An Approach for Entropy Assessment of Ring Oscillator-Based Noise Sources

InfoGard, a UL Company

Dr. Joshua Hill

A Short Introduction

- Randomness is necessary for cryptographic systems:
  - By Kerckhoffs' principle, we assume that adversaries know the design, thus know how secret values are selected.
  - Game theory tells us that in these circumstances, the random selection of parameters yields the least advantage for the attacker.
- We want some model-oriented way to assess the entropy production for noise sources.

Ring Oscillators

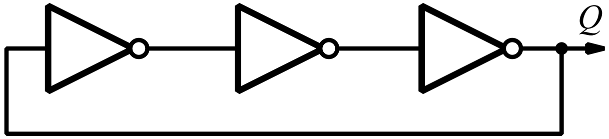

- Take an odd number of inverters (and any number of buffers) and arrange them in a loop. Sample the output of this ring oscillator at frequency f_s.
- Each logic gate has an expected propagation delay. These (along with the number of gates) establish the expected period of the ring oscillator, T.
- Variations from the per-gate expected propagation delay (*jitter*) come from the superposition of a variety of sources: [Fischer 2012]
  - local Gaussian thermal noise
  - local deterministic switching noise
  - global noise sources such as the power supply, electromagnetic emanations, etc.
- Of these sources, local Gaussian thermal noise is the only source that is difficult for an informed active attacker to predict or to have an undue influence on.

Ring Oscillator Statistics

- As the name suggests, “local Gaussian thermal noise” has a normal distribution. [HLL 1999]
- This component of the jitter is expected to be independent from gate to gate.
- The global noise (and some local switching noise) is not independent from gate to gate.
- The sum of independent normally distributed random variables is normally distributed.
- Today, we’ll simplify and track the local Gaussian normal thermal noise induced jitter on a per-cycle basis (This jitter is normally distributed, with mean 0 and variance σ^2).

Ring Oscillator Issues

- Ring oscillators are commonly very temperature and source voltage sensitive.
- This local Gaussian thermal noise may not dominate the observed jitter.
- Sometimes the structure of the deterministic switching noise and global noise produce undesirable emergent behavior:
  - Mean seeking
  - Mutual locking (“entrainment”)

Rough Draft

- There's clearly uncertainty in there somewhere, so we arrive at our first attempt.
  1. Get some ring oscillators.
  2. Assess their entropy production via statistical tests.
  3. Entropy.
- There are problems with this approach.
  - We only noted that there is uncertainty; we didn't provide a min-entropy bound.
  - We didn't resolve how to determine which part of the jitter is due to this local Gaussian thermal noise.
  - Most statistical testing doesn't work here…

Statistical Testing of Ring Oscillators

- Oscillators may have output that is principally deterministic.
  - An autocorrelation test can be used on output of individual ring oscillators. [BLMT 2012]
  - These don’t produce an easy bound for min-entropy.
- Simple combinations of oscillators have nearly ideal statistical properties.
  - The XOR of the output of a small number of jitter-free oscillators pass most statistical tests (even though there is absolutely no entropy present). [BBFV 2010]
  - Any cryptographic processing will make the output look nearly ideal even in the absence of uncertainty.

Underlying Approach

- We proceed using an approach that is similar to one specified in a paper by Sunar et al. [SMS 2007]
- We assume that there is no uncertainty associated with when the cycle starts, the global elements of jitter, or the deterministic elements of jitter. (We only credit uncertainty induced by local Gaussian thermal noise since the last sample).
- Instead of tracking the literal output (which can look good, even in the absence of non-determinism) we track changes from the expected deterministic output.
- The underlying distribution depends on the number of oscillator cycles since the last sample (there are $\lfloor 1/(f_s T) \rfloor$ of these).
- The standard deviation of the accumulated jitter is thus $\delta = \sigma\sqrt{\lfloor 1/(f_s T) \rfloor}$.

Welcome to Mathistan

To express this, we use the cumulative distribution function of the jitter:

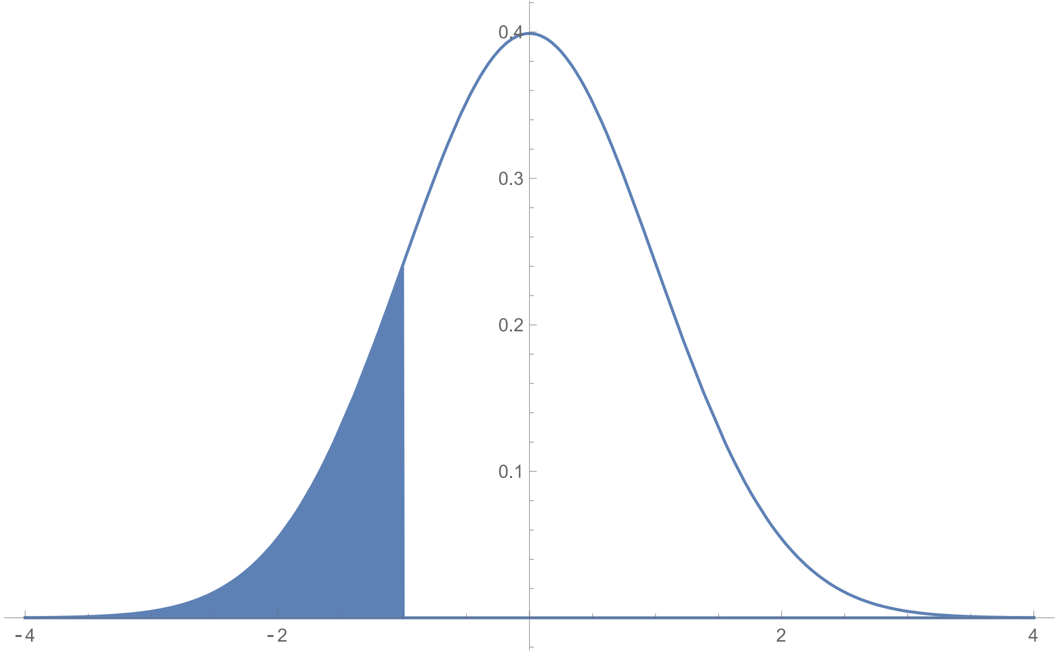

$$\mathrm{jCDF}_{\delta}(x) = \frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{x}{\delta\sqrt{2}}\right)\right]$$

where

$$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}}\int_{0}^{x} e^{-t^2}\,dt$$

“It’s only a model…”

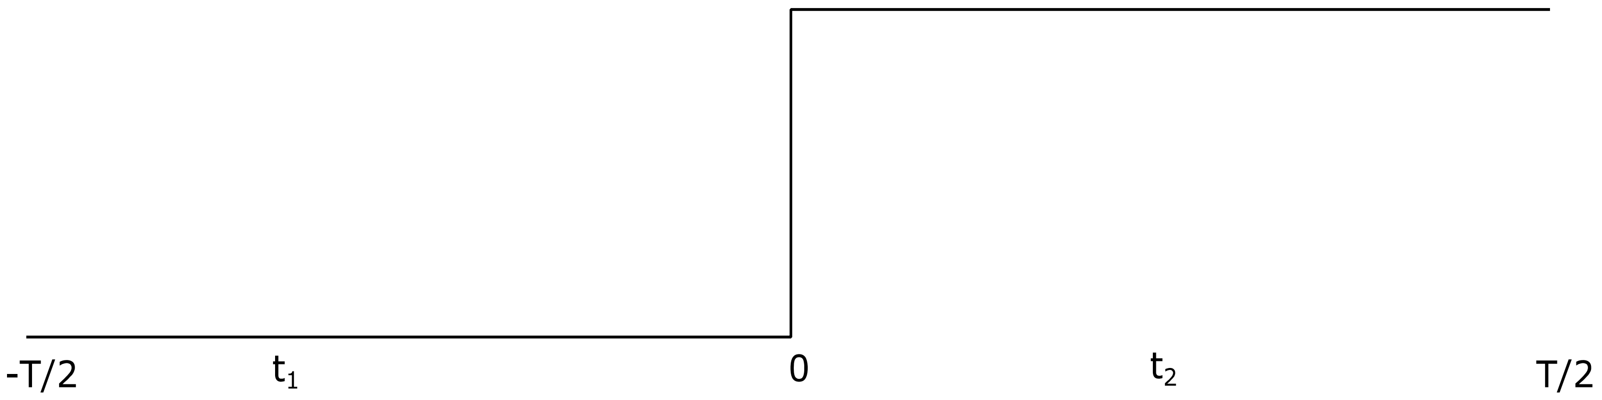

- We assume that the accumulated jitter between samples will be less than half the length of the period. [SMS 2007]
- At time $-\frac{T}{2} < t_1 < 0$, the probability of the output being a ‘1’ (rather than the expected ‘0’) is the same as the probability that the jitter is less than t_1, which is $\mathrm{jCDF}_{\delta}(t_1)$.
- At time $0 < t_2 < \frac{T}{2}$, the probability of outputting a ‘0’ rather than a ‘1’ is the same as the probability that the jitter is greater than t_2, which is based on the complementary CDF, $\left(1 - \mathrm{jCDF}_{\delta}(t_2)\right)$.

Buckets

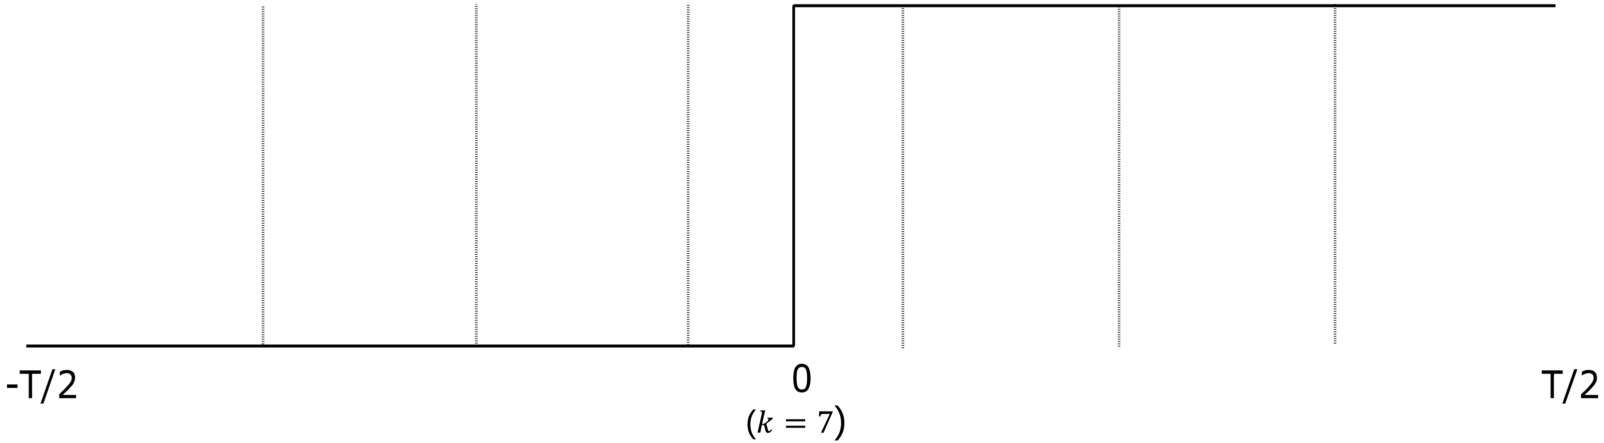

- Subdivide the period into k buckets.
- In a non-center bucket, the expected deterministic output is the most likely symbol
- In the center bucket the distribution of symbols is symmetric, so the choice of most likely symbol is free.
- For $-\frac{T}{2} < t < 0$, the maximal probability of a deterministic output (thus smallest min-entropy in the bucket) is obtained by choosing the left hand edge of the bucket.
- For $0 < t < \frac{T}{2}$, this is instead the right hand side of the bucket.

“Better get a Bucket”

- We assume that we sample a bucket chosen uniformly at random.
- This assumption makes sense, on average, for most designs.
  - Avoid clear integer relationships between the oscillator period and the sample period.
  - Exact relationships never really happen…
- We can then bound the probability of a deterministic output:

$$\Pr(\text{det. output}) \leq P_k = \frac{1}{k}\left[\sum_{\ell=1}^{(k+1)/2} \text{jCDF}_\delta\left(-\frac{T}{2}+\frac{(\ell-1)T}{k}\right)+\sum_{\ell=(k+3)/2}^{k}\left(1-\text{jCDF}_\delta\left(-\frac{T}{2}+\ell\frac{T}{k}\right)\right)\right]$$

thus

$$H_\infty(\text{output}) \geq -\log_2(P_k)$$

How Many Buckets?

- By choosing a low number of buckets, you are discarding assessed entropy.
- By choosing a very high number of buckets, you are making the “bucket is chosen uniformly at random” claim harder to justify.
- There is a limit to how much entropy you can get out of the system.
- Each sum above is just a Riemann Sum, so (accounting for symmetry) in limit you get:

$$P_{\infty} = \frac{2}{T}\int_{-T/2}^{0} \text{jCDF}_{\delta}(x)\, dx$$

So… How Many Buckets?

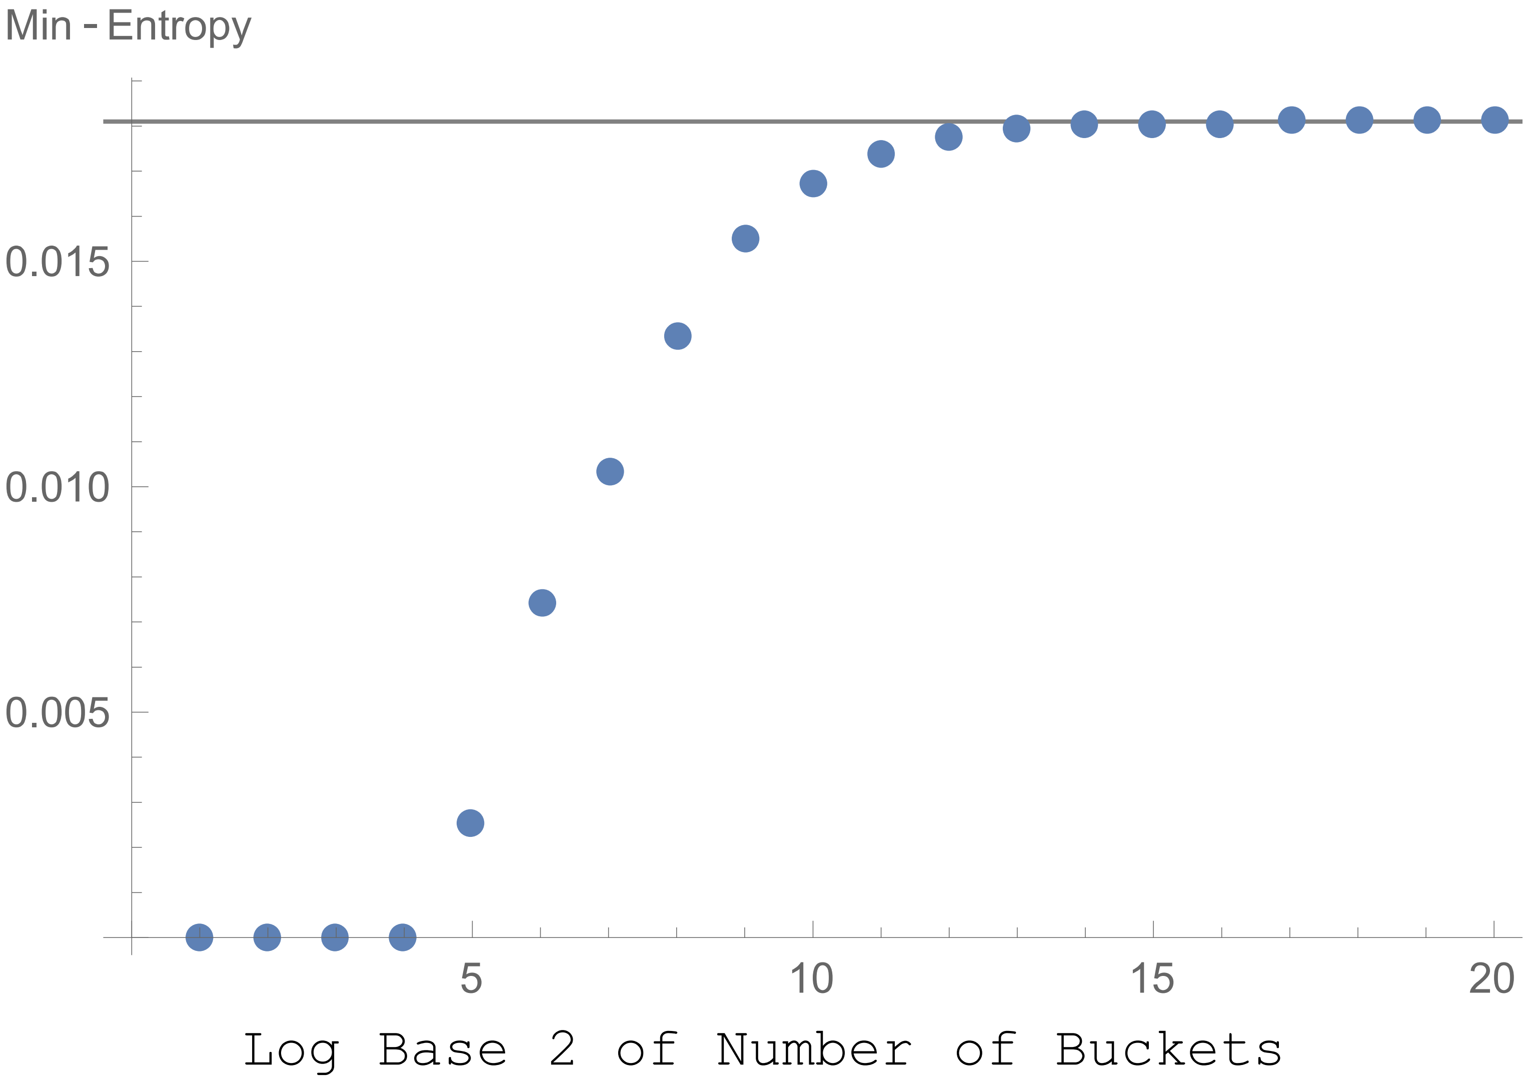

How to Arrive at Parameters?

- Direct measurement of cycle times can be accomplished
  - Internally (using a stable fast oscillator)
  - Externally (being careful not to introduce extra jitter through the sample method)
- Produce ring oscillators in matched pairs, and then subtract the measurements of the pairs
  - Global noise cancels
  - Some local switching noise can cancel (if the rings are sufficiently intertwined)
  - Correct for variance.
  - Can also use a (divided) version of one to trigger sampling of the other.
- Take measurements across all supported environmental and voltage ranges. Use the smallest measured jitter standard deviation.
- Direct continuous measurement of variance is also possible. [FL 2014]

Assumptions I

- The minimum standard deviation of the per-cycle jitter due to local Gaussian thermal noise is σ for any set of environmental conditions within the normal environmental operational conditions or other hardware components within the same core.
- Any accumulated jitter between samples will be less than half the length of the period. [SMS 2007]
- We assume that each bucket is equally likely to be sampled from. [SMS 2007].

Assumptions II

- The jitter originating from local Gaussian thermal noise is independent from sample to sample.
- The oscillators are sampled at a fixed frequency of f_s.
- All system design details are known by any attacker (Kerckhoffs' principle).
- There is no uncertainty associated with when the cycle starts, the global elements of jitter, or the deterministic elements of jitter. (We only credit uncertainty induced by local Gaussian thermal noise since the last sample).

System Parameters

Variable	Description
T	Period of the ring oscillator.
σ	Standard deviation of the per-cycle local Gaussian thermal noise induced jitter for the ring oscillator.
f_s	Sampling frequency.
k	Number of buckets used for the model.

- The values of T and σ are expected to vary across the operational voltage and temperature ranges. The values that result in the smallest assessed entropy should be chosen.
- Any change to the sampling frequency requires an updated assessment.

References

- [BLMT 2012] Baudet, Lubicz, Micolod, and Tassiaux. *On the Security of Oscillator-Based Random Number Generators*. Journal of Cryptology. Vol. 24, No. 2, 2011, pp 398 - 425.
- [BBFV 2010] Bochard, Bernard, Fischer, and Valtchanov. *True-Randomness and Pseudo-Randomness in Ring Oscillator-Based True Random Number Generators*. International Journal of Reconfigurable Computing, Vol. 2010.
- [Fischer 2012] Fischer. *A Closer Look at Security in Random Number Generators Design*. Third International Workshop on Constructive Side-Channel Analysis and Secure Design, COSADE 2012, May 2012.
- [FL 2014] Fischer and Lubicz. *Embedded evaluation of randomness in oscillator based elementary TRNG*. CHES 2014.
- [HLL 1999] Hajimiri, Limotyrakis, and Lee. *Jitter and Phase Noise in Ring Oscillators*. IEEE Journal of Solid-State Circuits, Vol. 34, No. 6, June 1999, pp. 790 - 804.
- [SMS 2007] Sunar, Martin, and Stinson. *A Provably Secure True Random Number Generator with Built-In Tolerance to Active Attacks*. IEEE Transactions on Computers, Vol. 56, No. 1, January 2007.

THANK YOU.