SIEMENS

Siemens Corporate Technology | June 2015

A Critical Look at Measurements of the Statistically Independent Components of Ring Oscillator Noise

Markus Dichtl, Pascale Böffgen

First published observations of deviant RO jitter accumulation

SIEMENS

At CryptArchi 2011, Richard Newell gave a talk „Measurement of FPGA ring oscillator noise and analysis using the Allan variance method".
He pointed out very clearly that the accumulated RO jitter does not meet the i.i.d. assumption.

At CryptArchi 2012, Patrick Haddad gave a talk „On the way to monitor random number generation", where he showed graphics of the variance of accumulated jitter quite equivalent to our next slide. He gave quantitative statements about the noise contributions from white and 1/f noise.

At Cryptarchi 2013, Markus Dichtl gave a talk „On Ring Oscillator Based True Random Number Generators and Some of their Variants", where he discussed how the memory effect in ring oscillators might work.

Our observation of jitter accumulation in ROs

Standard deviation of accumulated jitter (unit 50 ps)

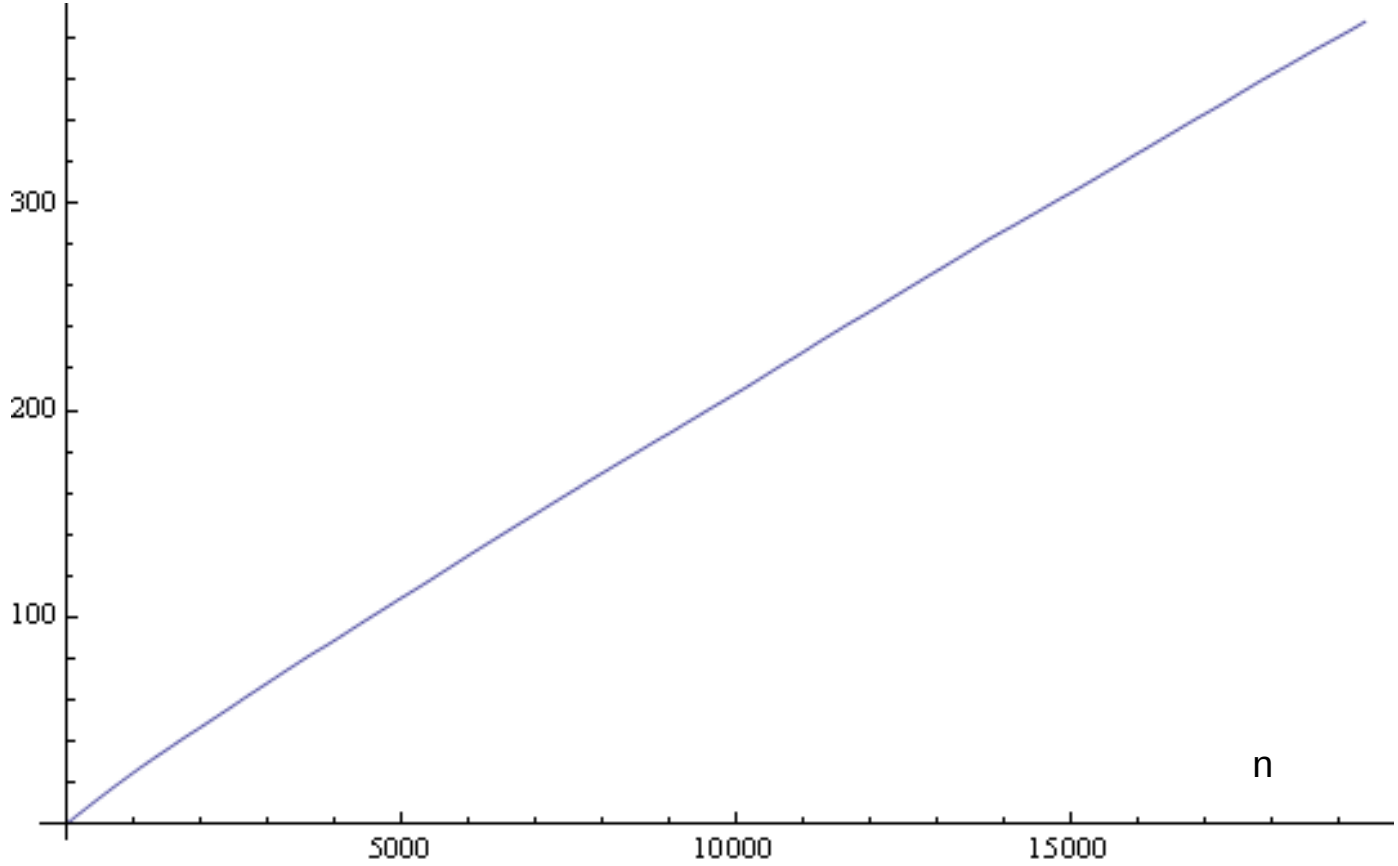

Jitter accumulation over n RO periods, determined from several thousand restarts of a RO of length 13 on Spartan 3

Mathematical background of the deviant jitter behaviour

SIEMENS

Central Limit Theorem

There is no such thing as a single central limit theorem of probability theory, but numerous variants of it.

Many of them imply that (under suitable assumptions) in the limit of $n \to \infty$, the standard deviation of the sum of some random variables $x_1 + x_2 + ... + x_n$ is $\sigma\sqrt{n}$, with σ the standard deviation of the x_i.

Classically, the x_i are assumed to be i.i.d., but these conditions can be weakened to allow non-identical distributions and also some forms of dependence.

Interpreted in ring oscillator terms, the central limit theorem indicates that, as the standard deviation of n accumulated jitter contributions does not increase proportionally to $\sqrt{n}$, but more or less proportionally to n, that the accumulated jitter contributions are dependent.

Variants of the central limit theorem would imply a standard deviation proportional to $\sqrt{n}$ even in the case of dependencies for temporally close jitter contributions only, so the dependencies between jitter contributions must exist for jitter contributions with a larger time period between them.

How to distinguish bad and good noise I

All the difficulties with deviant ring oscillator behaviour seemed to be solved when Haddad, Teglia, Bernard, and Fischer presented their paper
„On the assumption of mutual independence of jitter realizations in P-TRNG stochastic models" at DATE 2014.

Their approach:

The relative jitter of a first RO is measured after a predetermined fixed number N of periods of a second RO.

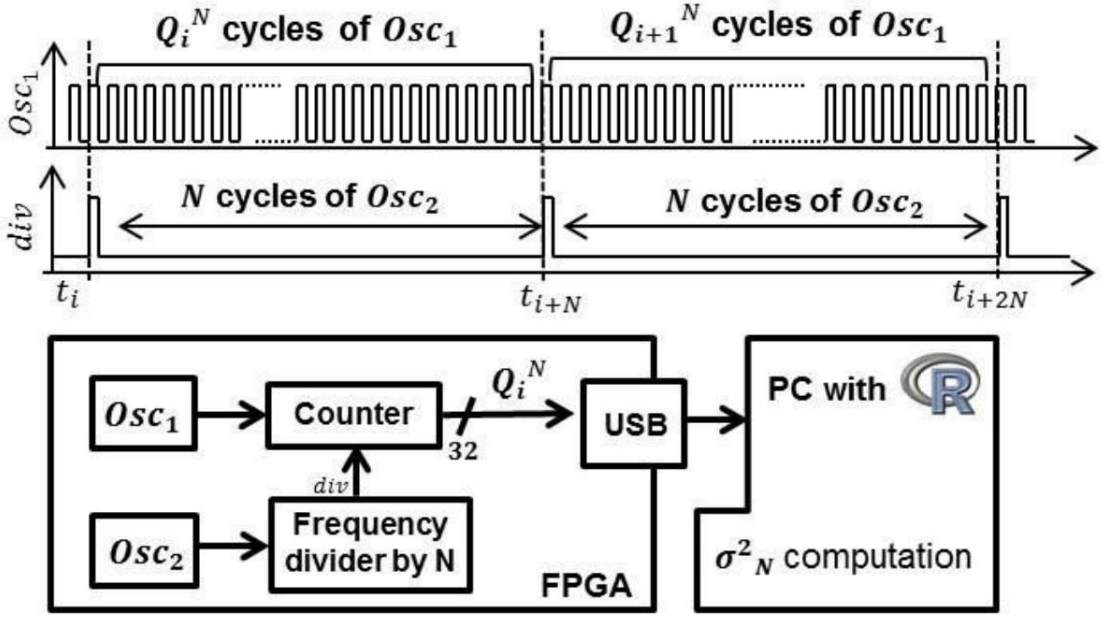

Figure 6 from the Haddad et al. paper

How to distinguish bad and good noise II

In order to eliminate long range memory effects of flicker noise, Haddad et al. consider differences of the jitter contributions of the first RO for two subsequent intervals of N periods of the second RO. This is equivalent to the concept of the Allan-Variance, already suggested by Richard Newell at Cryptarchi 2011.

Haddad et al. derived theoretically that the observed variance should depend on N like $bN + cN^2$, where b and c are determined by fitting the curve to the observed data.

The term proportional to N is the independent thermal noise, the one proportional to N^2 the flicker noise.

The experimental data from Haddad et al.

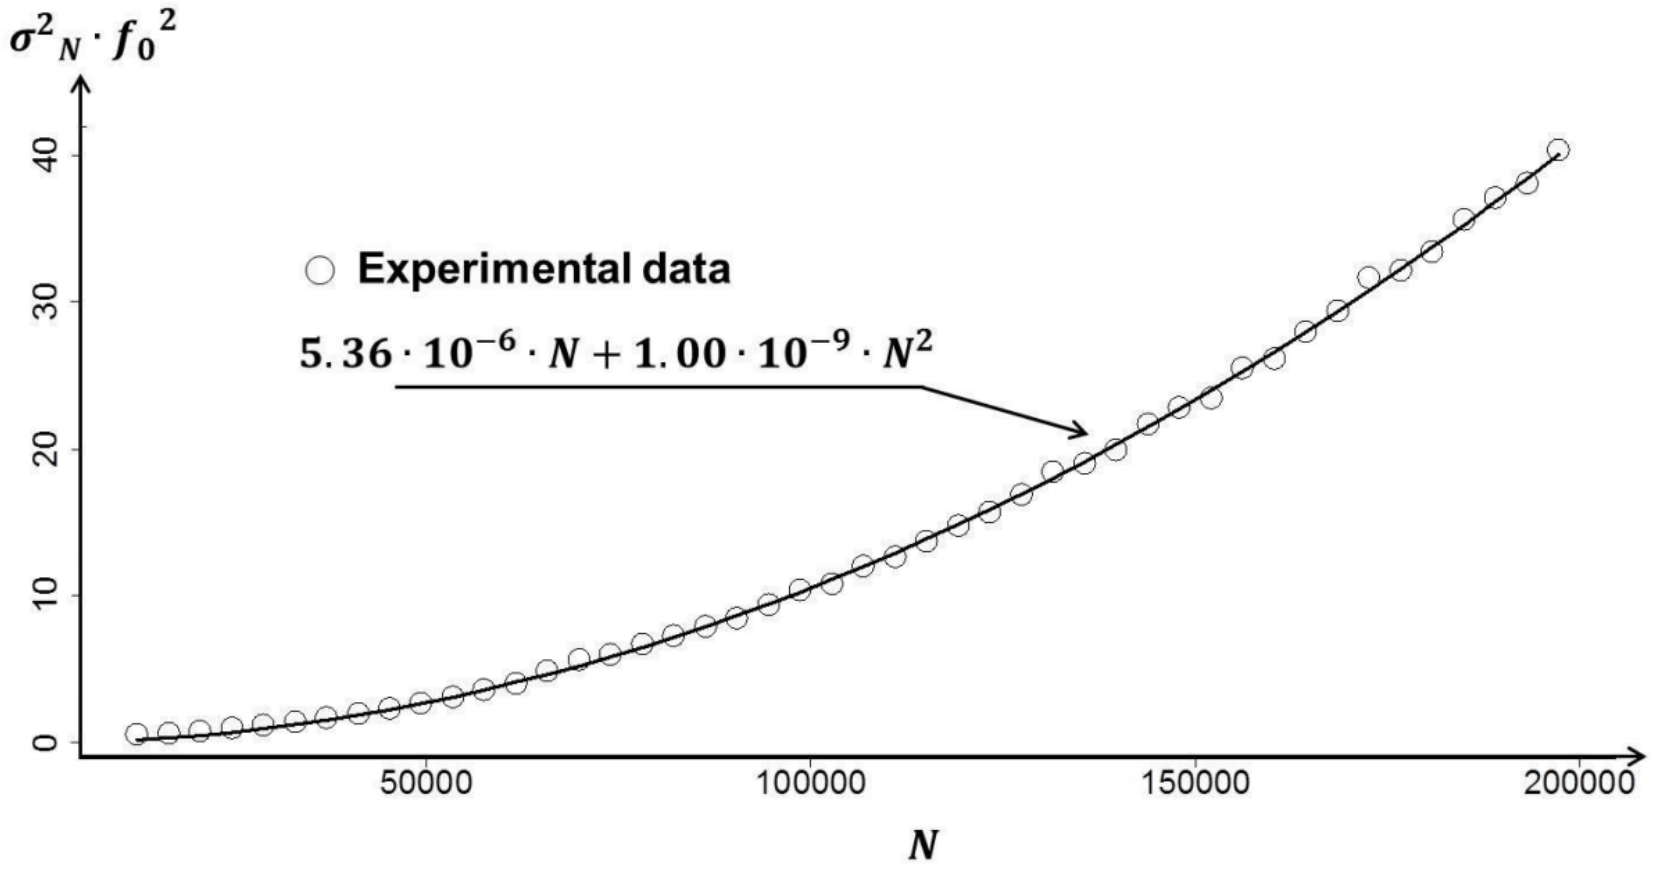

Our own measurements I

- setup almost as in Haddad et al. (see next slide)
- Xilinx Spartan-3 FPGA
- Ring oscillator with 31 inverters
- 2^18 = 262144 values for the calculation of the variances,
  as suggested by Haddad et al.
- N-values (x-coordinates) between 500 and 250000,
  increment of 500 for the measured data
  (N = number of periods of the 2nd ring oscillator before reading out the counter of the 1st ring oscillator)

Our own measurements II

We saw a minor problem with the circuitry suggested in Haddad et al.. The state of the period counter of the one ring oscillator is read asynchronously at a time determined by the oscillations of the other ring oscillator. There is a slight risk that at the time of sampling, the counter is in a transitory state, which could result in metastability or incorrect measurements.

We tried to solve this problem by implementing the counter as a synchronous **Gray code counter** (to be off by one at most) and by using 3-flip-flop synchronization to reduce the risk of metastability.

We tried both the original Haddad et al. design and our Gray code counters, but there were no substantial differences in the results.

Our subsequent results are from the Gray code counter design.

Approximation curve of the form b N + c N^2

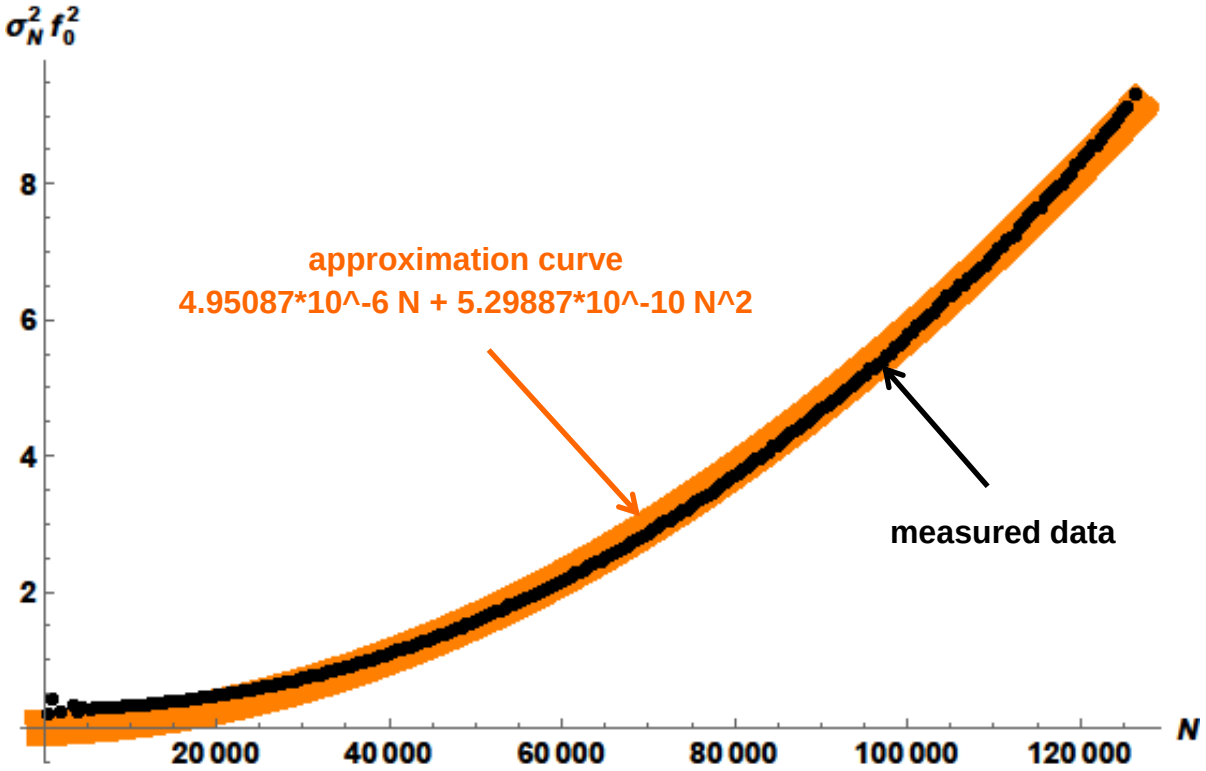

- bad approximation,
  sum of error squares = 2.9515
- # waiting periods = 403969

Details of Haddad‘s Curve

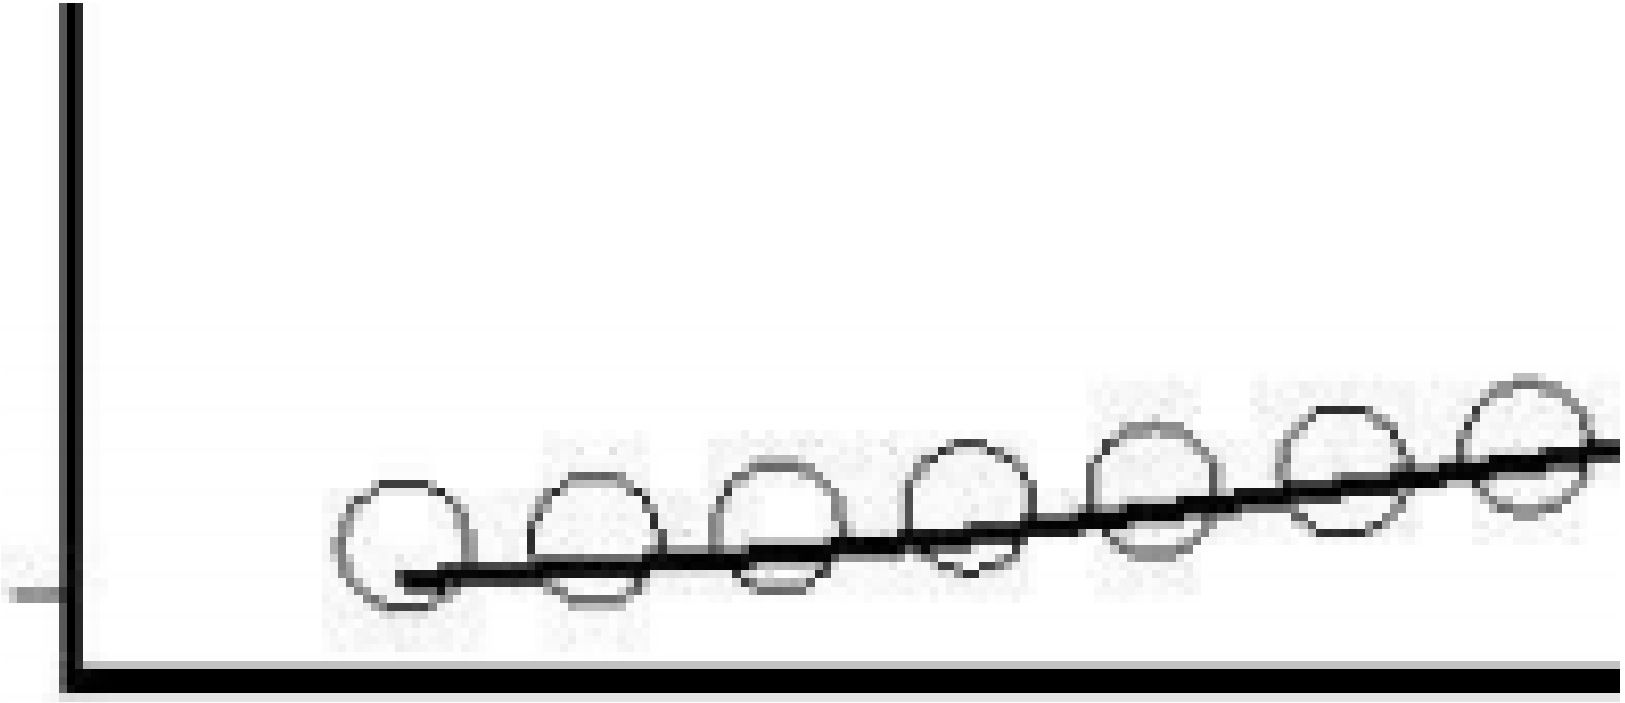

For the lowest values of N, the observed data are systematically above the theoretical curve.

This holds for Haddad‘s curve as well as for ours.

Quantization noise

The systematic errors for low values of N can be explained by quantization noise.

As the jitter of the first RO is only determined in units of an integer number of periods of the second RO, there is a considerable quantization contribution to the variance.

We suggest to include a constant term a to account for the quantization noise:

$$a + b\,N + c\,N^2$$

Approximation curve of the form a + b N + c N^2

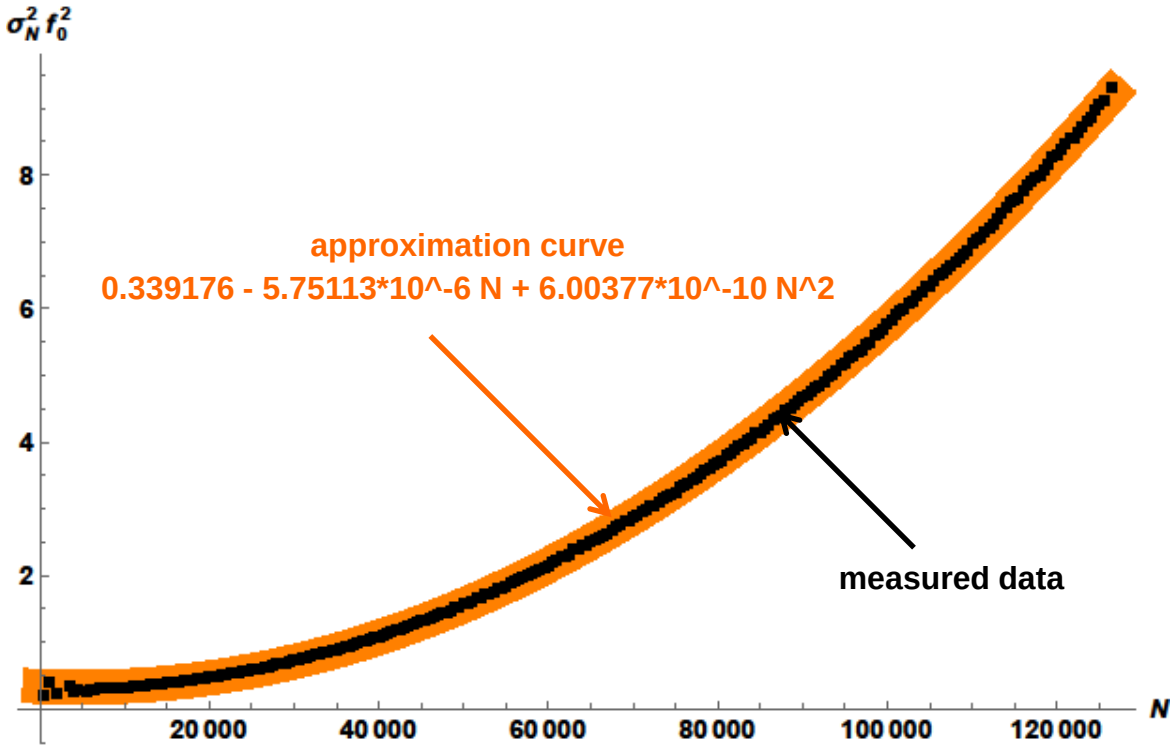

- better approximation,
  sum of error squares = 0.15996
- a represents the quantization noise.
- **but negative b !!**

Have we measured incorrectly?

Now, the negative thermal noise contribution seems so absurd that one may wonder whether something went wrong with the measurements?

We extracted the measurement data of Haddad et al. from the illustration in the paper and fitted a curve of the form
$a + b\,N + c\,N^2$.
Again, **b turned out to be negative**!

(When dropping the constant term for curve fitting, our values for b and c agreed well with the results from Haddad et al.)

Can the curve really tell us about the thermal noise?

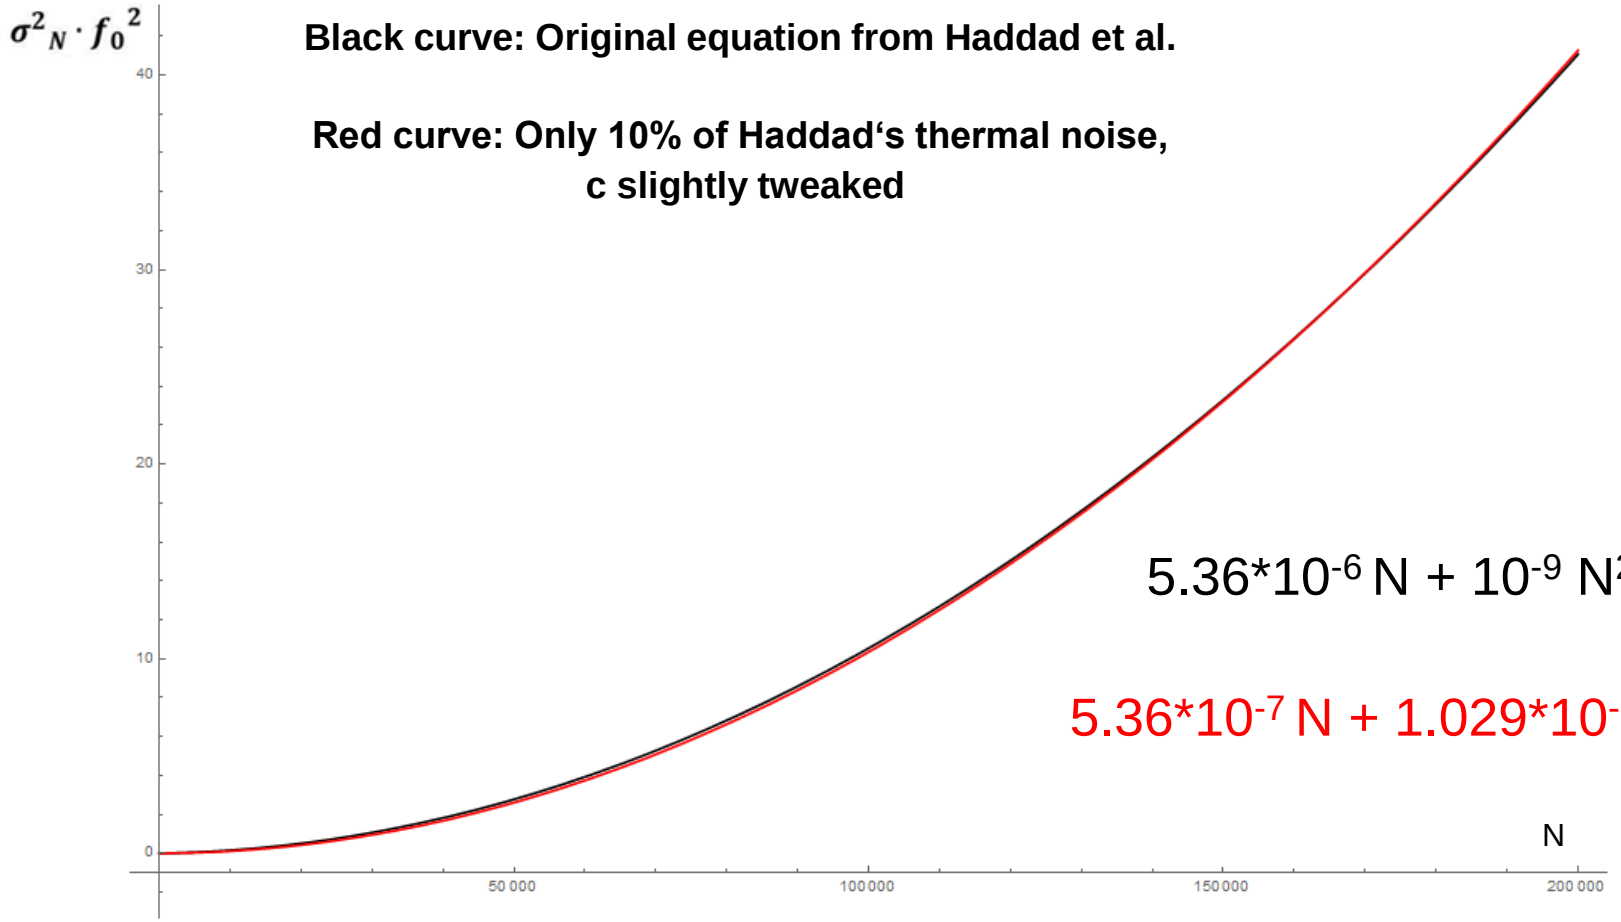

Black curve: Original equation from Haddad et al.

Red curve: Only 10% of Haddad‘s thermal noise, c slightly tweaked

$5.36*10^{-6}\,N + 10^{-9}\,N^2$

$5.36*10^{-7}\,N + 1.029*10^{-9}\,N^2$

What conclusions to draw from the very close two curves

SIEMENS

One can wonder, as the two curves of quite different thermal noise contributions are so close together in the previous illustration, whether it is possible to clearly distinguish the two kinds of noise in the range of values of N considered in the measurements.

Obviously, the smaller the value of N, the stronger the contribution of the thermal noise contribution.

However, we have a dilemma here:

For smaller N, the measured variances are dominated by **quantization noise**, as the approach suggested by Haddad et al. can measure accumulated jitter only in integer multiples of ring oscillator periods.

One way to escape from this dilemma is to change the measurement method in order to be able to determine much smaller jitter contributions.

However, also measurements with an oscilloscope (temporal resolution of 50 ps) and a logic analyzer did not result in consistent data to clearly distinguish between dependent and independent noise contributions.

Is the model wrong?

Quite frustrated from failing again and again, we started to wonder about the validity of the underlying noise model.

Our doubts are supported by the fact that there seems to be no generally accepted theory for the details of the origin of flicker noise in CMOS semiconductors.

In general, 1/f noise in many cases seems to behave not exactly as 1/f, but as $1/f^{\beta}$ with a value of β close to 1.

So, we tried to fit our experimental data with a curve of the form

$$a + b\,N + c\,N^{\alpha}$$

where the value of α is, like a, b, and c, also determined by fitting the curve to the experimental data.

Approximation curve of the form a + b N + c N^α

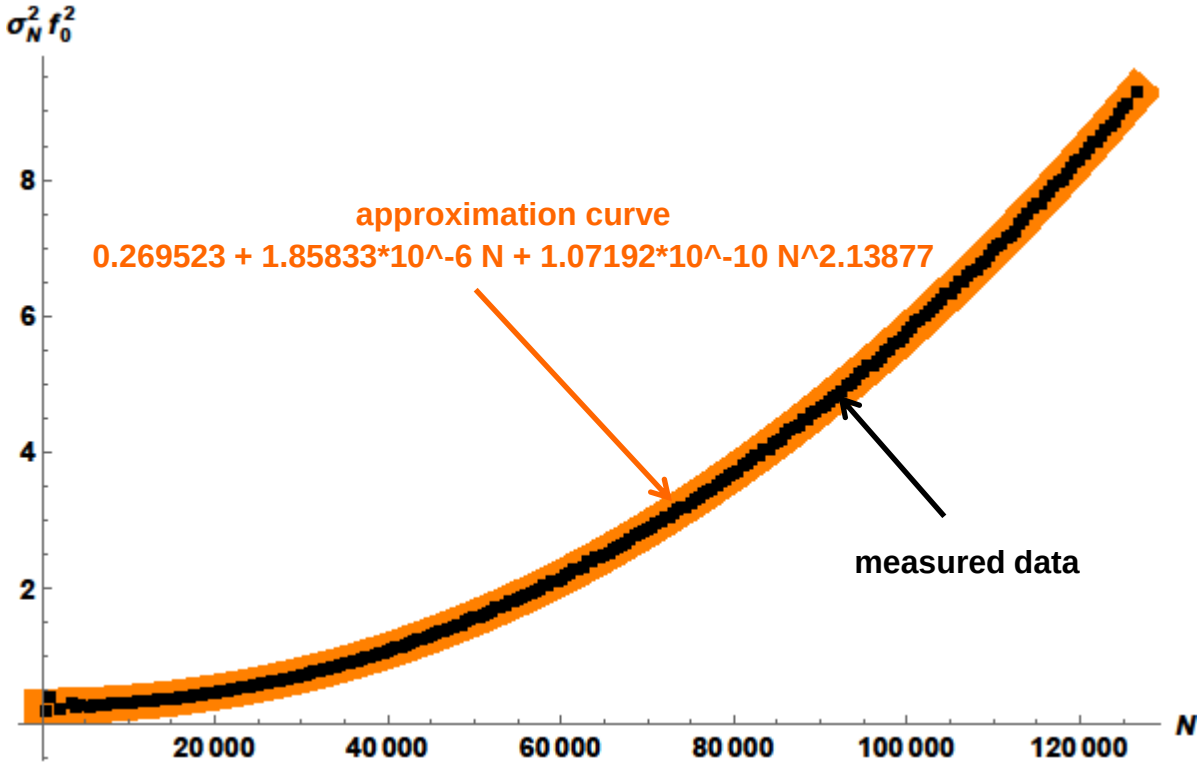

- best approximation,
  sum of error squares = 0,0777706
- positive b
- # waiting periods = 1076235
- How to interpret α > 2?

An excellent fit with the experimental data was achieved.

Approximation curve of the form b N + c N^α

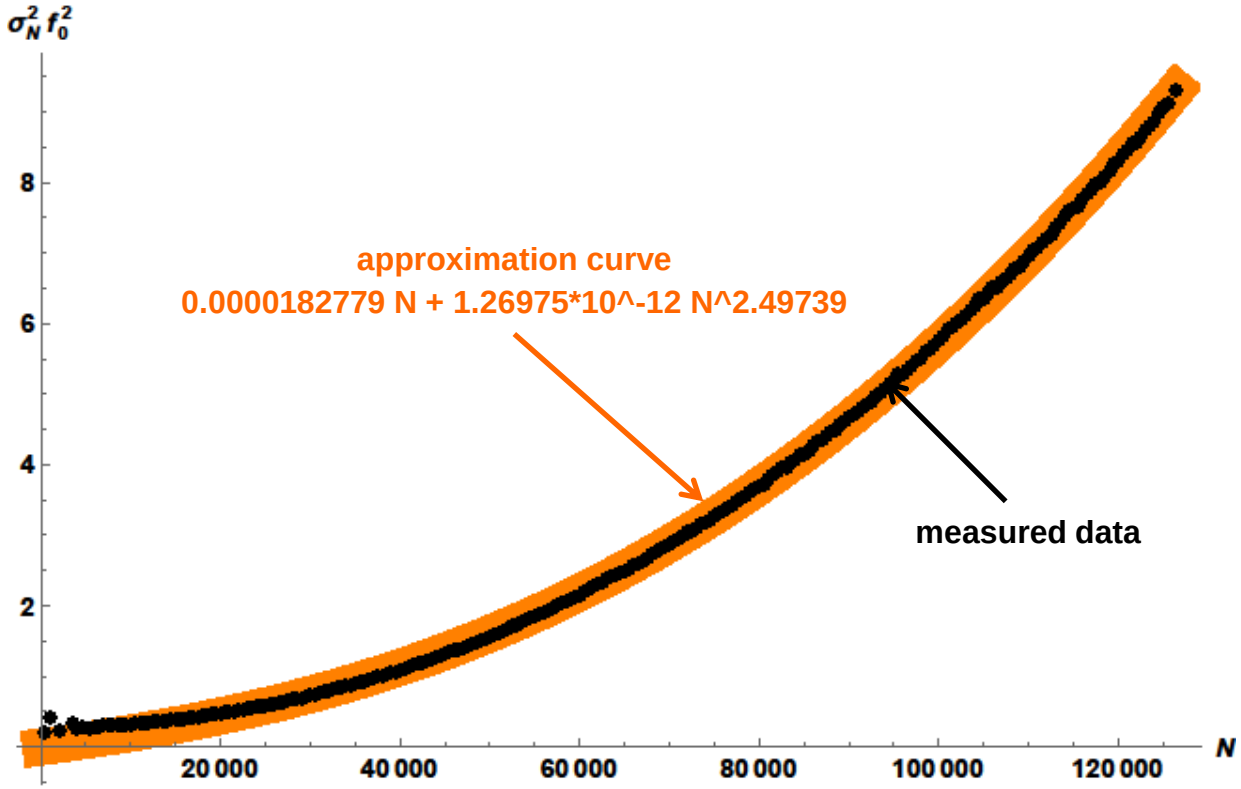

- approximation for small N not good, sum of error squares = 0,929359
- # waiting periods = 109422
- absolute term for quantization noise obviously missing

This just shows that the constant term a is still needed.

Conclusion

- The approach suggested by Haddad et al. to distinguish dependent and independent noise contributions does not work as suggested.
- Further investigating expressions of the form $a + b N + c N^{\alpha}$ and looking for a theoretical justification of α might be worth while.
- TRNGs on FPGAs remain a challenging topic!