Health Testing for Periodically Sampled Ring Oscillators

or “Fugues and Their States”

Joshua E. Hill, PhD

CMUF Entropy WG
20230530
(Presentation Version *20230602*)

Health Testing For Ring Oscillator Designs

In today's talk, I'll discuss:

- Ring Oscillator Parameters.
- Ring Oscillator Failure Modes.
- Ring Oscillator Health Test Cutoffs.
- Statistical Power for Health Tests
  - RCT
  - APT
  - Crosswise RCT

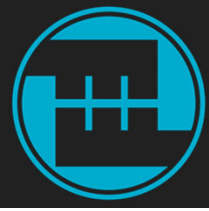

Ring Oscillator Parameters

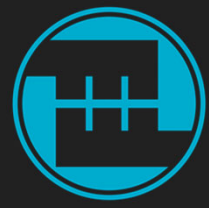

What Parameters?

- Stochastic models are commonly parameterized in terms of a single normalized accumulated jitter proportion (called “quality” in various of these papers).
- q denotes the jitter standard deviation, T_o $(= 1/f_o)$ is the ring oscillator’s period, and T_s $(= 1/f_s)$ is the sample period.

Jitter Proportion

$$q = \frac{\sigma_o}{T_o}\sqrt{\lfloor T_s/T_o \rfloor}$$

Number of RO cycles per sample

$$= f_o \sigma_o \sqrt{\lfloor f_o/f_s \rfloor}$$

What Parameters?

Another relevant parameter is the amount of expected inter-sample phase change.

- Formally, this is defined as

$$T_s = \left\lfloor \frac{T_s}{T_o} \right\rfloor T_o + \upsilon T_o,$$

where $\upsilon \in [0,1)$. This is equivalent to the statement $\upsilon \equiv \frac{T_s}{T_o} \bmod 1$.

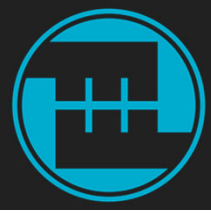

Reference Design

- 8 identical ring oscillator copies.
- Output is concatenated to an 8-bit raw symbol.
- Unless otherwise specified:

Parameter	Nominal Value
RO Frequency	1 GHz
Sample Frequency	9.985MHz
Jitter (per RO cycle)	10ps
q **(RO "quality")**	10%
υ **(expected inter-sample phase change)**	16%
Assessed Entropy per RO bit	0.07
Raw Data Assessed Entropy	0.575
α **(Targeted False Positive Rate)**	2^{-20}

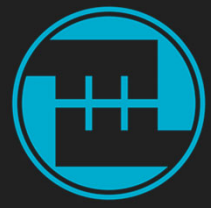

Health Testing in SP 800-90B

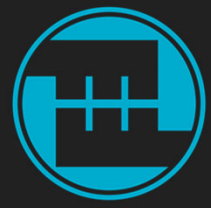

Health Testing General Requirements

- SP 800-90B does not require the use of specific health tests.
- It is commonly viewed as easier to progress through the SP 800-90B ESV program if the approved APT and RCT are performed on the raw noise produced by the noise source.
- Failure signaling is mediated by a specified abstract interface (**Get_entropy_input()** function).
- Reporting of a failure is required when entropy is requested. [Shall ID #66].
- Different types of errors (e.g., intermittent and persistent failures) can be described.
- Persistent errors must inhibit data output from the entropy source. [Shall ID #67-69]

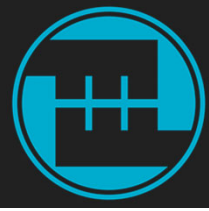

Ring Oscillator Failure Modes

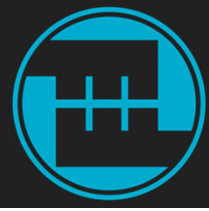

Failure to Oscillate

- If a ring experiences a total failure (i.e., it fails to oscillate), then its output becomes fixed.
  - It outputs a ‘0’ or ‘1’, depending on the phase on failure.

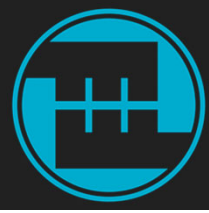

Oscillator Output Entrainment

- In Markettos and Moore oscillators were caused to entrain to an external signal injected using the power line. [MM 2009]
- Similar attacks have also been conducted by Bayon et al. using injected RF signals. [BBAFPRM 2012].
- Bochard et al. suggests that an attacker can cause as many as 25% of the oscillators to become mutually locked, and thus not contribute entropy. [BBFV 2010]
- Mureddu et al. notes that output entrainment tended to only occur when the nominal frequencies were close (<3% difference). [MBBF]
- When rings become output entrained, they change output with a fixed relationship.

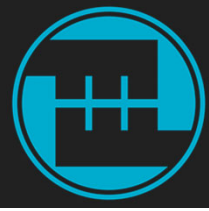

Failure Mode Commonalities

- Here, a failure leads to a substantial change in the probability for each output symbol.
- For failed rings the output is fixed.
- For output entrainment, there is a fixed relationship between the entrained rings.
- In both cases, this failure results in an output distribution change which increases the probability of the most common symbol.
- This makes both APT and RCT failures more likely.

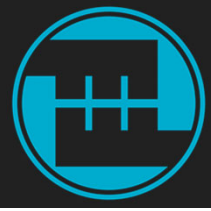

Health Testing Parameters

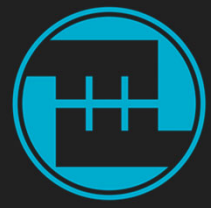

Cutoff Selection

- Health test cutoffs for each test can be established in several ways.
- Any of these are compatible with SP 800-90B.
- The cutoff selection procedure needs to be described within the Entropy Analysis Report.

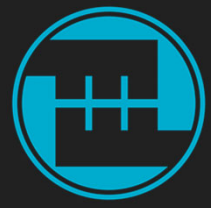

90B Cutoff Approach

- NIST allows flexibility for:
  - The targeted false positive rate,
  - its connection to the selected cutoffs, and
  - the entropy estimate used for this procedure.
- SP 800-90B recommends that targeted false positive rates be in the interval $[2^{-40}, 2^{-20}]$.
  - NIST allows for choices outside this range. [Shall ID #70]
- There are choices for the entropy estimate used for generating the cutoff values.
  - $H_{\text{submitter}}$
  - Lower than $H_{\text{submitter}}$ to reduce the effective false positive rates. [SP 800-90B Section 4, Footnote 5]
  - Higher than $H_{\text{submitter}}$ to increase the statistical power of the health testing.

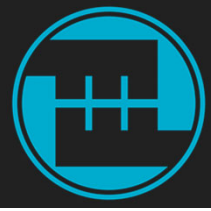

Cutoff Selection

Some cutoff selection options:

- Establishing cutoffs using the same approach as used in SP 800-90B using
  - The $H_{\text{submitter}}$ estimate (marked “90B” here), or
  - A different entropy estimate.
    - For the APT and RCT, we will see that a good candidate for this alternate estimate is the Most Common Value (MCV) entropy estimate.
    - Such cutoffs are marked “MCV” here.

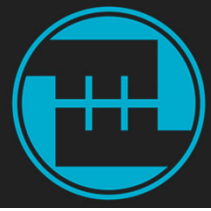

Cutoff Selection

Some (more) cutoff selection options:

- Empirical cutoffs
  - Can be generated with
    - Actual data, or
    - Simulated data
  - Empirical cutoffs are the smallest cutoff value such that the number of observed errors is less than or equal to the number of expected errors for that dataset size.
  - There are tools within [Theseus] that can be used to both simulate this source, and estimate the false positive rate for various cutoff settings.

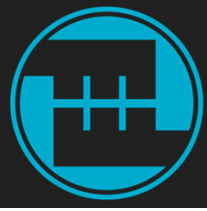

90B Cutoff Approach

- The procedures outlined for cutoff selection ([SP 800-90B Sections 4.4.1 and 4.4.2]) presume an IID noise source.
  - The only statistical defect for an IID source is bias.
  - The APT and RCT are sensitive to changes in bias, and insensitive to failure modes that do not have a substantial impact on the bias.
- Ring oscillator designs have substantial internal state (the per-oscillator phase with respect to the sampling clock), and so are not IID sources.
- A more reasonable cutoff / false positive estimate uses the result of the MCV estimator.
  - This characterizes the symbol bias.

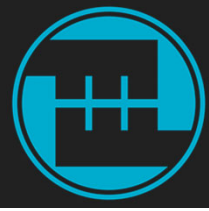

Cutoff Selection

- Here, empirical tests use simulated data.
- Simulated datasets are large (1 billion samples each) for each tested condition.

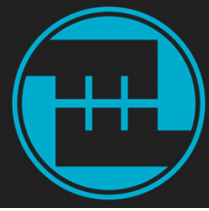

Health Tests

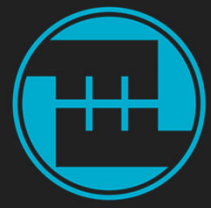

Repetition Count Test

- The RCT is a simple test that checks for a catastrophic failure of the noise source that results in a fixed output from the noise source.
- Cutoffs based on the example noise source:

Cutoff Approach	C
90B	36
MCV	4
Empirical	4

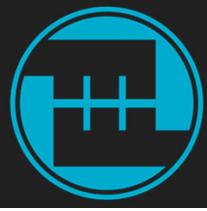

RCT Power

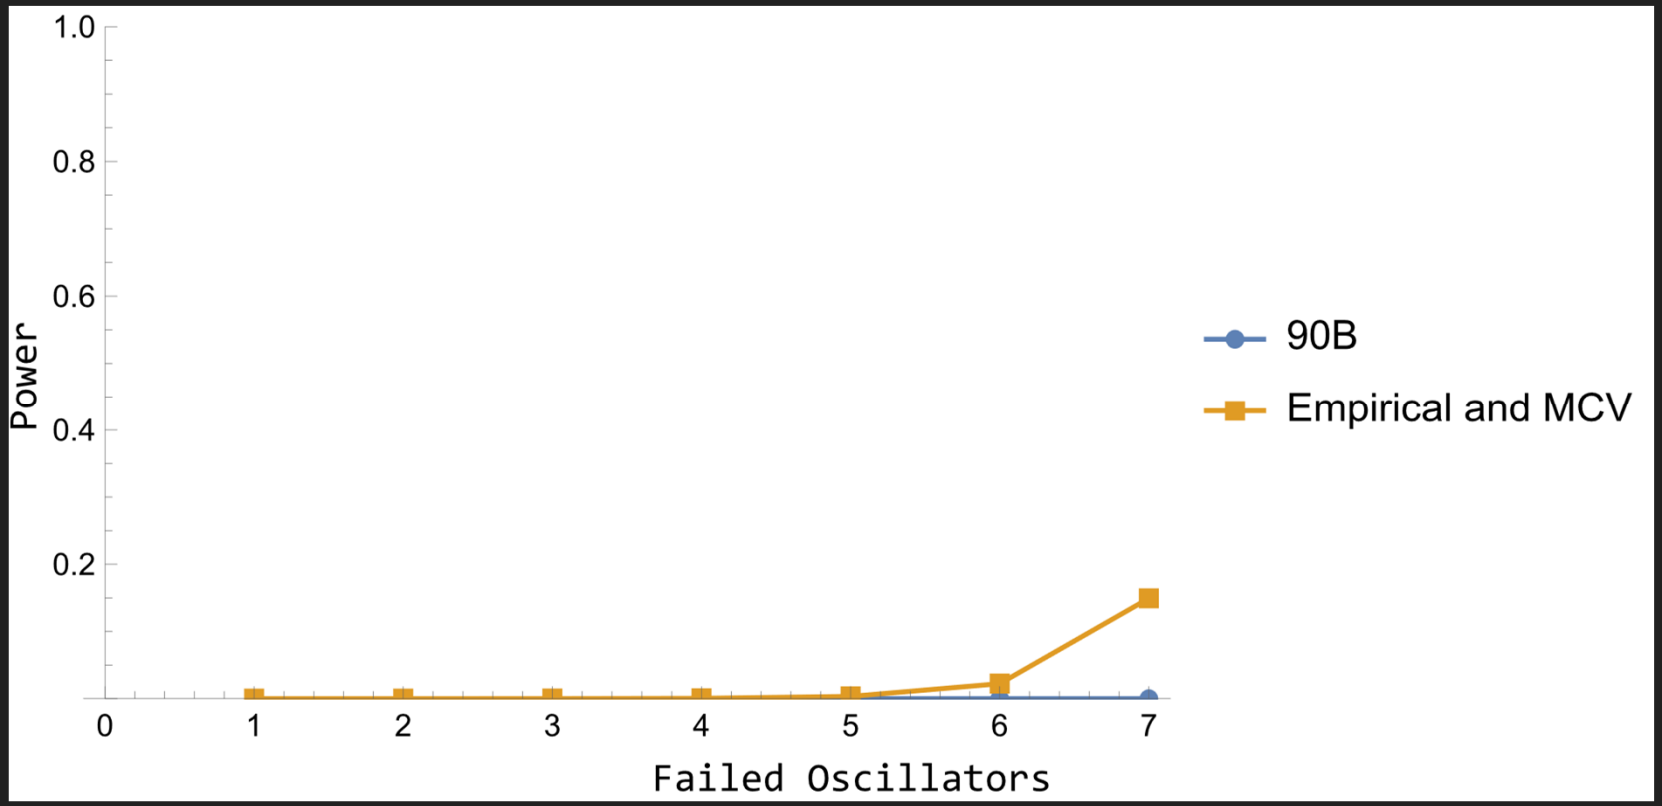
1.0
0.8
0.6
0.4
0.2
Power
0
1
2
3
4
5
6
7
Failed Oscillators
90B
Empirical and MCV

RCT Power

- The RCT is a health test that looks for catastrophic failure.
- This test is fairly insensitive to entrainment until 6 of the rings have failed (thus 7 of the 8 rings produce outputs with persistently fixed relationships).
- In this test, failure is not likely until all of the rings have become fixed (and in this failure mode, the power of the test is 100%).
- Here, the cutoffs selected using the empirical and MCV approaches are the same, so their power graphs coincide.

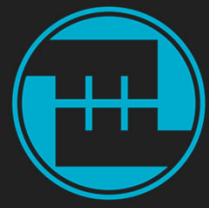

Adaptive Proportion Test

- The APT is a test that checks to see if the per-window reference value is more likely to occur within a window than expected.
- Cutoffs based on the example noise source:

Cutoff Approach	C
90B	394
MCV	14
Empirical	15

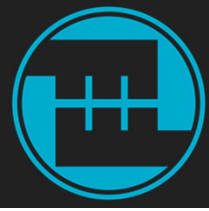

APT Power

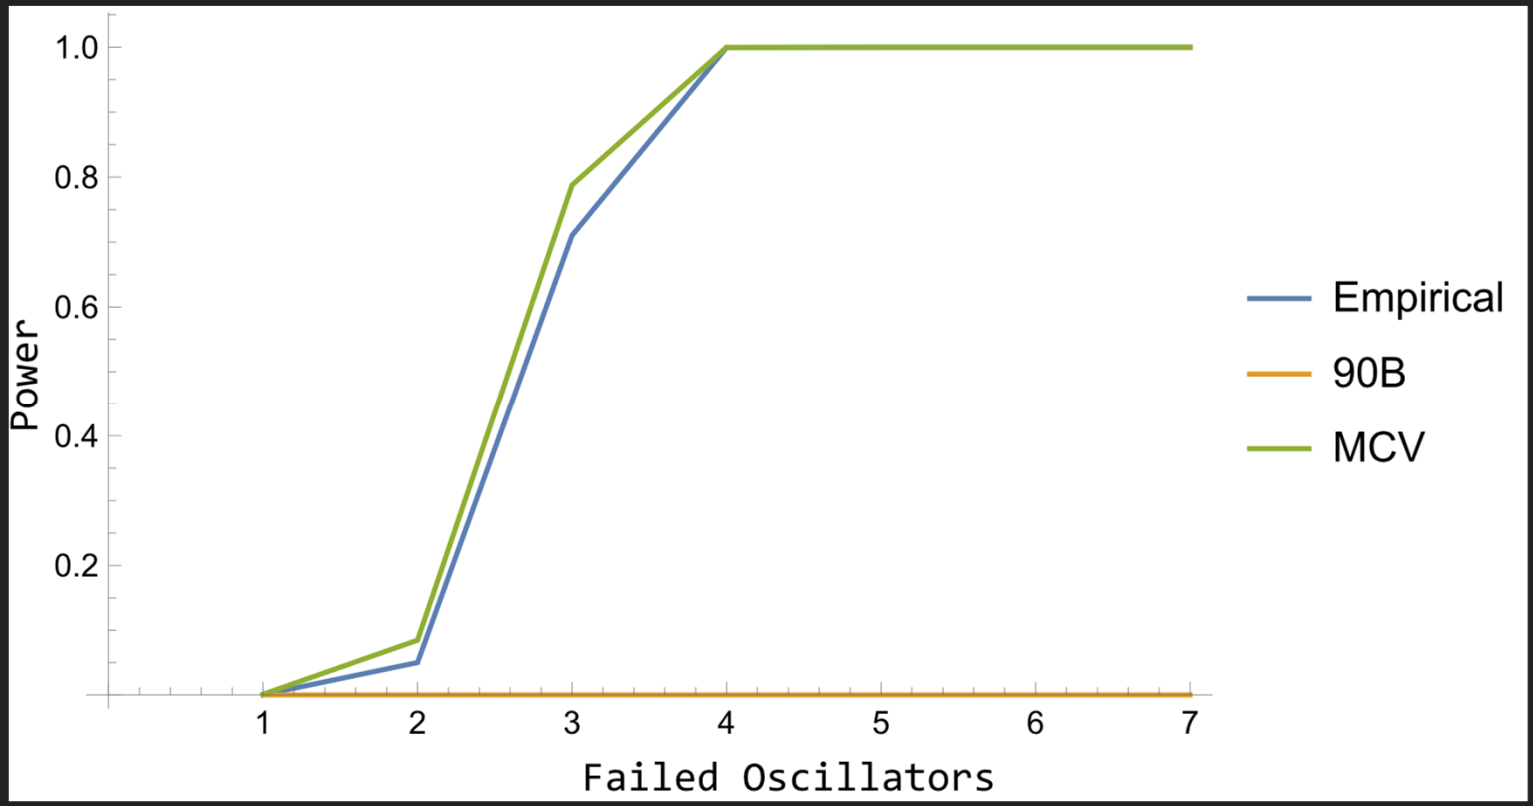
1.0
0.8
0.6
0.4
0.2
Power
1
2
3
4
5
6
7
Failed Oscillators
Empirical
90B
MCV

APT Power

- Abstracting SP 800-90B Section 4.5 Criterion 2, we examine the “4 failed oscillators” case.
- Here we produce approximately 50% of the assessed entropy.
- We can examine the power and false positive rate for the resulting APT test for various cutoff selections.

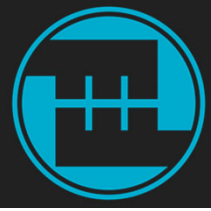

APT Power

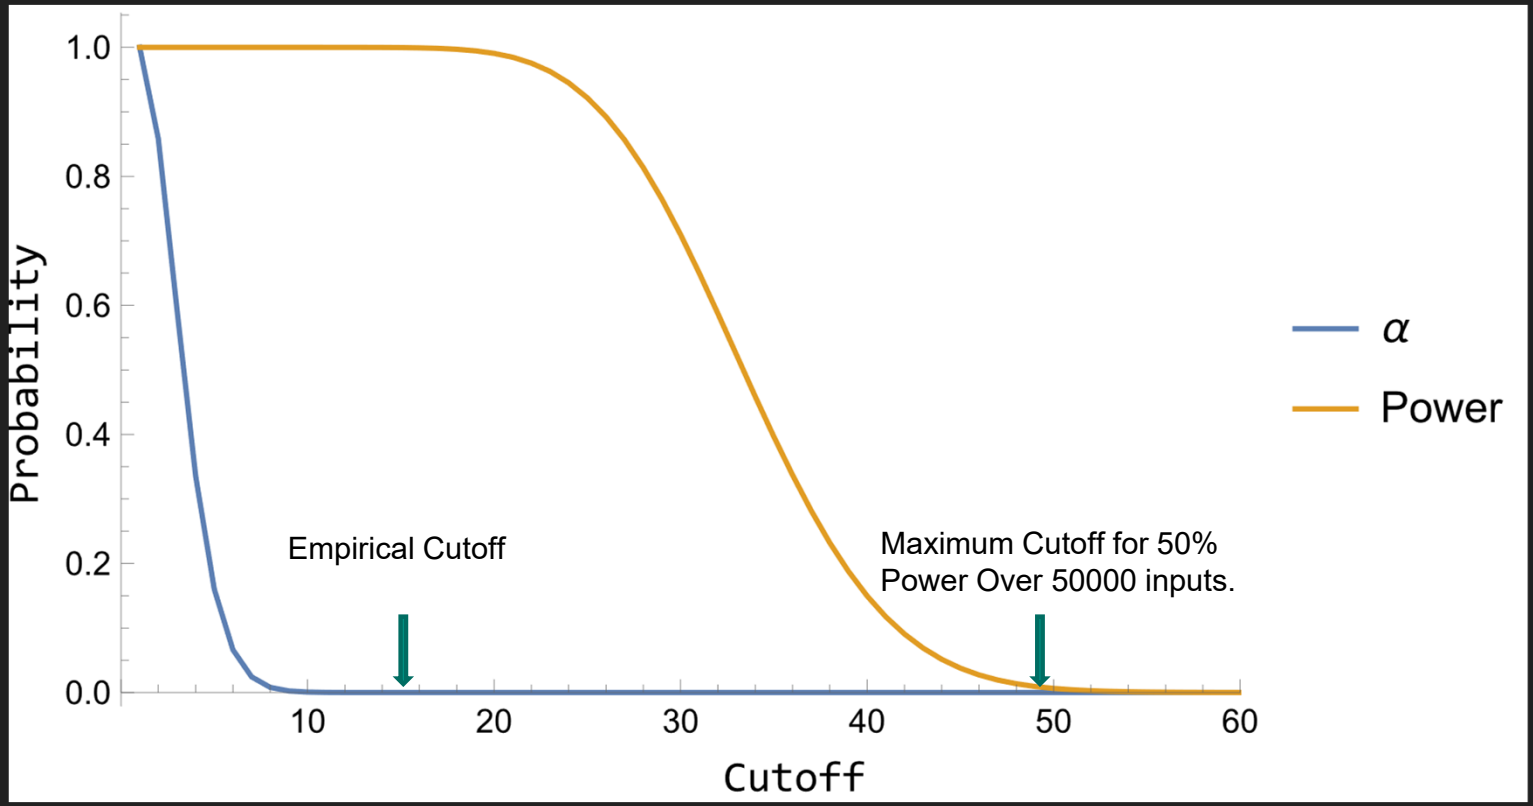
1.0
0.8
0.6
0.4
0.2
0.0
Probability
10
20
30
40
50
60
Cutoff
Empirical Cutoff
Maximum Cutoff for 50%
Power Over 50000 inputs.
α
Power

APT Power

- As ring failures occur, an APT failure becomes much more likely.
- The 90B cutoff has limited power for anything but a catastrophic failure.

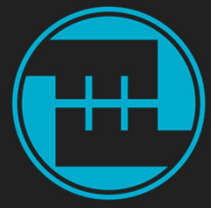

Crosswise RCT

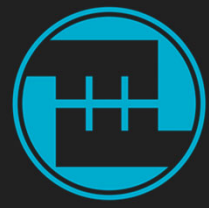

The Crosswise RCT: Mission Statement

If we are concerned about failed rings and output entrainment, then why not actually look for failed rings and output entrainment?

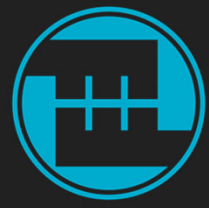

The Crosswise RCT

- The APT and RCT view each n-bit raw data sample as a single monolithic symbol.
- Here, each such n-bit symbol is the output of n noise source copies.
- These individual noise source copy outputs can be tested, both independently and in relationship to each other, in order to directly detect these anticipated failure modes.

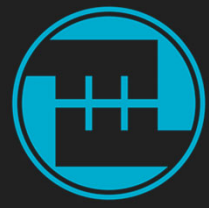

The Crosswise RCT

The Crosswise Repetition Count Test (CRCT) decomposes the raw input into n independent samples, and then:

1. Runs an RCT on the individual noise source copy bits (the “literal” RCTs). The literal RCTs detect total failure of the individual rings.
2. Runs an RCT on the XOR of every pair of distinct rings (the “cross” RCTs). The cross RCTs detect protracted runs of output where a noise source copy pair is either persistently equal, or persistently complements of each other.

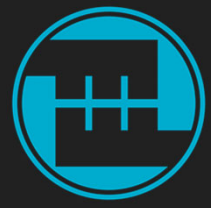

Crosswise RCT Cutoffs

- The literal and cross RCTs have different statistical behavior, so it is desirable for the CRCT to have distinct cutoffs for literal RCTs and cross RCTs.
- The relationship between the cutoffs depends on the system parameters.

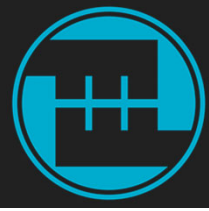

Crosswise RCT Cutoffs (Fixed $q = 10\%$, υ Varies)

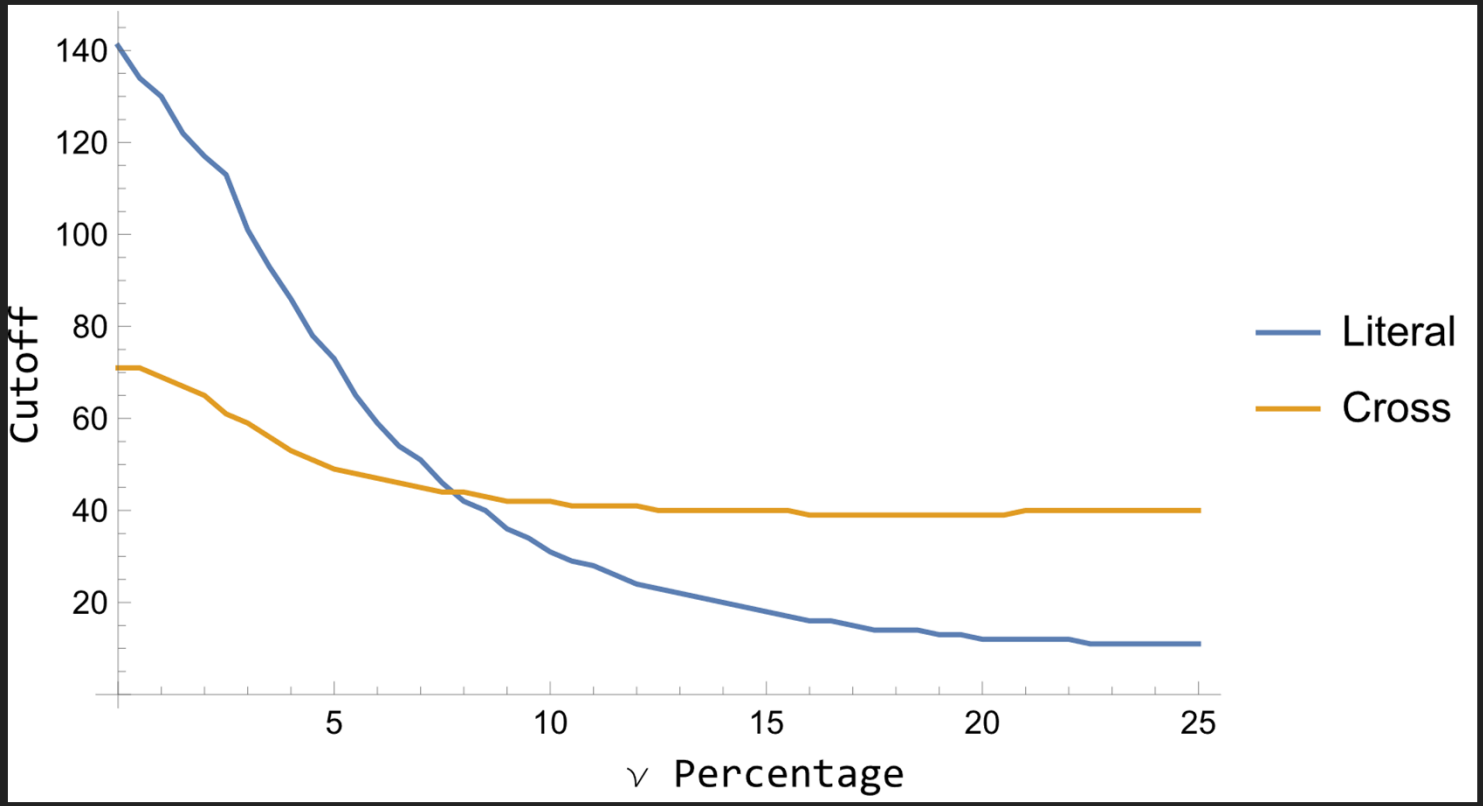

Crosswise RCT Cutoffs (Fixed $\upsilon = 0$, q Varies)

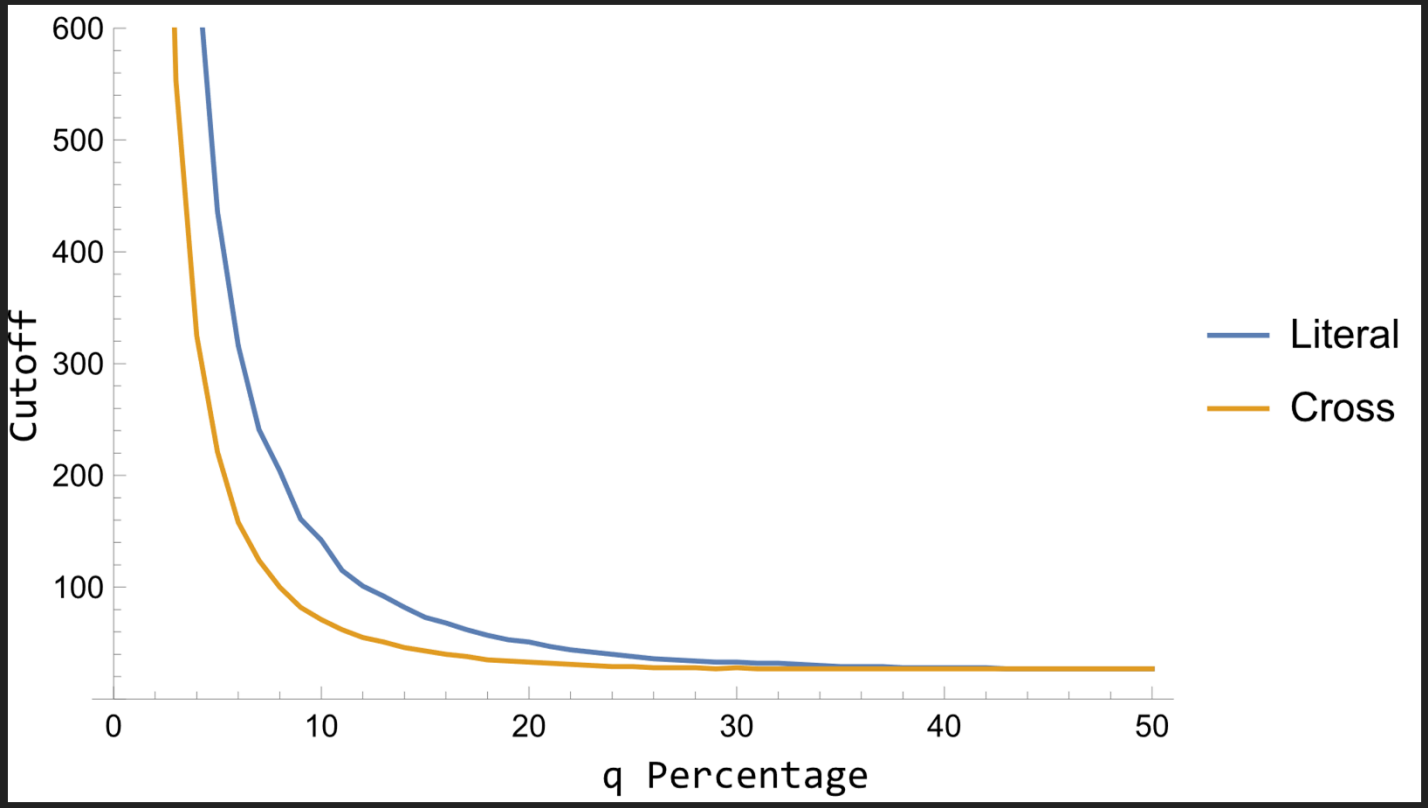

Crosswise RCT Cutoffs

It is useful to use one of the empirical approaches for establishing the cross RCT cutoffs.

Cutoff Approach	C
Literal	16
Cross	39

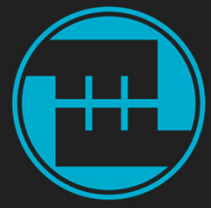

Crosswise RCT Specification

The CRCT is run as follows:

1. A=crossExpand(`next()`).
2. B=(1,…,1).
3. X=crossExpand(`next()`).
4. For i = 1 to T
   - If (X[i] == A[i]),
     - B[i]=B[i]+1.
     - If (B[i] ≥ C[i]), signal a failure.
   - else:
     - B[i]=1.
5. A=X.
6. Go to Step 3.

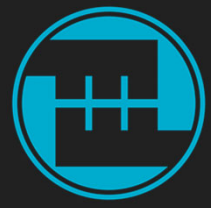

The Crosswise RCT

- In software the crossExpand() function can be implemented using a few circular shifts/XOR rounds, building up a wide symbol that is then bitwise assessed.
- In hardware this is a few XOR gates + as many 1-bit RCT implementations as needed.

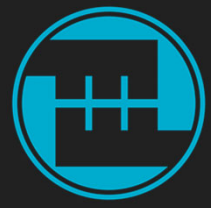

The Crosswise RCT False Positive Rate

- A Crosswise RCT is really a bunch of (mostly independent) 1-bit RCTs standing on each other's shoulders in a large trench coat.
- For a n-bit source there are:
  - n literal RCTs.
  - $\binom{n}{2}$ cross RCT tests.
- For the example 8-RO source, this yields 36 distinct sub-tests.
  - 8 literal RCTs.
  - 28 cross RCTs.
- The targeted per-sub-RCT false positive is then:

$$\alpha_s = 1 - (1 - \alpha)^{\frac{1}{36}}$$

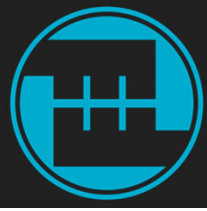

Crosswise RCT Power

- This test detects persistent output entrainment or fixed outputs with a power of 100%.
- A larger cutoff doesn’t reduce power for persistent failures, it only delays signaling.

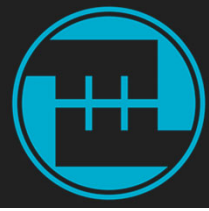

Conclusion

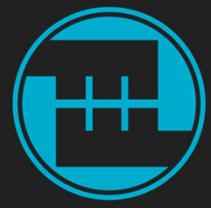

What Were We JUST Talking About?

- The RCT is only sensitive to total failure.
- The APT can be sensitive to output entrainment, so long as the cutoff is chosen reasonably.
- For this type of design, the described APT and RCT selection procedure:
  - Does not yield the targeted false positive rate.
  - Yields health tests sensitive only to catastrophic failure.
- Using the MCV assessment produces a fairly reasonable cutoff.
- Empirical cutoffs produce a fairly reasonable cutoff.

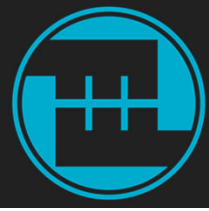

What Were We JUST Talking About?

- The Crosswise RCT is a health test that effectively detects the two failure modes outlined here.
- Appropriate cutoffs should be chosen empirically.

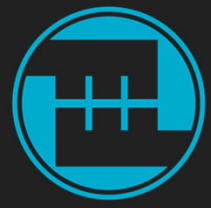

References

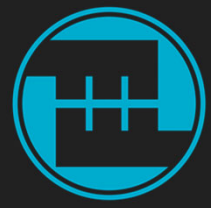

References (Failure Modes)

- [BBAFPRM 2012] Pierre Bayon, Lilian Bossuet, Alain Aubert, Viktor Fischer, François Poucheret, Bruno Robisson, Philippe Maurine. *Contactless Electromagnetic Active Attack on Ring Oscillator Based True Random Number Generator*.
- [BBFV 2010] Nathalie Bochard, Florent Bernard, Viktor Fischer, and Boyan Valtchanov. *True-Randomness and Pseudo-Randomness in Ring Oscillator-Based True Random Number Generators*. International Journal of Reconfigurable Computing, Vol. 2010.
- [MM 2009] Markettos and Moore. “The Frequency Injection Attack on Ring-Oscillator -Based True Random Number Generators” CHES 2009.
- [MBBF] Mureddu, Bochard, Bossuet, and Fischer. *Experimental Study of Locking Phenomena on Oscillating Rings Implemented in Logic Devices.* 2019.

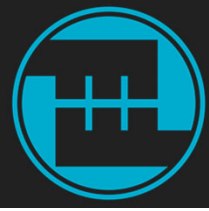

References (Program)

- [Shall ID] *90B-Shall-Statements*. https://csrc.nist.gov/CSRC/media/Projects/cryptographic-module-validation-program/documents/esv/90B%20Shall%20Statements.xlsx.
- [SP 800-90B] Meltem Sönmez Turan, Elaine Barker, John Kelsey, Kerry A. McKay, Mary L. Baish and Mike Boyle. *Recommendation for the Entropy Sources Used for Random Bit Generation*. January 2018.

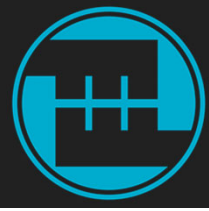

References (Software)

- [Theseus] https://github.com/KeyPair-Consulting/Theseus