A Metric for Measuring the Evenness of Timing System Rep-Rate Patterns

E. Bjorklund

Abstract

When computing repetition rate patterns for SNS timing events, it is sometimes useful to be able to measure the quality of a proposed pattern. The quality of a pattern is determined primarily by the evenness of its pulse distribution. In this tech-report, we describe the "Ugliness" function – which measures the degree to which a pattern deviates from perfect symmetry. An "Individual Ugliness" metric is also developed which allows us to identify the specific pulse or pulses that are the major contributors to a pattern's collective ugliness. The general ugliness function is computable in $\mathcal{O}(n^2)$ time, where n is the repetition rate. The individual ugliness function is computable in $\mathcal{O}(n)$ time.

Introduction

Computing super-cycle patterns for various repetition rate values is a fairly efficient and straight-forward process as long as there are no constraints on which cycles are available for scheduling the pulses (for the details of this process, see [1] and [2]). Once you introduce constraints into the system (for example, "Gate A may only occur on cycles that also contain Gate B, but not on cycles that contain Gate C…") you can quickly run into scenarios in which computing the most evenly distributed pattern that meets the given constraints is cost-prohibitive. You then need to resort to heuristic solutions that can only guarantee "relatively good" patterns. As with most optimization problems, the heart of a good heuristic lies in having a good, easily computed metric to let you know whether a proposed solution is better or worse than the other candidates.

The Properties of Ugliness

It is easy to recognize a good pattern when you see one. The trick is to get a computer to recognize it. For example, if the pattern width (super-cycle length) is 8 and your repetition rate (number of pulses) is 2, then the pattern:

10001000

is clearly an optimal distribution. Between every '1' in the pattern, there are three '0's. Since timing systems repeat their super-cycles *ad nauseam*, any rotation of the above pattern is also an optimal distribution. When the rep-rate does not evenly divide the super-cycle length, or when there are constraints on where you can place the '1's, the situation becomes less obvious. In the two patterns below:

01010100	01001100
Pattern A	**Pattern B**

It is clear that Pattern A is more evenly distributed than Pattern B. We would like our metric to distinguish between these two patterns as well.

Apart from just measuring evenness, there are some other properties we would like our metric to have:

- **The metric must be invariant under rotations**. Since the super-cycle is repeated *ad nauseam*, it is not really important where the pattern begins or ends. Therefore, the two patterns:

01010100 00010101

should have the same value.

- **The metric should be efficient to compute**. Given the SNS super-cycle length of 600, we would like to restrict the computational complexity to no more than $\mathcal{O}(n^2)$.
- **The metric should return 0 for a perfectly even pattern**. The value should increase as the symmetry of the pattern decreases. This just makes the metric easier to use than if we made the other, less convenient, choice of assigning ∞ to a perfectly even pattern.

The last constraint implies that our metric is actually measuring the lack of symmetry in our pattern. Since the ancient Greeks equated symmetry with beauty, we call our metric the "*Ugliness*" of the pattern.[1]

Some Methods That Don't Work (and Why)

Before describing the metric we eventually adopted, it is useful to gain insight into the problem by exploring some techniques that did not meet our criteria, and the reasons for their failure. In this section we also develop some of the nomenclature that we will be using to describe the ugliness function (so you can't skip over this section – in spite of the title).

The standard statistical tools for measuring distributions are "Variance" (σ^2), and "Mean Average Deviation" (MAD). Variance is computed as:

$$\text{var} = \sigma^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^2 \quad (1)$$

where $\bar{x}$ is the average value:

$$\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i \quad (2)$$

The Mean Average Deviation is similar to the variance, but uses absolute value (instead of squaring) to remove the "signedness" of the deviations:

$$MAD = \frac{1}{n}\sum_{i=1}^{n}|x_i - \bar{x}| \quad (3)$$

Because of the greater range in the variance (squaring the difference produces larger values than taking the absolute value does), the "variance" approach usually produces better results.

One obvious first attempt at measuring the ugliness of a pattern would be to measure either the variance or MAD of the "forward distances" between the '1's in the pattern.

[1] One implication of this is, of course, is that you can only be so beautiful, but there is no limit on your potential for ugliness.

Let $d(i, j)$ be the distance, measured in a "forward" direction (i.e. left-to-right), with wraparound, between the *i*th and *j*th elements of the pattern. $d(i, j)$ is computed as follows:

$$d(i, j) = (j - i) \bmod m \tag{4}$$

where *m* is the length of the pattern (super-cycle) and the "mod" function is defined such that it always returns a positive value, e.g:

$$(-x) \bmod m = (m - x) \bmod m \tag{5}$$

Now define the "Pulse Set" (*P*) of a pattern to be the ordered set of the index values of each pulse in the pattern (i.e. "all the ones"). Formally put, we have:

$P = \{i \mid \text{the } i\text{th slot in the pattern contains a } 1\}$

For example, if the pattern is:

01101101

then

$P = (1,2,4,5,7)$

Note that the first element of the pattern has index value 0. This convention allows us to use modular arithmetic in the indices without requiring further alterations to our definition of the "mod" function. Note also that *P* is an "<u>ordered</u>" set (this will be important later on).

Finally, we define $\delta(i)$ to be the forward distance between the *i*th element of *P* and its nearest neighbor in the forward direction (with wraparound). $\delta(i)$ is therefore computed by the equation:

$$\delta(i) = d\left(P_i, P_{(i+1) \bmod n}\right) \tag{6}$$

The "variance" of a rep-rate pattern, can then be expressed as:

$$\text{var} = \frac{1}{n} \sum_{i=0}^{n-1} \left(\delta(i) - \bar{\delta}\right)^2 \tag{7}$$

where *n* is the number of '1's in the pattern.

So our previous pattern, 01101101, which is represented by the pulse set, P=(1,2,4,5,7), generates the δ set, δ =(1,2,1,2,2), from which we compute a pattern variance of 0.24.

At first glance, the variance method appears to have all the characteristics we are looking for in a good metric. It is invariant over rotations, it is zero for a perfectly symmetrical pattern, and it is computable in linear time. However, to see how the variance method fails, consider the pattern:

01110101

which also has a variance of 0.24, but is clearly less desirable than our first pattern! To see why this is the case, consider the *P* and δ sets of the two patterns:

$P1$ = (1,2,4,5,7) $\delta 1$ = (1,2,1,2,2)
$P2$ = (1,2,3,5,7) $\delta 2$ = (1,1,2,2,2)

Both δ sets contain two 1's and three 2's, so the sum and mean are the same. The variance calculation sums the same terms for both patterns (although in a different order) and so it arrives at the same value.

The key is to notice that the δ set for the first pattern has a more even distribution of 1's and 2's than the δ set of the second pattern. Clearly, we need to look at more than just the distances between a pulse and its immediate neighbor. We also need to consider the distribution of those distances, and the distribution of those distributions, *etc*. To do this, we need to look at the relationship between each pulse and every other pulse in the pattern, which leads to the algorithm we finally adopted as our "Ugliness" metric.

The Ugliness Equation

The algorithm we adopted to measure a pattern's "Ugliness" is as follows. First take the variance of the forward distances between adjacent pulses. Then take the variance between the forward distances of every other pulse, the variance between the forward distances of every third pulse, *etc*. The mean value of all these variances is defined to be the pattern's "Ugliness".

We define $\delta_j(i)$ to be the forward distance between pulse *i* and the *j*th pulse after *i*.

$$\delta_j(i) = d\left(P_i,\ P_{(i+j) \bmod n}\right) \qquad (8)$$

An interesting and useful observation on the δ_j function is the relation:

$$\sum_{i=0}^{n-1} \delta_j(i) = jm \qquad (9)$$

where *m* is the size of the entire pattern, and *n* is the number of pulses (1's). To understand the reason for this relationship, consider that when you sum the forward distances between each pulse and its nearest neighbor in the forward direction (δ_1), you are, in effect, traversing the entire pattern. When you sum the forward distances between each pulse and its second nearest neighbor in the forward direction (δ_2), you traverse the entire pattern twice, *etc*. A direct application of this relationship is that the mean of a δ_j function can be expressed simply as:

$$\bar{\delta}_j = \frac{jm}{n} \qquad (10)$$

The mean of the variances of the δ_j functions (the "Ugliness") can then be expressed as:

$$\frac{1}{n-1}\sum_{j=1}^{n-1}\left[\frac{1}{n}\sum_{i=0}^{n-1}\left(\delta_j(i) - \frac{jm}{n}\right)^2\right] \qquad (11)$$

Equation (11) can be computed in $\mathcal{O}(n^2)$ time, which fits our criteria for computational complexity. We can reduce the computation time by roughly one half, however, if we consider the symmetry between the δ_j and δ_{n-j} functions.

Let D be an $n \times (n-1)$ matrix constructed from P as follows:

$$D_{i,j} = \delta_j(i) - \frac{jm}{n} \tag{12}$$

For example, given the pattern 01101101, with P = (1,2,4,5,7), then:

$$D = \begin{bmatrix} -0.6 & -0.2 & -0.8 & -0.4 \\ 0.4 & -0.2 & 0.2 & 0.6 \\ -0.6 & -0.2 & 0.2 & -0.4 \\ 0.4 & 0.8 & 0.2 & 0.6 \\ 0.4 & -0.2 & 0.2 & -0.4 \end{bmatrix}$$

Note that the columns of D represent the differences between $\delta_j(i)$ and $\overline{\delta}_j$. In other words, the variance of a δ_j function is given by the mean of the sum of the squares of the elements of the jth column of D.

$$\operatorname{var}(\delta_j) = \frac{1}{n}\sum_{i=0}^{n-1}\left(D_{i,j}\right)^2 \tag{13}$$

Also note that the mean of all the $\operatorname{var}(\delta_j)$ functions is the “Ugliness” function specified by (11).

$$Ugliness = \frac{1}{n-1}\sum_{j=1}^{n-1}\left[\frac{1}{n}\sum_{i=0}^{n-1}\left(D_{i,j}\right)^2\right] \tag{14}$$

The following theorem will allow us to cut our computation time in half by only considering the first n/2 columns of D in the computation of equation (14)[2].

Theorem:

$$D_{i,j} = -D_{(i+j)\bmod n,\ (n-j)}$$

The “English Translation” of this theorem is that the last column of D is just a rotation and negation of the first column of D. The penultimate column of D is just a rotation and negation of the second column of D, *etc*. This relationship can be seen in the example above.

[2] At this point, the astute (or “nit-picky”) reader will probably have noticed the inconsistency in our indexing scheme. The row indices of D begin with 0, and the column indices of D begin with 1. This is an artifact of the definition of D owing to the fact that pulse and slot numbers (i) begin with 0 (so that we can use modular arithmetic on them) but forward distances (j) begin with 1 (the forward distance between a pulse and itself is of no interest). We chose to maintain this convention because it makes the following proof a lot cleaner.

Proof of Theorem:
By definition (12):

$$D_{(i+j)\,\mathrm{mod}\,n,\,(n-j)} = \delta_{n-j}\big((i+j)\,\mathrm{mod}\,n\big) - \frac{(n-j)m}{n} \tag{15}$$

By definition (8):

$$\delta_{n-j}\big((i+j)\,\mathrm{mod}\,n\big) = d\Big(P_{(i+j)\,\mathrm{mod}\,n},\ P_{((i+j)\,\mathrm{mod}\,n+(n-j))\,\mathrm{mod}\,n}\Big) \tag{16}$$

However, by the properties of modular arithmetic, and the fact that *j* is always less than *n*:

$$P_{((i+j)\,\mathrm{mod}\,n+(n-j))\,\mathrm{mod}\,n} = P_{((i+j)+(n-j))\,\mathrm{mod}\,n} = P_{(i+n)\,\mathrm{mod}\,n} = P_i \tag{17}$$

Backing up to equation (15), we now manipulate the second term of the subtraction to get:

$$\frac{(n-j)m}{n} = \frac{mn-mj}{n} = m - \frac{jm}{n} \tag{18}$$

Substituting (17) and (18) back into (15) gives us:

$$D_{(i+j)\,\mathrm{mod}\,n,(n-j)} = d\Big(P_{(i+j)\,\mathrm{mod}\,n},\ P_i\Big) - m + \frac{jm}{n} \tag{19}$$

Referring back to the definition of *d* given in (4), and our definition of the "mod" function given in (5), we can show, using the properties of modular arithmetic, that:

$$d(i,j) = m - d(j,i) \tag{20}$$

Substituting (20) back into (19), we get:

$$\begin{aligned} D_{(i+j)\,\mathrm{mod}\,n,(n-j)} &= m - d(P_i,\ P_{(i+j)\,\mathrm{mod}\,n}) - m + \frac{jm}{n} \\ &= -\delta_j(i) + \frac{jm}{n} \\ &= -D_{i,j} \end{aligned}$$

Q.E.D.

As a result of the above theorem, we know that the columns in the second half of the *D* matrix contain exactly the same numbers as the columns in the first half of the *D* matrix. Since the variance calculation squares the column values, the information in the second half of the *D* matrix is completely redundant and therefore does not need to be re-computed.

By using only the first $n/2$ columns of the D matrix, multiplying by 2, factoring out the averaging, and simplifying, we obtain the following “streamlined” ugliness equation:

$$Ugliness = \frac{2}{(n/2)} \sum_{j=1}^{n/2} \left[\frac{1}{n} \sum_{i=0}^{n-1} \left(\delta_j(i) - \frac{jm}{n} \right)^2 \right]$$

$$= 2\frac{2}{n}\left(\frac{n}{2}\frac{1}{n}\right) \sum_{j=1}^{n/2} \left[\sum_{i=0}^{n-1} \left(\delta_j(i) - \frac{jm}{n} \right)^2 \right]$$

$$= \frac{2}{n} \sum_{j=1}^{n/2} \left[\sum_{i=0}^{n-1} \left(\delta_j(i) - \frac{jm}{n} \right)^2 \right] \tag{21}$$

Individual Versus Collective Ugliness

Equation (21) gives the collective ugliness of the entire pattern. It is also sometimes useful to know how much (or little) an individual pulse contributes to the overall ugliness of a pattern.

Let $u(i)$ be the individual ugliness value of pulse i in pattern P. It turns out that $u(i)$ can be computed by taking the variance of the forward distances between pulse i and all the other pulses in the pattern. Using our previously developed nomenclature, we have:

$$u(i) = \frac{1}{n-1} \sum_{j=1}^{n-1} \left(\delta_j(i) - \overline{\delta(i)} \right)^2 \tag{22}$$

where:

$$\overline{\delta(i)} = \frac{1}{n-1} \sum_{j=1}^{n-1} \delta_j(i) \tag{23}$$

For an example, consider one of our previous patterns, 01110101, with P=(1,2,3,5,7). Using equation (22) on this pattern give us:

$u(0) = 3.6875$
$u(1) = 5.0000$
$u(2) = 3.6875$
$u(3) = 2.1875$
$u(4) = 2.1875$

We see that the ugliest pulse in the pattern is P_1=2, which corresponds to the pulse in slot 2 (the third slot) of the pattern. P_0=1 and P_2=3 are the next ugliest pulses (owing to their proximity to P_1). P_3=5 and P_4=7 are the least ugly pulses in the pattern. And in fact, if you remove P_1 from the pattern, you get 01010101 – which is a perfectly even pattern (Ugliness = 0).

Acknowledgements
The majority of the work reported on in this Tech Note was performed during the summer of 1990 at the Los Alamos Neutron Scattering Center (LANSCE) by Eric Bjorklund (PI), Joseph Cavanaugh, and Jose Antonio Castaneda Becerra Solorio Cabello Lopez.

The work on individual ugliness, and the application to the SNS timing system was performed during the fall of 2002.

References
[1] E.Bjorklund, “Algorithms for Optimally Distributed Timing Pulses”, LANL Report, LA-UR-89-3558, (1989).
[2] E.Bjorklund, “The Theory of Rep-Rate Pattern Generation in the SNS Timing System”, SNS ASD Tech Note, SNS-NOTE-CNTRL-99, (2003).