Centripidity

Pulse Train

Introduction

The Pulse Train is a random trigger/gate generator that provides a great deal of control over the nature of the randomness.

Trigger pulses are output with a random inter-arrival time that follows a user-chosen probability distribution with adjustable parameters.

The available distributions are:

- Fixed Value
- Uniform
- Triangular
- Exponential
- Erlang
- Normal
- Log-Normal
- Weibull
- Gamma
- Beta
- Binomial

In addition to the trigger output there is a gate output which rises at the same time as the trigger but stays open for a period determined by a different, closed probability distribution. The available gate time distributions are:

- Fixed Value
- Uniform
- Triangular
- Beta

The sequence of triggers is repeatable and storing of initial conditions is possible using the Phone Bank module. Starting and stopping the Pulse Train is possible either manually or via the application of external trigger pulses.

The Interface

The figure below highlights all of the main controls and displays available on the Pulse Train.

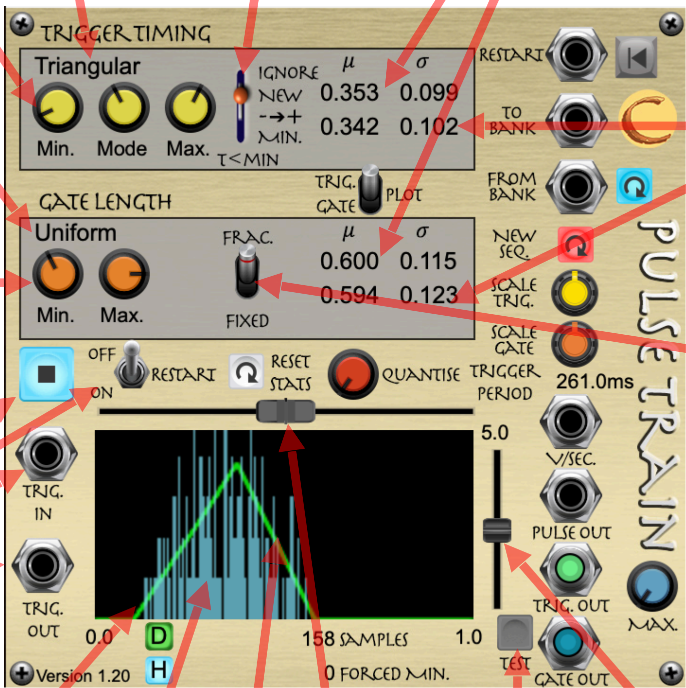

Trigger Distribution
This drop-down menu allows choice of the distribution to be used.

Trigger Distribution Controls
These knobs, up to three may appear, specify the parameters of the statistical distribution that controls trigger generation.

Gate Distribution
This drop-down menu allows choice of the distribution to be used to control gate time.

Gate Distribution Controls
These knobs, up to three may appear, specify the parameters of the statistical distribution that controls gate duration.

Run Controls
Start and stop the generation of triggers and gates either via a button press or an incoming trigger. Using the button also generates a trigger output.

The Restart switch determines whether the next Play will restart or resume the pulse sequence.

Low Value Switch
For many distributions the user may specify a minimum value that can be produced and this switch determines what will happen if the given distribution should produce a value below that minimum.

Theoretical Stats
The expected mean and standard deviation for the chosen distribution and parameter values.

Sample Stats
The mean and standard deviation of the trigger and gate times generated.

Gate Mode
In the upper position the gate length generated will represent a fraction of the trigger inter-arrival time. So, if the time to the next trigger is 0.6 seconds and the gate time is 0.25, the gate will stay open for 0.15 sec.

In the lower position the generated gate time is the number of seconds that the gate will stay open.

Distribution Plot
The chosen probability density function is plotted here in green.

Histogram Plot
A plot (in grey) of the samples so far chosen from the distribution.

Horizontal Scale
Sets the duration of the horizontal plot axis.

Density Function Plot Area

Test Button
In order to get a clearer idea of what the long-term histogram will look like, engaging this button will generate 100,000 trigger intervals to fill out the histogram. The generated trigger times are otherwise ignored by the module. They do not result in the generation of trigger pulses.

This can be particularly useful to see the effect of altering the minimum interval time in accordance with the Low Value Switch.

Histogram Scale
If the sample data is not from pure distribution, if for example a minimum generated value has been set, the histogram may not directly scale with the plotted distribution. In this case, the histogram can be scaled accordingly to check how well it fits.

This control effects only the vertical scale of the histogram, not the plotted density function.

Inter-Arrival Time
This display shows the current inter-arrival time for the next trigger/gate.

Plot Switch
Determines whether the plot area will display the trigger or gate distribution or (in centre position) nothing.

Restart
A button press or a trigger pulse will reset the Pulse Train to its initial condition and, assuming no controls have been changed, reproduce the same pulse sequence.

Inter-Arrival Time as Voltage
The voltage presented here is numerically equivalent to the trigger inter-arrival time in seconds.

Phone Bank Interface
See the Phone Bank manual for more details.

Reset Stats
Clears out all samples and restarts building the histogram.

New Sequence
Reinitialise the Pulse Train to produce a new, different sequence of pulse times.

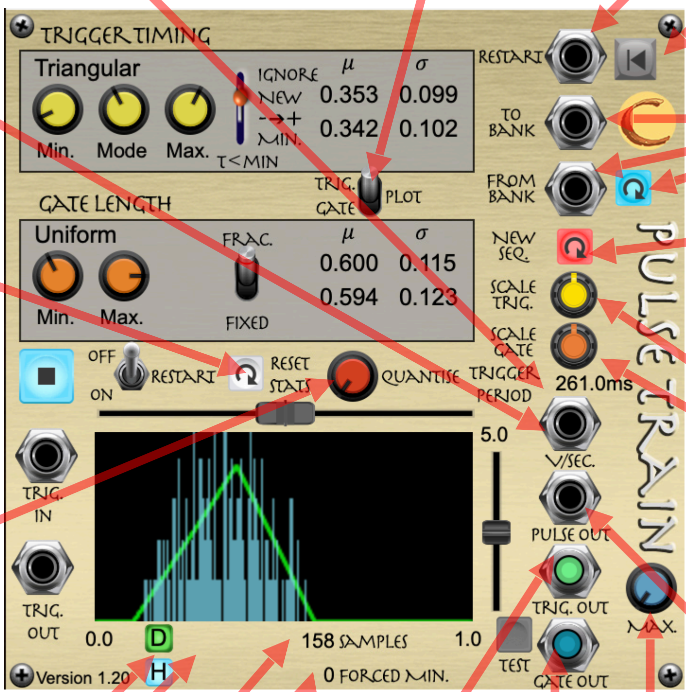

Output Scale Controls
The timing of trigger inter-arrivals and gate durations generated by the distributions can be scaled up or down with these knobs.

Quantise Knob
The triggers generated are statistically random but often it is useful to generate triggers/gates that have an underlying pulse. If this knob is raised above the *Off* position then all generated inter-arrival times and gate times will be an integer multiple of the quantisation time set here.

Pulse Out
When quantising is switched on this output sends out a pulse with a period equal to the quantisation time.

Trigger and Gate Outputs

Display Buttons
Turn the density function and histogram plots on and off independently.

Forced Min. Count
When the Low Value Switch is in MIN. mode this counter increments every time a value below the minimum allowed time is generated and replaced with that minimum allowed value.

Max Control
At its minimum setting this control is off. When raised it sets, in ms, the maximum allowed time between triggers.

If the chosen distribution produces an inter-arrival time greater than this value it is rejected and a new time is generated.

Samples Generated
Displays the total number of trigger inter-arrival times since the most recent reset.

Distribution Parameters

The plotted graph for each of these distributions is the corresponding probability density function. To interpret these graphs the important point to note is that the area under the graph between two specific time values is the probability that the next generated time value will lie within that range.

The following descriptions are largely about the technical nature of each distribution. They are provided for those who are interested but are in no way necessary for the operation of the Pulse Train.

Constant

This distribution has only the single parameter labelled as *Value*. Unsurprisingly this distribution repeatedly produces the same inter-arrival or gate time - that which is set as its value.

Uniform

The uniform distribution has two parameters, *Min.* And *Max.*. The distribution produces random numbers between these two parameters with equal probability across the range.

Triangular

The three parameters of this distribution specify the two base points of the triangle and the peak (*Mode*). Values close to the mode will occur more frequently than those closer to the lower corners of the triangle.

Exponential

This single parameter distribution is good for simulating many natural arrival phenomenon such as the time between customers arriving at a bank or photons striking a camera sensor under even illumination. The single control sets the average value that will be produced over time.

In this implementation there is a second parameter labelled *Min.*. This parameter allows the user to stop generated times from being too short. If the distribution should generate a time lower than this limit the result will be dependent upon the setting of the *Low Value Switch* which is described below.

Erlang

This distribution is a sequence of Exponential events where the *Beta* parameter is the mean of each of the corresponding Exponentials and the *Order* specifies the number of exponential events in the sequence.

This distribution also has a *Min.* Parameter the effect of which is described below in the *Low Value Switch* section.

Normal

The normal distribution is a good approximation in the variability of many naturally occurring phenomenon. There are two parameters: *Mean* and *Standard Deviation*. The standard deviation is a measure of how wide the distribution is. 68.2% of generated values will lie within 1 standard deviation and 99.7% within 3 standard deviations of the mean.

Note that this distribution is open in both directions, meaning that it can produce both negative and positive values without limit.

This distribution also has a *Min.* Parameter the effect of which is described below in the *Low Value Switch* section.

Log-Normal

This is just the natural logarithm of the normal distribution and, as a consequence, it is non-symmetric and, although it has no upper bound, is limited to producing positive values only.

This distribution also has a *Min.* Parameter the effect of which is described below in the *Low Value Switch* section.

Weibull

The Weibull distribution is used in reliability modelling and is a good statistical representation for the failure time of a device or machine. It will not generate values below zero but it has no positive bound. The first parameter, *Beta*, controls the scale of the distribution while the second (*Alpha*) parameter determines the overall shape.

This distribution also has a *Min.* Parameter the effect of which is described below in the *Low Value Switch* section.

Gamma

Often used to model the distribution of times required to complete a task, this distribution will not produce negative numbers but has no upper bound. The first parameter, *Beta*, controls the scale of the distribution while the second (*Alpha*) parameter determines the overall shape.

This distribution also has a *Min.* Parameter the effect of which is described below in the *Low Value Switch* section.

Beta

Both of the parameters of the beta distribution (*Alpha* 1 and *Alpha* 2) are shape parameters and so it can take a variety of shapes but, having no scale parameter, it is bounded within the range 0 to 1.

Binomial

The binomial distribution differs from all of the other distributions in that it is discrete. That is, it cannot produce values over a continuous range but only a limited number of specific values.

The first parameter of this distribution, *Period*, can be thought of as a timing clock. The values produced will always be multiples of this period. The binomial distribution is an even pulse generator except that the second parameter, *Prob.*, is the probability that the pulse will actually appear in each period.

So, the inter-arrival time of each pulse will always be an integer multiple of the period set by the first parameter.

Other Parameters

Low Value Switch (LVS)

To prevent the time between triggers from being too short those distributions that need it have a *Min.* Parameter which can limit how low a value the distribution can generate. However, the effect of this parameter is dependent upon the setting of another switch: the *Low Value Switch*, labelled on the interface as T<MIN.

- In the upper position (IGNORE), the *Min.* Parameter will be ignored and the distribution generated value will be used. However, if this value is less than the sample period (2.0833 x 10^{-5}), or in the case of the Normal distribution, negative, the trigger will also be ignored and a new value generated.
- In the NEW position the generated value will be rejected and a new random variate sought until it is above or equal to the specified minimum.
- In the -➔+ position any produced value that is below the minimum will be reflected back above the minimum. This alteration to the generated values will distort the distribution and hence effect the histogram display and the sample statistics.
- In the MIN. Position, any value below the specified minimum will be set equal to the specified minimum. As a consequence, that low value will be disproportionately represented in the sample and hence will distort the histogram and the sample statistics.

The effect of this can be viewed in the histogram display. The following diagrams show the resultant histogram for a Normal distribution with $\mu = 0.4$ and $\sigma = 0.1$, in each case wth the Min. Knob set to 0.3 seconds.

With the LVS set to IGNORE.

In this mode the Min. Knob is ignored and so the histogram closely follows the theoretical probability density curve.

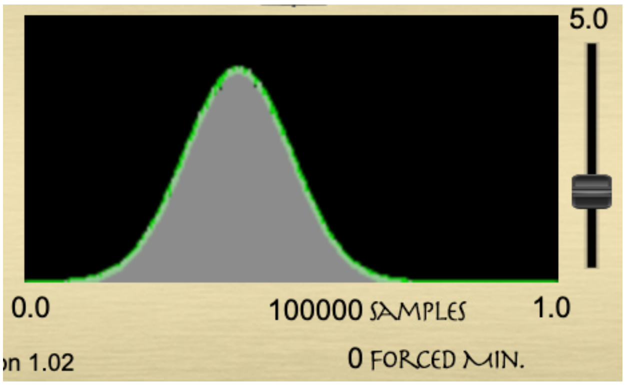

With the LVS set to NEW.

In this mode the Min. Knob is enforced by rejecting all values that fall below the minimum and generating a new interval time. This continues until the generated value is larger than the set minimum and so the histogram closely follows the theoretical probability density curve above that minimum.

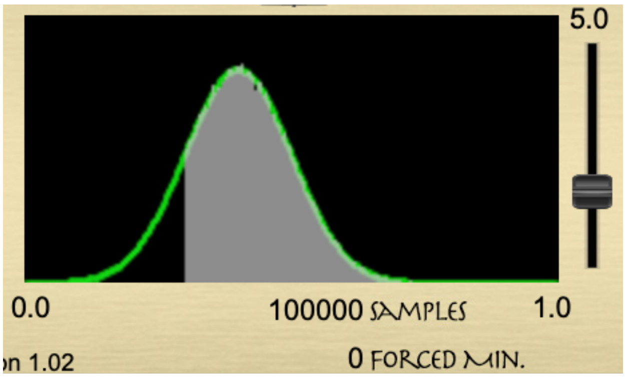

With the LVS set to -➔+.

In this mode, any generated value that is below the minimum will be reflected back across the minimum value. As an example, with the minimum value set to 0.3 a generated value of 0.2 would be returned as 0.4; a generated 0.05 would become 0.55. This causes the histogram to bulge upwards for values close to the minimum.

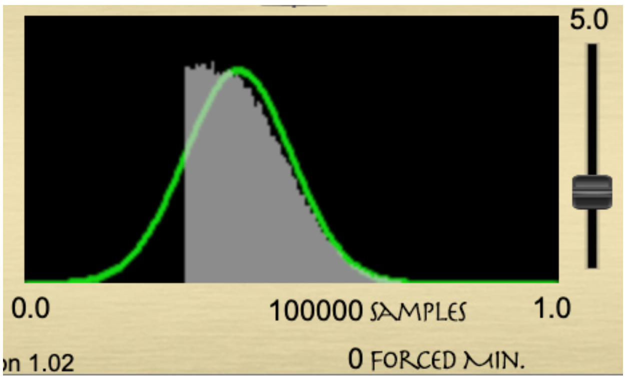

This has the effect of altering the vertical scale of the histogram relative to the density plot but adjusting the vertical scale slide can compensate for this. The slide has no effect on the timing data it just makes it possible to align the two plots for easier comparison. This is particularly useful when the "reflected" bulge in the histogram goes off scale.

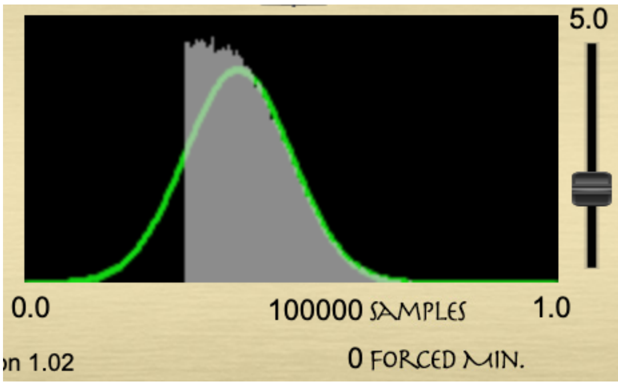

With the LVS set to MIN.

This will cause a large spike at the minimum value which is quite likely to go off scale in the histogram plot. Once again, adjusting the histogram's vertical scale can make this possible to view better. In this mode the *Forced Min.* counter below the display will give a real indication of how many times a sample has been raised to meet the minimum allowed value.

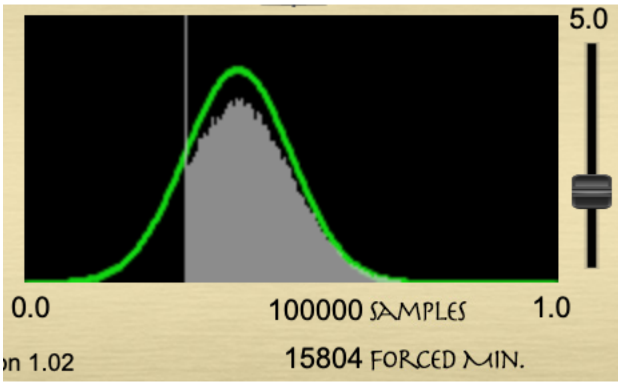

Vertical Scale = 1.0

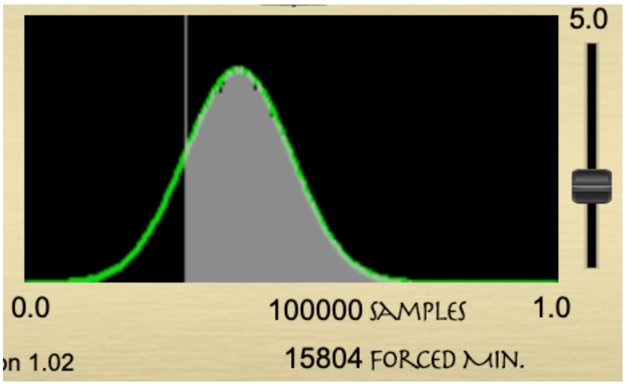

Vertical Scale = 1.17

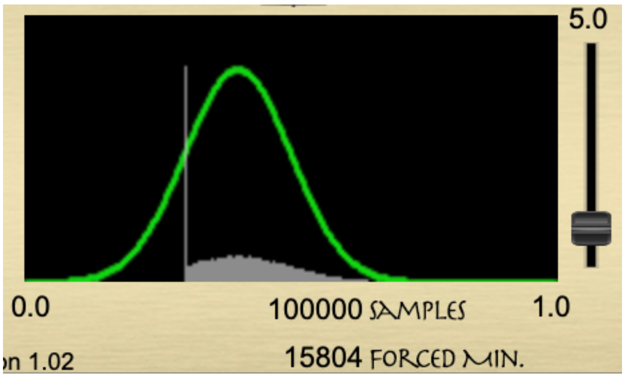

Vertical Scale = 0.14

Quantise Knob

As mentioned above, the binomial, and indeed the constant, distribution allow all generated inter-arrival times for triggers to be the integer multiple of an underlying pulse. In fact, this is possible for any of the distributions using the *Quantise Knob*.

The nine continuous distributions can generate inter-arrival times anywhere within their range but if you want all triggers to line up with an underlying pulse timing, this knob will achieve that.

In its *Off* position it has no effect but once elevated to a specific pulse duration (selectable between 1 - 1000ms) all subsequent triggers are guaranteed to be generated at a time that is some multiple of that underlying quantised pulse value.

Frac./Fixed Switch

Each gate output opens to correspond with the appearance of a trigger pulse at the trigger output. The chosen gate distribution will determine the length of time that the gate is open. However, the generated gate time has two separate interpretations and this switch determine which interpretation is active.

- FRAC. - the generated time represents the proportion of the trigger inter-arrival time for which the gate will be open. If this generated time reaches or exceeds one then the gate will not close during this pulse.

- FIXED - the generated value is the time, in seconds, that the gate will remain open. If this time is greater than the trigger inter-arrival time then the gate will not close during this pulse.