

REDUCING HARMONIC DISTORTION BY DITHERING

PRIOR APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Patent Application No. 62/914,487 filed on October 13, 2019, which is incorporated by reference in its entirety.

[0002] The prior application, U.S. Application No. 16/600,496 filed on October 12, 2019, is incorporated by reference in its entirety.

FIELD OF THE DISCLOSURE

[0003] The present disclosure relates generally to signal processing that produces a substantially error-free signal preserving the amplitude and phase relative to a constant base frequency while allowing the carrier frequency to vary.

BACKGROUND

[0004] Carrier frequency is defined herein as the consequential instantaneous frequency of the output digital signal pulses. Base frequency is defined herein as the center frequency that is described by an unmoving or constant phase signal. It is known that any phase change in the signal also constitutes a frequency shift. In this case that is realized by the carrier frequency shifting away from the base frequency. To preserve compatibility with the input intended for the elements of a phased array, the input data to the method is considered to be the phase and duty cycle of a pulse-width modulation. This is measured with respect to a steady reference signal that is a fixed source at the base frequency.

[0005] Phased array systems rely on the production of an exact carrier frequency to function. To simplify systems, it is often assumed that the carrier frequency is emitted during all relevant times so that the system can be treated as time

invariant. This time invariance is necessary for the input signals to the array element transducers to be treated as complex values.

[0006] Generating a constant frequency pulse-width modulated (PWM) digital signal with a given phase offset for all relevant times is trivial. But changing the state of a phased array system often involves changing the phase angle of the elements, which violates the time-invariance requirement. This results in many side-effects, including a shift in frequency. Since the digital signal generation assumes that the base frequency (the frequency with which the primitive phase angles are specified relative to) is equal to the carrier frequency for all relevant times, this causes errors in the digital signals output to each array element transducer. Thus, it is necessary for the development of a signal generation system that is capable of producing a digital signal using the free selection of amplitude and phase. This is used to produce a substantially error-free signal that preserves the amplitude and phase relative to a constant base frequency while allowing the carrier frequency to vary.

SUMMARY

[0007] The hardware-efficient method of generating the up-sampling of the phase represented as the evaluation of high-order polynomial interpolant is novel. The aim is to produce a PWM output that respects and correctly interprets changes in frequency while also preserving absolute phase and phase changes. Without loss of generality, this technique may be also restated with phase delays that produces a “sign flip” in angle from the technique described.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The accompanying figures, where like reference numerals refer to identical or functionally similar elements throughout the separate views, together with the detailed description below, are incorporated in and form part of the specification, serve to further illustrate embodiments of concepts that include the claimed invention and explain various principles and advantages of those embodiments.

[0009] Figure 1 shows geometric behavior of a system having $3/2$ of a base frequency.

[0010] Figure 2 shows geometric behavior of a system having exactly the base frequency.

[0011] Figure 3 shows geometric behavior of a system having a slowly increasing phase relative to the base frequency.

[0012] Figure 4 shows geometric behavior of a system having an arbitrary function of phase angle relative to the base frequency.

[0013] Figures 5A–5D show visualizations of the key parameters for a 25% duty cycle linear phase and duty realization of the disclosed method.

[0014] Figures 6A–6D show visualizations of the key parameters for a 75% duty cycle linear phase and duty realization of the disclosed method.

[0015] Skilled artisans will appreciate that elements in the figures are illustrated for simplicity and clarity and have not necessarily been drawn to scale. For example, the dimensions of some of the elements in the figures may be exaggerated relative to other elements to help to improve understanding of embodiments of the present invention.

[0016] The apparatus and method components have been represented where appropriate by conventional symbols in the drawings, showing only those specific details that are pertinent to understanding the embodiments of the present invention so as not to obscure the disclosure with details that will be readily apparent to those of ordinary skill in the art having the benefit of the description herein.

DETAILED DESCRIPTION

[0017] I. Introduction

[0018] The moving phase angle may be considered as an equivalent formulations for a phase-frequency modulated wave:

$$\cos\left(\omega t + \int \frac{d\theta'(t)}{dt} dt + \theta\right) = \cos(\theta'(t) + \omega t) = \cos\left(\theta + \int \omega'(t) dt\right),$$

where $\theta'(t)$ is a time-dependent function of phase and $\omega'(t)$ is a time-dependent function of frequency. It can be seen that $d\theta'(t)/dt$ is a measure of deviation of the carrier frequency from the base frequency. This can be simplified by normalizing both angle and ω (divide through by the base frequency and 2π radians, θ now being measured in revolutions), yielding $\omega = 1$.

[0019] Describing the phase delay $\theta'(t)$ may be achieved by interpolating phase offsets generated in subsequent base frequency steps by a polynomial, since it is beneficial for the frequency to be defined and continuous on the endpoints. The frequency is defined as:

$$\frac{d\theta'(t)}{dt} + \omega = \omega'(t),$$

where the first-time derivatives of phase angle also contribute to the instantaneous carrier frequency and thus form two derivative constraints:

$$\frac{d\theta'(0)}{dt} + \omega = \omega'(0) = \omega + \theta_0 - \theta_{-1},$$

$$\frac{d\theta'(1)}{dt} + \omega = \omega'(1) = \omega + \theta_1 - \theta_0.$$

The two endpoints of the interval in angle also have further constraints:

$$\theta'(0) = \theta_0,$$

$$\theta'(1) = \theta_1,$$

which together with the constraints on carrier frequency make four in total. This necessitates a cubic polynomial interpolation for this level of continuity. As shown, defining $\omega'(0)$ and $\omega'(1)$ can be achieved using backwards differences, thus limiting the number of samples required in the future direction and reducing latency. This also reduces the total number of immediately available samples

required from four to three precomputed samples of the phase angle and duty cycle of the intended signal.

[0020] The cubic form of the interpolating spline polynomial formed from backwards differences is:

$$\theta'(t) = (-\theta_{-1} + 2\theta_0 - \theta_1)t^3 + (2\theta_{-1} - 4\theta_0 + 2\theta_1)t^2 + (\theta_0 - \theta_{-1})t + \theta_0,$$

which is repeated for every interval.

[0021] Further, the phase may be also represented by a lower degree polynomial. Although this would imply sacrificing some of the continuity conditions, the reasonable approach is to produce discontinuities in frequency anyway (but importantly, phase continuity is preserved as only the time derivatives of the phase are discontinuous). Even with frequency discontinuities, the technique using this interpolant enjoys a significant accuracy improvement over the standard technique. The linear interpolant for such a method may be stated as:

$$\theta'(t) = (\theta_1 - \theta_0)t + \theta_0.$$

[0022] Although the complexity of the implementation increases, higher order interpolation polynomials may equally be used without loss of generality. The on time of a digital signal is described by the duty cycle, which is assumed proportional to the amplitude of the signal. This motivates the name “pulse-width modulation.” This can be realized here by adding an interpolation on the duty cycle value Δ of the signal encoded as a pulse-width percentage at the base frequency:

$$\Delta'(t) = (\Delta_1 - \Delta_0)t + \Delta_0.$$

Defining the output signal going into the element as a digital approximation to:

$$\cos\left(\omega t + \int \frac{d\theta'(t)}{dt} dt + \theta\right) = \cos(\theta'(t) + \omega t),$$

so a time-varying θ phase offset with respect to the base frequency may also be viewed as a deviation from the base signal frequency ω , effectively $d\theta'(t)/dt$. To

search for the locations of the pulses, zeroes (also multiples of 2π) of the angle input to the cosine function must be found. These correspond to peaks in the wave and high points in the digital signal. To achieve this, both angle and ω are normalized (divide through by the frequency and 2π radians, all θ now being measured in revolutions), yielding $\omega = 1$. Therefore, the condition being searched for is:

$$\omega t - \theta'(t) = t - \theta'(t) = 0.$$

This describes the center of the pulse at each step.

[0023] To find the extent of the pulse around the center point, the value $|t - \theta'(t)|$ is computed. If it is smaller than a given value representing an amplitude, then the point in time is within the pulse, in the high region of the digital signal. Otherwise, the point in time is outside the pulse and in the low region of the digital signal.

[0024] II. Hardware Techniques to Create Suitable PWM Output

[0025] The hardware-efficient method of generating the up-sampling of the phase represented as the evaluation of high-order polynomial interpolant is also novel. The aim is to produce a PWM output that respects and correctly interprets changes in frequency while also preserving absolute phase and phase changes. Without loss of generality, this technique may be also restated with phase delays that produces a “sign flip” in angle from the technique described.

[0026] The moving phase angle may be considered as an equivalent formulations for a phase-frequency modulated wave:

$$\cos\left(\omega t + \int \frac{d\theta'(t)}{dt} dt + \theta\right) = \cos(\theta'(t) + \omega t) = \cos\left(\theta + \int \omega'(t) dt\right),$$

where $\theta'(t)$ is a time-dependent function of phase and $\omega'(t)$ is a time-dependent function of frequency. It can be seen that $d\theta'(t)/dt$ is a measure of deviation of the carrier frequency from the base frequency. This can be simplified by

normalizing both angle and ω (divide through by the base frequency and 2π radians, θ now being measured in revolutions), yielding $\omega = 1$.

[0027] Describing the phase delay $\theta'(t)$ may be achieved by interpolating phase offsets generated in subsequent base frequency steps by a polynomial, since it is beneficial for the frequency to be defined and continuous on the endpoints. The frequency is defined as:

$$\frac{d\theta'(t)}{dt} + \omega = \omega'(t),$$

where the first-time derivatives of phase angle also contribute to the instantaneous carrier frequency and thus form two derivative constraints:

$$\frac{d\theta'(0)}{dt} + \omega = \omega'(0) = \omega + \theta_0 - \theta_{-1},$$

$$\frac{d\theta'(1)}{dt} + \omega = \omega'(1) = \omega + \theta_1 - \theta_0.$$

The two endpoints of the interval in angle also have further constraints:

$$\theta'(0) = \theta_0,$$

$$\theta'(1) = \theta_1,$$

which together with the constraints on carrier frequency make four in total. This necessitates a cubic polynomial interpolation for this level of continuity. As shown, defining $\omega'(0)$ and $\omega'(1)$ can be achieved using backwards differences, thus limiting the number of samples required in the future direction and reducing latency. This also reduces the total number of immediately available samples required from four to three precomputed samples of the phase angle and duty cycle of the intended signal.

[0028] The cubic form of the interpolating spline polynomial formed from backwards differences is:

$$\theta'(t) = (-\theta_{-1} + 2\theta_0 - \theta_1)t^3 + (2\theta_{-1} - 4\theta_0 + 2\theta_1)t^2 + (\theta_0 - \theta_{-1})t + \theta_0,$$

which is repeated for every interval.

[0029] Further, the phase may be also represented by a lower degree polynomial. Although this would imply sacrificing some of the continuity conditions, the reasonable approach is to produce discontinuities in frequency anyway (but importantly, phase continuity is preserved as only the time derivatives of the phase are discontinuous). Even with frequency discontinuities, the technique using this interpolant enjoys a significant accuracy improvement over the standard technique. The linear interpolant for such a method may be stated as:

$$\theta'(t) = (\theta_1 - \theta_0)t + \theta_0.$$

[0030] Although the complexity of the implementation increases, higher order interpolation polynomials may equally be used without loss of generality. The on time of a digital signal is described by the duty cycle, which is assumed proportional to the amplitude of the signal. This motivates the name “pulse-width modulation.” This can be realized here by adding an interpolation on the duty cycle value Δ of the signal encoded as a pulse-width percentage at the base frequency:

$$\Delta'(t) = (\Delta_1 - \Delta_0)t + \Delta_0.$$

Defining the output signal going into the element as a digital approximation to:

$$\cos\left(\omega t + \int \frac{d\theta'(t)}{dt} dt + \theta\right) = \cos(\theta'(t) + \omega t),$$

so a time-varying θ phase offset with respect to the base frequency may also be viewed as a deviation from the base signal frequency ω , effectively $d\theta'(t)/dt$. To search for the locations of the pulses, zeroes (also multiples of 2π) of the angle input to the cosine function must be found. These correspond to peaks in the wave and high points in the digital signal. To achieve this, both angle and ω are normalized (divide through by the frequency and 2π radians, all θ now being measured in revolutions), yielding $\omega = 1$. Therefore, the condition being searched for is:

$$\omega t - \theta'(t) = t - \theta'(t) = 0.$$

This describes the center of the pulse at each step.

[0031] To find the extent of the pulse around the center point, the value $|t - \theta'(t)|$ is computed. If it is smaller than a given value representing an amplitude, then the point in time is within the pulse, in the high region of the digital signal. Otherwise, the point in time is outside the pulse and in the low region of the digital signal.

[0032] Figures 1 through 4 geometrically demonstrate how testing that this value is less than $\Delta'(t)/2$ generates the appropriate pulse.

[0033] Figure 1 shows the geometrical behavior 400 for the edge case of $3/2$ of the base frequency. In this graph, Δ is $\Delta(t)$, the y-axis 410 represents normalized angle (in revolutions) θ , and the x-axis 420 represents normalized time (in base frequency periods) t . This Figure 1 is a geometric interpretation of the PWM generation when applied to a slowly decreasing phase (with derivative $-1/2$, negative slope down and to the right), relative to the base frequency represented by the diagonal phase lines 440a, 440b, 440c. The diagonal phase lines in this and in all other figures herein are graphed on a periodic domain and therefore “wrap around” the plots as shown. Thus, many of the apparent discontinuities represent the same curve.

[0034] The distance between the two sets of repeating curves crosses the threshold where it is less than $\Delta/2$ (defined as half the duty cycle quantity) distance in a number of places 450a, 450b, 450c, 450d, 450e, 450f, 450g, 450h that repeat in time. These two sets of curves are the constant phase versus timelines 460a, 460b, 460c, 460d, 460e ($\theta = t$ or $\theta = \omega t$, but wrapped around in rotations and base frequency periods since ω is normalized to one). This travels up and to the right of the diagram that represent the base frequency with zero phase offset behavior. The interpolated phase curves ($\theta'(t)$) that represent the desired behavior that are an addition to this signal in phase 440a, 440b, 440c. Where the two curves “match” in phase closely enough (less than $\Delta/2$), these

regions represent the pulse parts of the pulse signals 430. The dashed vertical lines projected from the $\Delta/2$ distance factors 450a, 450b, 450c, 450d, 450e, 450f, 450g, 450h show the places on the PWM signal 430 where the binary state is changed inducing pulse edges due to the $\Delta/2$ distance factor being reached.

[0035] The constant phase versus timelines ($\theta = t$ or $\theta = \omega t$) travelling up and to the right of the diagram that represent the base frequency with zero phase offset behavior are repeated for every period of the base frequency. The repetition in the vertical direction shows that it is true for all integer numbers of rotations in angle. Thus, it is true even considering numerical wrap-around of the counters used to implement the method. This generates a PWM signal with a carrier frequency that is three-halves the base frequency (where the frequency multiplier is obtained by subtracting the instantaneous derivative of the interpolated phase lines $\theta'(t)$ ($-\frac{1}{2}$) from the derivative of the constant phase versus timelines $\theta = \omega t$ (1), so $1 - (-\frac{1}{2}) = \frac{3}{2}$). At the bottom is the final digital signal 430 that is to drive the element made up of all of the points where the two sets of curves are less than $\Delta/2$ distance apart.

[0036] Figure 2 shows the geometrical behavior 500 for exactly the base frequency. In this graph, Δ is $\Delta(t)$, the y-axis 510 represents normalized angle (in revolutions) θ , and the x-axis 520 represents normalized time (in base frequency periods) t . Shown is a geometric interpretation of the PWM generation when applied to a flat constant phase angle $\theta'(t)$ (horizontal lines with derivative zero in time) that does not change relative to the base frequency represented by the diagonal phase lines 540a, 540b, 540c.

[0037] The distance $\Delta/2$ 550a, 550b, 550c, 550d, 550e, 550f again represents the transition points between the two states in the pulse signal. Thus, the two curves cross over exactly once per base frequency period because the interpolated phase curve is horizontal and represents a constant phase angle. This generates a PWM signal with a carrier frequency that is exactly equal to the base frequency (where the frequency multiplier is again obtained by subtracting the instantaneous derivative of the interpolated phase lines $\theta'(t)$ (0) from the derivative of the

constant phase versus timelines $\theta = \omega t$ (1), so $1 - 0 = 1$). The dashed lines show the pulse edges in the pulsed signal. At the bottom is the final digital signal 530 that is to drive the element made up of all of the points where the two sets of curves are again less than $\Delta/2$ distance apart.

[0038] Figure 3 shows the geometrical behavior 600 for the edge case of $1/2$ the base frequency. In this graph, Δ is $\Delta(t)$, the y-axis 610 represents normalized angle (in revolutions) θ , and the x-axis 620 represents normalized time (in base frequency periods) t . The distance $\Delta/2$ 650a, 650b again represents the transition points between the two states in the pulse signal.

[0039] Shown is a geometric interpretation of the PWM generation when applied to an increasing phase $\theta'(t)$ (with derivative $1/2$), relative to the base frequency represented by the diagonal phase lines 640a, 640b, 640c. This generates a PWM signal with a carrier frequency that is half the base frequency (where the frequency multiplier is obtained by subtracting the instantaneous derivative of the interpolated phase lines $\theta'(t)$ ($+1/2$) from the derivative of the constant phase versus time lines $\theta = \omega t$ (1), so $1 - (+1/2) = 1/2$). At the bottom is the final digital signal 630 that is to drive the element made up of all of the points where the two sets of curves are again less than $\Delta/2$ distance apart.

[0040] Figure 4 shows the geometrical behavior 700 as to how an example interpolated function where the gradient and thus frequency changes significantly over time fits into this geometric description. In this graph, Δ is $\Delta(t)$, the y-axis 710 represents normalized angle (in revolutions) θ , and the x-axis 720 represents normalized time (in base frequency periods) t . Shown is a geometric interpretation of the PWM generation when applied to a more arbitrarily defined function of phase angle, relative to the base frequency represented by the diagonal phase lines $\theta'(t)$ or $\theta = \omega t$ 740a, 740b, 740c, 740d, 740e, 740f, 740g, 740h, 740j, 740k. The distance $\Delta/2$ 750a, 750b, 750c, 750d, 750e, 750f, 750g, 750h, 750j, 750k, 750m, 750n, 750p, 750q, 750r, 750s, 750t, 750u, 750v, 750w again represents the edges in the pulse signal. But here while they represent the same $\Delta/2$ distance on the y-axis, they correspond to varying pulse length on the x-axis.

The wavy horizontal lines 745a, 745b, 745c are the interpolated phase lines $\theta'(t)$ in this example.

[0041] The variation in the derivative of $\theta'(t)$ moves between positive derivative that generates longer pulses at a lower frequency and negative derivative that generates shorter pulses at a higher frequency. This is due to the crossings between y-axis distances smaller than $\Delta/2$ and larger than $\Delta/2$ changing their relative distance apart. At the bottom is the final digital signal 730 that is to drive the element, wherein pulse edges are induced when the signal y-axis distance crosses the $\Delta/2$ threshold.

[0042] It can also be proven that the duty cycle value $\Delta'(t)/2$ when used in this way scales appropriately with frequency for this scheme.

[0043] III. Conditional Test and Dithering

[0044] In order to obtain reduced harmonic distortion, it is beneficial to introduce dithering to alleviate artifacts generated by the discretization of the signal into precise digital sub-intervals (ticks). But in order to do so and not introduce further problems, it is necessary to ensure that each dithering operation on each rising and falling pulse edge on each element is uncorrelated. This is difficult to achieve in a hardware system, so some rearrangement of the scheme is required.

[0045] Having obtained $\theta'(t)$, $\Delta'(t)$ and t with counters, where $\theta'(t)$ and $\Delta'(t)$ can be arbitrary degree polynomials, these can now be used to evaluate:

$$|t - \theta'(t)| < +\Delta'(t)/2.$$

to test when this key condition that drives which Boolean state the pulse output is emitting is true when wrap-around is allowed in the function quantities, as described earlier. For efficiency in hardware systems and in order to control the beginning and end edge of each pulse individually, these are worked into two counters that have different tests applied to them as the negative and positive parts of the above test. This may be expressed as:

$$t - \theta'(t) \geq -\frac{\Delta'(t)}{2} \text{ and } t - \theta'(t) < +\frac{\Delta'(t)}{2}.$$

[0046] As only the sign is important to conduct this test, an extra bit may be added to the counters to remove the divide by multiplying through by two:

$$2t - 2\theta'(t) \geq -\Delta'(t) \text{ and } 2t - 2\theta'(t) < +\Delta'(t).$$

[0047] Then the counters are rearranged to make this a sign test:

$$\begin{aligned} & -\left((2t - 2\theta'(t)) + \Delta'(t)\right) \leq 0 \text{ and } (2t - 2\theta'(t)) - \Delta'(t) < 0, \\ & -\left((2t - 2\theta'(t)) + \Delta'(t)\right) - \varepsilon \leq -\varepsilon \text{ and } (2t - 2\theta'(t)) - \Delta'(t) < 0, \\ & -\left((2t - 2\theta'(t)) + \Delta'(t)\right) - \varepsilon < 0 \text{ and } (2t - 2\theta'(t)) - \Delta'(t) < 0, \end{aligned}$$

where the physical manifestation of ε is generally just the least significant bit and the space of numbers in which the counter operates is reworked to be signed.

[0048] As a result, two final counters may be constructed for these quantities, wherein the sign bits of the counters denote whether or not the test succeeded as the result of the logical AND operator applied to both of their sign bits. An implementation may choose to use the AND operator for less than 50% duty cycle ($\Delta'(t) \leq \frac{1}{2}$) and the logical OR operator for greater than 50% duty cycle ($\Delta'(t) \geq \frac{1}{2}$), where at a 50% duty cycle both operators produce the same result. This is further complicated by the fact that the counters swap roles when the duty cycle crosses 50%, at over 50% duty cycle, the ‘begin’ counter controls the falling edge and the ‘end’ counter controls the rising edge.

[0049] For brevity, the following describes a degree one (linear) polynomial implementation of $\theta'(t)$ and $\Delta'(t)$. The required counters for the edges ‘begin’ and ‘end’ may be expanded into:

$$\begin{aligned} \text{begin} &:= -\left(\left((\Delta'_1 - \Delta'_0)t + \Delta'_0\right) + \left(2t - 2((\theta'_1 - \theta'_0)t + \theta'_0)\right) + \varepsilon\right), \\ \text{end} &:= -\left(\left((\Delta'_1 - \Delta'_0)t + \Delta'_0\right) - \left(2t - 2((\theta'_1 - \theta'_0)t + \theta'_0)\right)\right) \end{aligned}$$

Because this is a linear example, these may be broken down into sub-counters by evaluating and taking the derivative of each counter and evaluating at the start of the interval to determine the starting values and constants to be computed. For

greater degree polynomials, the counters may be used as described in the previous section, but for this linear example, only two initializing constants and two derivatives are required:

$$\begin{aligned} \text{begin}_0 &:= -(\Delta'_0 - \theta'_0) - \varepsilon, \\ \frac{d\text{begin}}{dt} &:= -\left((\Delta'_1 - \Delta'_0) - 2(1 - (\theta'_1 - \theta'_0))\right), \\ \text{end}_0 &:= -(\Delta'_0 + \theta'_0), \\ \frac{d\text{end}}{dt} &:= -\left((\Delta'_1 - \Delta'_0) + 2(1 - (\theta'_1 - \theta'_0))\right). \end{aligned}$$

[0050] As before, rewriting this to use real digital logic with counters and bit shifting means replacing t in the interval $[0, 1)$ with k that represents the number of discrete ticks passed since the interval began. Defining the number of ticks per interval as 2^p , this then makes the initial counter values (assuming the inputs are in the interval $[0, 1)$ and not already in fractions of 2^p , which would be in the interval $[0, 2^p - 1]$):

$$\begin{aligned} \text{begin}_0 &:= -2^{2p}(\Delta'_0 - 2\theta'_0) - 1, \\ \frac{d\text{begin}}{dt} &:= -2^p \left((\Delta'_1 - \Delta'_0) - (2 - (2\theta'_1 - 2\theta'_0))\right), \\ \text{end}_0 &:= -2^{2p}(\Delta'_0 + 2\theta'_0), \\ \frac{d\text{end}}{dt} &:= -2^p \left((\Delta'_1 - \Delta'_0) + (2 - (2\theta'_1 - 2\theta'_0))\right), \end{aligned}$$

where ε has been replaced with a digital 1, the smallest non-zero value in the representation.

[0051] These given assignments are then the initial value of the ‘begin’ and ‘end’ counters and their respective delta increments which are applied at the beginning of each digital tick, as the evaluations required are:

$$\begin{aligned} \text{begin}(t) &:= \text{begin}_0 + \frac{d\text{begin}}{dt}k, \\ \text{end}(t) &:= \text{end}_0 + \frac{d\text{end}}{dt}k. \end{aligned}$$

[0052] As described before, then applying logical bit operators to the sign bit of both ‘begin’ and ‘end’ is used to produce the final state for the pulse output.

[0053] IV. Unweighted Dither

[0054] As stated, due to artefacts introduced by the discretization into digital ticks it is necessary to add dither to the system to ameliorate the effects of the signal quantization.

[0055] To alleviate the quantization effects, reduce the sidebands and thus distortion in the output, each pulse edge has a probabilistic dither added to it. But because of the way the positions of the edges are determined parametrically, it is difficult to add dither that varies the edge position by a single tick when the frequency of each interval (or each tick for higher order polynomials) may be different over time.

[0056] Adding dither to each edge in t (whose crossing and thus edge position is determined parametrically) would require significant extra operations (at least a multiplication, which would be expensive). This would determine how to move the edge by a prescribed amount in t , likely by adding a weighted dither factor to the parametrically defined edge counter that can be viewed as moving primarily in θ .

[0057] An alternative approach, is to not weight the dither and apply it directly to the counter (that is use ‘unweighted’ dither), thereby dithering in the ‘phase angle’ space θ , rather than directly in time t . This can be viewed as dithering by moving or jittering one of the curves in Figures 1, 2, 3 and 4 in the ‘phase angle’ axis θ instead of moving or jittering the curve along the time axis t . This means that the dither is then projected through the parametric definition of the edge location. Effectively this just means that more dither is applied when the instantaneous carrier is at lower frequencies and less is applied when the instantaneous carrier is at higher frequencies, as a movement in θ generates a greater or lesser movement in t respectively in those situations.

[0058] To add this approach onto the two edges in the method, two dither values are produced: one for the ‘begin’ counter $\text{begin}(t)$ and one for the ‘end’ counter $\text{end}(t)$. These are defined as the integer values:

$$\begin{aligned}\text{begin}_\delta &:= 2(2^p \times \text{uniformRandom}([0,1))), \\ \text{end}_\delta &:= 2(2^p \times \text{uniformRandom}([0,1))),\end{aligned}$$

where the result in each case is ideally a uniformly distributed random integer value in the interval $[0, 2^{p+1} - 1]$ (but if implemented as a linear-feedback shift register (LFSR) can be reasonably $[1, 2^{p+1} - 1]$).

[0059] The evaluations are now:

$$\begin{aligned}\text{begin}_d(t) &:= \text{begin}_0 + \frac{d\text{begin}}{dt}k - \text{begin}_\delta, \\ \text{end}_d(t) &:= \text{end}_0 + \frac{d\text{end}}{dt}k + \text{end}_\delta,\end{aligned}$$

where the dithers are applied with opposite sign to each so as not to bias the final output of the device.

[0060] Expanding these $\text{begin}_d(t)$ and $\text{end}_d(t)$ out we have:

$$\begin{aligned}\text{begin}_d(t) &:= -2^{2p} \left(\Delta'_0 - 2 \left(\theta'_0 - \frac{\text{begin}_\delta}{2^{p+1}} \right) \right) - 2^p \left((\Delta'_1 - \Delta'_0) - (2 - (2\theta'_1 - 2\theta'_0)) \right) k - 1, \\ \text{end}_d(t) &:= -2^{2p} \left(\Delta'_0 + 2 \left(\theta'_0 - \frac{\text{end}_\delta}{2^{p+1}} \right) \right) - 2^p \left((\Delta'_1 - \Delta'_0) + (2 - (2\theta'_1 - 2\theta'_0)) \right) k,\end{aligned}$$

First, this effectively moves the phase by up to one tick on each edge. Second, as the expected value of a uniform distribution on the interval $[0,1)$ is a half, it is expected that this causes bias on the phase of every pulse edge by half a sample but does not bias the duty cycle. Thus, the amplitude is not modified overall.

Biasing every pulse by half a sample, as every element in the system is driven by the same electronics, does not have an effect and so can be neglected.

[0061] A final complication in implementing this approach is that since multiple rising or falling edges may occur in a single base frequency period and the edges may occur anywhere in the timeline. Thus, the time at which the random dither values begin_δ and end_δ are changed must be chosen carefully. It is necessary to find a way to reinitialize each of the random dither values begin_δ and end_δ at time where edges can be guaranteed not to fall. This ensures there are no race conditions which could cause spurious pulses if the dither amount is modified on the same clock cycle as it modifies where the edge falls. It would also be inadvisable to allow the same dither value to apply for more than one edge as it would correlate these edges.

[0062] A system that achieves this is one where it can be guaranteed that a jump in the dither values begin_δ and end_δ cannot affect the final pulse shape. This occurs during the parts of the counter cycle where the most significant bits for each counter place them outside of any region where their sign may affect the location of an edge. Assume that the values of the $\text{begin}(t)$ and $\text{end}(t)$ counters are interpreted as to be in the signed interval $[-\frac{1}{2}, +\frac{1}{2})$ or the upper bits may be neglected to place it in this range. To ensure that it is not near a sign change, the easiest method is to test and reinitialize the dither value if, and only if, the counter lies in the interval $[+\frac{1}{8}, +\frac{3}{8})$ or $[-\frac{3}{8}, -\frac{1}{8})$ to ensure sufficient time to reinitialize the random values. These intervals in particular can be determined by checking the sign bit and asserting that the next two most significant bits must XOR to a Boolean true.

[0063] A simple method for achieving this is to create a linear feedback shift register (LFSR) that embody the begin_δ and end_δ values. By only allowing the register to be clocked by one bit while the above conditions are true, which as long as the value is much shorter in bits than one quarter of the ticks in the

interval, this effectively make the jitter randomly selected for each rising and falling edge.

[0064] Turning to Figures 5A, 5B, 5C and 5D, shown are visualizations of the key parameters for a 25% duty cycle linear phase and duty realization of the method used to produce the jitter or dither effect on the output signal.

[0065] Figure 5A shows a plot 100 of $\theta = \theta'(t)$ and $\theta = t$ to show the key distance parameter $(\Delta'(t))/2$ and how its projection produces the edges in the final pulse signal. The x-axis 104 shows the again normalized time (in base frequency periods) t , while the y-axis 102 shows normalized phase angle (in revolutions) θ . The signal 105 now shows a 25% duty cycle to illustrate how the counter mechanism preserves duty cycle percentage whilst the phase and frequency varies. The dashed vertical lines 130 show how the distance test of the difference between $\theta = \theta'(t)$ and $\theta = t$ being less than $(\Delta'(t))/2$ (which is 1/8 in the illustration) maps onto a signal with a duty cycle of 25%. The horizontal wavy lines 111 are the linearly interpolated θ values from (in this embodiment) each base frequency tick, $\theta = \theta'(t)$, a function representing a sinusoidal variation in phase which gives rise to a sinusoidal variation in frequency of the final output pulse width (amplitude), phase and frequency modulated signal. The diagonal lines 110 as before show the lines $\theta = t$. The vertices 125 show the locations where the difference between $\theta = \theta'(t)$ and $\theta = t$ is zero, corresponding to the center point of the high region of the output signal. The vertical arrows 127 show the extent of the difference between the lines depicting $\theta = \theta'(t)$ and $\theta = t$ where the magnitude of this difference is $(\Delta'(t))/2$, denoting the edge of the output signal.

[0066] Figure 5B shows a plot 200 of the counters “begin(t) “ and “end(t) “ to show how the sign bits of the counters with wrap-around, when a logical AND is applied (the region that is the intersection and is between both dashed lines 232 and 234 but also at the same time between dotted lines 231 and 233 is the region which describes the high part of the final signal) produce the same digital signal. The x-axis 202 shows the again normalized time (in base frequency periods) t , but

in this figure the y-axis 204 shows now T which can be interpreted as the fractional number of periods of the output signal. The dashed diagonal lines 209 show the function “begin(t)” which is $\theta'(t) - t$ shifted by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the rising edge of the output signal. The dotted diagonal lines 211 show the function “end(t)” which is $\theta'(t) - t$ shifted in the opposite direction by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the other edge, correspondingly the falling edge of the output signal. The vertical dashed line 232 corresponds to the point where the dashed diagonal line 209 crosses $T=0$, the x-axis, which project down onto the rising edge of the output signal. The vertical dotted line 233 corresponds to the point where the dotted diagonal line 211 crosses $T = 0$, the x-axis. The vertical dashed line 220 corresponds to the point where the dashed diagonal line crosses the $T = (+/-)1/2$ which is the antipodean point to the rising edge and a location in time wherein applying a probabilistic jitter (a discontinuity) to the “begin(t)” counter cannot affect the output signal and so cannot generate runt pulses or other signal artifacts. (As used herein, “antipodean” means a point 180 degrees in phase away from a point of interest in a digital signal.) The vertical dotted line 222 corresponds to the point where the dotted diagonal line crosses the $T = (+/-)1/2$, which is the antipodean point to the falling edge and a location in time wherein applying a probabilistic jitter (a discontinuity) to the “end(t)” counter cannot affect the output signal and so cannot generate runt pulses or other signal artifacts. The signal 207 shows the output pulse width (amplitude), phase and frequency modulated signal.

[0067] Figure 5C shows a plot 300 of “begin(t)” (dashed) and “end(t)” (dotted) with an exaggerated extent of the unweighted dither shown as paler dashed (“begin(t)”) and paler dotted (“end(t)”) and how this might modify the output signal shown by introducing uncertainty in the edge locations. The x-axis 304 shows the again normalized time (in base frequency periods) t . The y-axis 302 shows T which can be interpreted as the fractional number of periods of the output signal. The darker dashed diagonal lines 310 show the function “begin(t)” which is $\theta'(t) - t$ shifted by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the rising edge of the

output signal. The paler dashed diagonal lines 320 show the function “begin(t)”, but this time shifted by the maximum quantity of dither to define between the diagonal lines 310 and 320 the range of possible locations for a dithered diagonal line representing the function “begin(t)”. The darker dotted diagonal lines 340 show the function “end(t)” which is $\theta'(t) - t$ shifted in the opposite direction by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the other edge, correspondingly the falling edge of the output signal. The paler dotted diagonal lines 330 show the function “end(t)”, but this time shifted by the maximum quantity of dither to define between the diagonal lines 340 and 330 the range of possible locations for a dithered diagonal line representing the function “end(t)”. The darker vertical dashed line 370 corresponds to the point where the darker dashed diagonal line 310 crosses $T = 0$, the x-axis, which projects down onto the latest possible rising edge of the output signal. The paler vertical dashed line 365 corresponds to the point where the paler dashed diagonal line 320 crosses $T = 0$, the x-axis, which projects down onto the earliest possible rising edge of the output signal. The darker vertical dotted line 382 corresponds to the point where the darker dotted diagonal line 340 crosses $T = 0$, the x-axis, which projects down onto the latest possible falling edge of the output signal. The paler vertical dotted line 381 corresponds to the point where the paler dotted diagonal line 330 crosses $T = 0$, the x-axis, which projects down onto the earliest possible falling edge of the output signal. The darker vertical dashed line 384 corresponds to the latest point in time and the paler vertical dashed line 383 corresponds to the earliest point in time where the dashed diagonal line crosses the $T = (+/-)^{1/2}$, which is the antipodean point to the rising edge and defines a location in time wherein applying a probabilistic jitter (a discontinuity) to the “begin(t)” counter cannot affect the output signal and so cannot generate runt pulses or other signal artifacts, even though this is dependent on the very counter to which dither is applied. The darker vertical dotted line 355 corresponds to the latest point in time and the paler vertical dotted line 350 corresponds to the earliest point in time where the dotted diagonal line crosses the $T = (+/-)^{1/2}$, which is the antipodean point to the

falling edge and defines a location in time wherein applying a probabilistic jitter (a discontinuity) to the “end(t)” counter cannot affect the output signal and so cannot generate runt pulses or other signal artifacts, even though this is dependent on the very counter to which dither is applied. The signal 362 is for reference an output signal created from the latest possible (un-dithered) counters, showing that there is a consistent systematic probabilistic bias of half of the dither width which may be neglected.

[0068] Figure 5D shows a plot 1400 of the reinitialization times (hatched boxes) of the dither values $[\text{“begin”}]_{\delta}$ and $[\text{“end”}]_{\delta}$, note that hatched boxes for the “begin(t)” cannot occur during an edge whose location may be changed by different value for $[\text{“begin”}]_{\delta}$ and vice versa. The x-axis 1402 shows the again normalized time (in base frequency periods) t . The y-axis 1405 shows T which can be interpreted as the fractional number of periods of the output signal. The dashed diagonal lines 1440 show the function “begin(t)” which is $\theta'(t) - t$ shifted by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the rising edge of the output signal, but at other values of T (such as between the dashed horizontal lines $T = +1/8$ 1498 and $T = +3/8$ 1492) can be used to determine an appropriate window in which to apply a discontinuity in the form of probabilistic jitter to the counter “begin(t)” and thus the rising edge. The dotted diagonal lines 1430 show the function “end(t)” which is $\theta'(t) - t$ shifted by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the falling edge of the output signal, but at other values of T (such as between the dotted horizontal lines $T = -3/8$ 1493 and $T = -1/8$ 1499) can be used to determine an appropriate window in which to apply a discontinuity in the form of probabilistic jitter to the counter “end(t)” and thus the falling edge. The vertical dashed line 1475 and the co-located dotted and dashed line 1480 denote the beginning and end of the region of time generated by the projection of T , in which it is valid to discontinuously modify the jitter of the function “begin(t)” safely, illustrated as left-side lined boxes 1490. The co-located dotted and dashed line 1480 and vertical dotted line 1470 denote the beginning and end of the region of time generated by the projection of T , in which it is valid to discontinuously

modify the jitter of the function “end(t)” safely, illustrated as right-side lined boxes 1415.

[0069] Turning to Figures 6A, 6B, 6C and 6D, shown are visualizations of the key parameters for a 75% duty cycle linear phase and duty realization of the method.

[0070] Figure 6A shows a plot 1500 of $\theta = \theta'(t)$ and $\theta = t$ to show the key distance parameter $(\Delta'(t))/2$ and how its projection produces the edges in the final pulse signal. The x-axis 1520 shows the again normalized time (in base frequency periods) t , while the y-axis 1510 shows normalized phase angle (in revolutions) θ . The signal 1540 now shows a 75% duty cycle to illustrate how the counter mechanism preserves duty cycle percentage whilst the phase and frequency varies. The dashed vertical lines 1575 show how the distance test of the difference between $\theta = \theta'(t)$ and $\theta = t$ being less than $(\Delta'(t))/2$ (which is $3/8$ in the illustration) maps onto a signal with a duty cycle of 75%. The horizontal wavy lines 1535 are the linearly interpolated θ values from (in this embodiment) each base frequency tick, $\theta = \theta'(t)$, a function representing a sinusoidal variation in phase which gives rise to a sinusoidal variation in frequency of the final output pulse width (amplitude), phase and frequency modulated. The diagonal lines 1525 as before show the lines $\theta = t$. The vertices 1576 show the locations where the difference between $\theta = \theta'(t)$ and $\theta = t$ is zero, corresponding to the center point of the high region of the output signal. The vertical arrows 1550 show the extent of the difference between the lines depicting $\theta = \theta'(t)$ and $\theta = t$ where the magnitude of this difference is $(\Delta'(t))/2$, denoting the edge of the output signal.

[0071] Figure 6B shows a plot 1600 of the counters “begin(t) “ and “end(t) “ to show how the sign bits of the counters with wrap-around, when a logical OR is applied (the region that is the union and is formed jointly from the interval between both dashed lines 251 and 253 but also at the same time between dotted lines 252 and 254 is the region which describes the high part of the final signal) produce the same digital signal. The x-axis 1610 shows the again normalized time (in base frequency periods) t , but in this figure the y-axis 1620 shows now T which can be interpreted as the fractional number of periods of the output

signal. The dashed diagonal lines 1640 show the function “begin(t)” which is $\theta'(t) - t$ shifted by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the rising edge of the output signal. The dotted diagonal lines 1630 show the function “end(t)” which is $\theta'(t) - t$ shifted in the opposite direction by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the other edge, correspondingly the falling edge of the output signal. The vertical dashed line 1660 corresponds to the point where the dashed diagonal line 1640 crosses $T = 0$, the x-axis, which project down onto the rising edge of the output signal. The vertical dotted line 1650 corresponds to the point where the dotted diagonal line 1630 crosses $T = 0$, the x-axis. The vertical dashed line 1663 corresponds to the point where the dashed diagonal line crosses the $T = (+/-)^{1/2}$, which is the antipodean point to the rising edge and a location in time wherein applying a probabilistic jitter (a discontinuity) to the “begin(t)” counter cannot affect the output signal and so cannot generate runt pulses or other signal artifacts. The vertical dotted line 1665 corresponds to the point where the dotted diagonal line crosses the $T = (+/-)^{1/2}$, which is the antipodean point to the falling edge and a location in time wherein applying a probabilistic jitter (a discontinuity) to the “end(t)” counter cannot affect the output signal and so cannot generate runt pulses or other signal artifacts. The signal 1670 shows the output pulse width (amplitude), phase and frequency modulated signal.

[0072] Figure 6C shows a plot 1700 of “begin(t)” (dashed) and “end(t)” (dotted) with an exaggerated extent of the unweighted dither shown as paler dashed (“begin(t)”) and paler dotted (“end(t)”) and how this might modify the output signal shown by introducing uncertainty in the edge locations. The x-axis 1710 shows the again normalized time (in base frequency periods) t . The y-axis 1720 shows T which can be interpreted as the fractional number of periods of the output signal. The darker dashed diagonal lines 1728 show the function “begin(t)” which is $\theta'(t) - t$ shifted by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the rising edge of the output signal. The paler dashed diagonal lines 1729 show the function “begin(t)”, but this time shifted by the maximum quantity of dither to define between the diagonal lines 1729 and 1728 the range of possible locations for a dithered

diagonal line representing the function “begin(t)”. The darker dotted diagonal lines 1724 show the function “end(t)” which is $\theta'(t)-t$ shifted in the opposite direction by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the other edge, correspondingly the falling edge of the output signal. The paler dotted diagonal lines 1722 show the function “end(t)”, but this time shifted by the maximum quantity of dither to define between the diagonal lines 1722 and 1724 the range of possible locations for a dithered diagonal line representing the function “end(t)”. The darker vertical dashed line 1776 corresponds to the point where the darker dashed diagonal line 1728 crosses $T = 0$, the x-axis, which projects down onto the latest possible rising edge of the output signal. The paler vertical dashed line 1774 corresponds to the point where the paler dashed diagonal line 1729 crosses $T = 0$, the x-axis, which projects down onto the earliest possible rising edge of the output signal. The darker vertical dotted line 1732 corresponds to the point where the darker dotted diagonal line 1724 crosses $T = 0$, the x-axis, which projects down onto the latest possible falling edge of the output signal. The paler vertical dotted line 1734 corresponds to the point where the paler dotted diagonal line 1722 crosses $T = 0$, the x-axis, which projects down onto the earliest possible falling edge of the output signal. The darker vertical dashed line 1738 corresponds to the latest point in time and the paler vertical dashed line 1736 corresponds to the earliest point in time where the dashed diagonal line 1728, 1729 crosses the $T = (+/-) \frac{1}{2}$, which is the antipodean point to the rising edge and defines a location in time wherein applying a probabilistic jitter (a discontinuity) to the “begin(t)” counter cannot affect the output signal and so cannot generate runt pulses or other signal artifacts, even though this is dependent on the very counter to which dither is applied. The darker vertical dotted line 1756 corresponds to the latest point in time and the paler vertical dotted line 1754 corresponds to the earliest point in time where the dotted diagonal lines 1722, 1724 crosses the $T = (+/-) \frac{1}{2}$, which is the antipodean point to the falling edge and defines a location in time wherein applying a probabilistic jitter (a discontinuity) to the “end(t)” counter cannot affect the output signal and so cannot generate runt pulses or other signal artifacts,

even though this is dependent on the very counter to which dither is applied. The signal 1730 is for reference an output signal created from the latest possible (un-dithered) counters, showing that there is a consistent systematic probabilistic bias of half of the dither width which may be neglected.

[0073] Figure 6D shows a plot 800 of the reinitialization times (hatched boxes) of the dither values $[\text{"begin"}]_{\delta}$ and $[\text{"end"}]_{\delta}$, note that hatched boxes for the “begin(t)” cannot occur during an edge whose location may be changed by different value for $[\text{"begin"}]_{\delta}$ and vice versa. The x-axis 802 shows the again normalized time (in base frequency periods) t . The y-axis 804 shows T which can be interpreted as the fractional number of periods of the output signal. The dashed diagonal lines 830 show the function “begin(t)” which is $\theta'(t) - t$ shifted by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the rising edge of the output signal, but at other values of T (such as between the dashed horizontal lines $T = +1/8$ 898 and $T = +3/8$ 890) can be used to determine an appropriate window in which to apply a discontinuity in the form of probabilistic jitter to the counter “begin(t)” and thus the rising edge. The dotted diagonal lines 840 show the function “end(t)” which is $\theta'(t) - t$ shifted by $(\Delta'(t))/2$ to produce at zero ($T = 0$) the falling edge of the output signal, but at other values of T (such as between the dotted horizontal lines $T = -3/8$ 892 and $T = -1/8$ 899) can be used to determine an appropriate window in which to apply a discontinuity in the form of probabilistic jitter to the counter “end(t)” and thus the falling edge. The co-located dotted and dashed line 844 and vertical dashed line 846 denote the beginning and end of the region of time generated by the projection of T , in which it is valid to discontinuously modify the jitter of the function “begin(t)” safely, illustrated as left-side lined boxes 820. The vertical dotted line 842 and the co-located dotted and dashed line 844 denote the beginning and end of the region of time generated by the projection of T , in which it is valid to discontinuously modify the jitter of the function “end(t)” safely, illustrated as right-side lined boxes 810. The boxes 810 and 820, it should be noted, now occur in the opposite order to Figure 5D.

[0074] III. Conclusion

[0075] While the foregoing descriptions disclose specific values, any other specific values may be used to achieve similar results. Further, the various features of the foregoing embodiments may be selected and combined to produce numerous variations of improved haptic systems.

[0076] In the foregoing specification, specific embodiments have been described. However, one of ordinary skill in the art appreciates that various modifications and changes can be made without departing from the scope of the invention as set forth in the claims below. Accordingly, the specification and figures are to be regarded in an illustrative rather than a restrictive sense, and all such modifications are intended to be included within the scope of present teachings.

[0077] Moreover, in this document, relational terms such as first and second, top and bottom, and the like may be used solely to distinguish one entity or action from another entity or action without necessarily requiring or implying any actual such relationship or order between such entities or actions. The terms “comprises,” “comprising,” “has,” “having,” “includes,” “including,” “contains,” “containing” or any other variation thereof, are intended to cover a non-exclusive inclusion, such that a process, method, article, or apparatus that comprises, has, includes, contains a list of elements does not include only those elements but may include other elements not expressly listed or inherent to such process, method, article, or apparatus. An element preceded by “comprises ...a”, “has ...a”, “includes ...a”, “contains ...a” does not, without more constraints, preclude the existence of additional identical elements in the process, method, article, or apparatus that comprises, has, includes, contains the element. The terms “a” and “an” are defined as one or more unless explicitly stated otherwise herein. The terms “substantially”, “essentially”, “approximately”, “about” or any other version thereof, are defined as being close to as understood by one of ordinary skill in the art. The term “coupled” as used herein is defined as connected, although not necessarily directly and not necessarily mechanically. A device or structure that is “configured” in a certain way is configured in at least that way but may also be configured in ways that are not listed.

[0078] The Abstract of the Disclosure is provided to allow the reader to quickly ascertain the nature of the technical disclosure. It is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. In addition, in the foregoing Detailed Description, various features are grouped together in various embodiments for the purpose of streamlining the disclosure. This method of disclosure is not to be interpreted as reflecting an intention that the claimed embodiments require more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive subject matter lies in less than all features of a single disclosed embodiment. Thus, the following claims are hereby incorporated into the Detailed Description, with each claim standing on its own as a separately claimed subject matter.