_2 Interface

So flicker noise is all about memory. How long then is the memory of a MOSFET? Bloom writes in [3] that *the long-time memory of the processes that produce* $1/f$ *noise is associated with the long occupation time constants of the interface traps [. . .] which could account for relaxation times distributed between, say,* 10^{-5} *and* 10^8*s*. Do you realize that 10^8 seconds are three *years*? I certainly have never had a MOSFET switched on for so long, but this is what it would take until the flicker-noise process becomes stationary. In theory. Measuring this is impossible, to measure noise one needs approximately ten times the period of the frequency, which in this case would be thirty years.

This brings us back to the question we already have answered for lakes: what is the lowest frequency at which we inject energy into the MOSFET? The surprising answer is: this depends on how long we switch the MOSFET on. Essentially, if you switch a transistor on for 10 seconds, this lower frequency corner is around 0.1 Hz. Any attempt to say *when* the energy is injected at this low frequency is futile, because to inject anything at 0.1 Hz requires the whole time span of 10 seconds. So the question "where is the lower cut-off frequency of flicker noise" does not actually make practical sense; if you attempt to measure lower frequencies, you must switch the system on for a longer time, and therefore you always move the cut-off frequency out of the range of frequencies you can measure. It is a bit like the unreachable case of fairy gold buried near the end of the rainbow. What is more, due to the non-stationary nature of flicker noise, your MOSFET will change it's noise behavior, although this effect is easier to describe than to measure. If a MOSFET whose flicker noise magnitude at 1 Hz is $1\ \mu V/\sqrt{Hz}$ is switched on for three years, then the flicker noise integrated up to 1 MHz still only amounts to 4.5 mV_{RMS}.

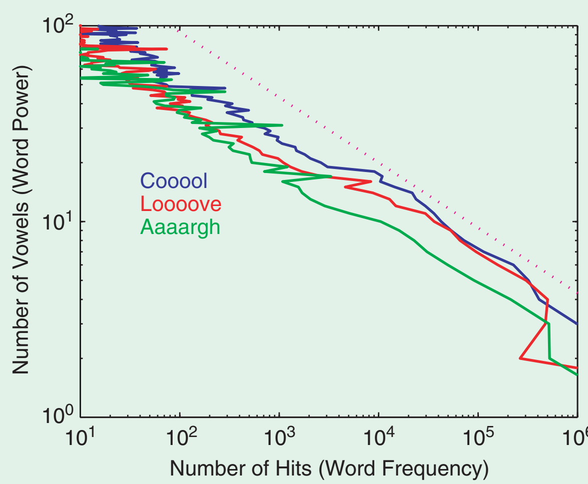

Figure 1. Number of Yahoo hits for a word with a certain number of vowels. The red line is "love", the green line "argh", the blue line "cool", and the blue dotted line is a plain $1/f^x$ behavior.

The Designer's Way to Run Away from Flicker Noise

Nevertheless, flicker noise can be very problematic in different applications; in audio and sensor applications, for example, it contributes directly to noise in the signal band; in RF applications it is often modulated into the phase noise of oscillators, to which I will come back below. So naturally designers have developed different ways to run away from flicker noise.

The easiest method is often used in audio amplifiers (for example in hearing aids); the magnitude of MOSFET flicker noise is inversely proportional to the MOSFET area, so one just makes the MOSFETs extremely wide. Doing so, audio and sensor amplifier designers often find themselves with differential pairs operating in deep weak inversion.

There are other methods, though. Flicker noise has a very long memory, correlates well with itself, so one could regard it as the extension of DC offset to higher frequencies. And indeed, all methods that can be used to compensate offset, such as auto-zeroing, correlated double sampling, or chopping, also reduce flicker noise [6], in the latter case simply by modulating it up to some high frequency where it can safely be ignored.

All of these methods only reduce the effect of flicker noise on the system behavior, though, but there is also one interesting method that reduces the flicker noise at its source: the memory of the MOSFET.

How to Make a FET Forget

Several authors (e.g., [3], [7], [8]) have presented circuits in which they reduce the flicker noise generated by the MOSFET by switching transistors off periodically. This can be done very easily in ring oscillators, for example, and will reduce the phase noise of such an oscillator. If a transistor is switched off, one would expect that the long-term memory of the flicker-noise process is deleted. This is not so.

Eric Klumperink [7] made measurements where he periodically switched the gate of a transistor from the operating point to a voltage low enough to stop the current flowing, and then back. While this removes almost all mobile carriers in the channel, it does not remove all charges that can go in and out of traps. A lot of memory remains, and the power of the flicker noise below the switching frequency is attenuated by only 5 dB. Switching the transistor to zero gate-source-voltage gives better performance, but still the attenuation is only 10 dB. Intuitively, this is not surprising, the built-in potentials of the MOSFET cause trappable carriers to be present at the interface even if no voltages are applied, and these carriers happily jump in and out of SiO_2 interface traps and flicker from the moment the transistor is manufactured. To delete most of the long-term memory, one needs to pull the gate beyond the source and drive the transistor into accumulation, counteracting the built-in potentials [3]. For ring oscillators, this means that if resonance is used such that the switching transients at the inverter gates swing considerably beyond the rails, there will be less phase noise in the oscillator, and the possible gain is more pronounced in modern MOS processes.

Horrifyingly, this also means that building a MOSFET is already sufficient for it to start flickering. It might be interesting to make measurements of a batch of transistors manufactured a long time ago, half of which are used as is, and the other half are first driven into accumulation, and then check whether they behave differently. Intellectually, we all hope to find no difference, but somehow we feel that the behavior of anything electronic does not become more predictable if it is left switched on for a loooooong time. And maybe all this even gives us an idea why Windows computers generally seem to be less stable in the long run than Linux computers: because, by design, Linux has a much weaker memory of what the user has already done, and different applications have independent memories, so less flickering is bound to occur.

Is There Still Something to Find Out for Us?

My love-hate relationship with flicker noise will continue, because there is still a lot we don't know about it and might want to find out. For one, it is still not agreed what the best mathematical model is, and I have not yet found a mathematical model that both includes the non-stationarity of $1/f$ noise and is also useful for designers. On the to-do list are also better ways to calculate and simulate it; things like auto-zeroing, correlated double sampling and other switched circuits are really cumbersome to analyze and to simulate in IC design tools (if it can be done at all). We also should creatively use Klumperink's and Bloom's switching ideas to reduce the effect of flicker noise in many different ways. The fourth missing element is teaching; flicker noise is so omni-present and important, but I yet have to find a university graduate who knows more about it than that it's there (and in most cases not even this). And these are all things that could also be contributed by us, the members of the IEEE CAS society.

I would like to thank my brother Martin, whom I got hooked on flicker noise and who helped me a lot understanding it better, and I really hope that I've got you hooked on it too, because what I like even more than writing about flicker noise is discussing it. Maybe in a coffee break of the next ISCAS? See you there.

References

[1] H. Schmid, "Theory and Practice—Thinking Styles in Engineering and Science," *Australian Journal on Information Systems*, special issue on knowledge management, pp. 106–115, Dec. 2001.

[2] Y. Tsividis, *Mixed analog-digital VLSI devices and technology*, McGraw-Hill, 1996.

[3] I. Bloom and Y. Nemirovsky, "$1/f$ noise reduction of metal-oxide-semiconductor transistors by cycling from inversion to accumulation," *Appl. Phys. Lett.*, vol. 58, no. 15, pp. 1664–1666, Apr. 1991.

[4] M. Keshner, "$1/f$ Noise," *Proc. IEEE*, vol. 70, no. 3, pp. 212–218, Mar. 1982.

[5] D. Slepian, "On Bandwidth," *Proc. IEEE*, vol. 64, no. 3, pp. 292–300, Mar. 1976.

[6] C. Enz and E. Vittoz, "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, Nov. 1996.

[7] E. Klumperink et. al., "Reducing MOSFET $1/f$ Noise and Power Consumption by Switched Biasing," *IEEE J. Solid-State Circuits*, vol. 35, no. 7, pp. 994–1001, July 2000.

[8] S. Gierkink et. al., "Intrinsic $1/f$ Noise Reduction and Its Effect on Phase Noise in CMOS Ring Oscillators," *IEEE J. Solid-State Circuits*, vol. 34, no. 7, pp. 1022–1025, July 2000.

Hanspeter Schmid (h.p.schmid@ieee.org) received the diploma in electrical engineering in 1994, the post-graduate degree in information technologies in 1999, and the degree Doctor of Technical Sciences in 2000, all from the Swiss Federal Institute of Technology (ETH Zürich), Switzerland.

He joined the Signal and Information Processing Laboratory of the ETH Zürich as a teaching assistant in 1994 and worked there as a lecturer and research assistant in the field of analog integrated filters.

From 2000–2005 he was an Analog-IC Designer with Bernafon AG, Switzerland, where he was part of a design team who developed a new mixed-signal IC platform for hearing aids. Now he is a Research Fellow at the Institute of Microelectronics (http://www.ime.technik.fhnw.ch/) at the University of Applied Sciences Northwestern Switzerland, and a senior lecturer at ETH Zürich (Analog Signal Processing and Filtering). His research interest lies in low-power high-speed circuits and systems.

Hanspeter Schmid presently is the secretary and designated chair of the Analog Signal Processing Committee of the IEEE CAS Society, and he serves as a reviewer for several journals and conferences.