User manual extension for the vdP-oscillator

Peter Pabon, Dec 2020.

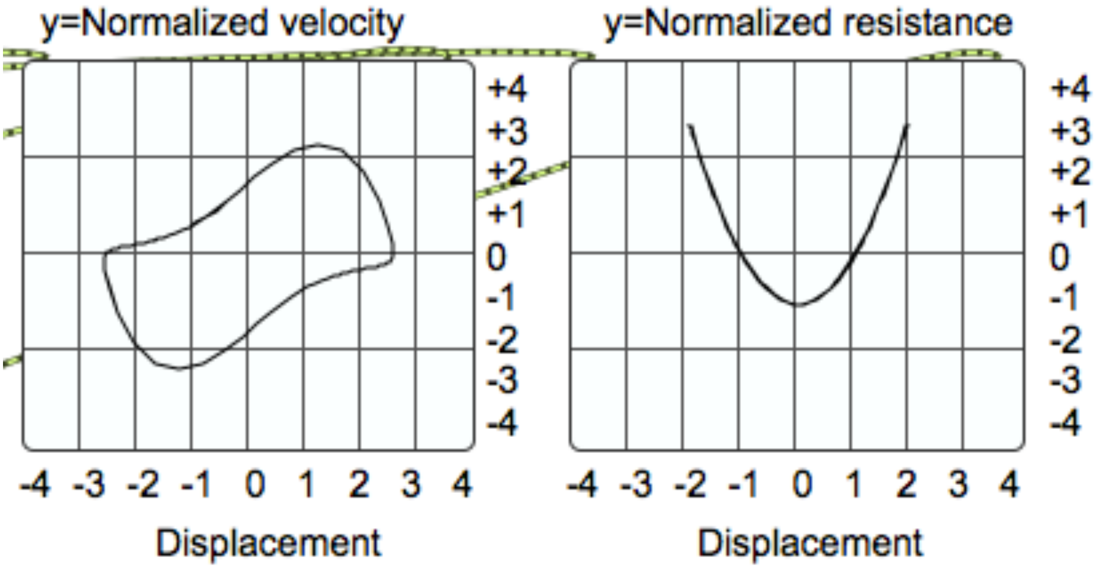

Figure 0. Phase plane, or phase portrait (left) and resistance portrait of the vdP oscillator

Implemented behaviour

Although named a vdP-oscillator, this module is not an oscillator in the classical sense. The implemented circuit is best considered a feedback patch that is able to expose very different, characteristic behaviours. Depending on its control settings, the circuit can behave as a filter, an oscillator, a frequency multiplier (ring-modulator), a phase-locked-loop (PLL), a noise shaper, or even a tone-burst generator. Furthermore, the circuit is able to suddenly and chaotically switch between these various behaviours, but it can also be steered into settings where it fluently and predictably morphs between the different output modalities.

Inside

The heart of the patch implements a state-variable band-filter circuit, that in its basic electrical response models the mechanical response of a physical mass-spring system driven by an external force. As the vdP-oscillator is an extension on this band-filter circuit, the output of the module will also be associated with one characteristic “resonance” frequency. However, contrary to a normal band filter, the special internal feedback mechanism that is employed is self-balancing on the edge of distortion and thereby favours a vibration pattern that derives from the sinusoidal path and inclines to a typical vdP-waveform. This vdP-waveform is fairly comparable to a square wave, but one with rounded edges and plateaus that gradually fall over time (top channel in Figures 1 & 2). Note the up/down alternating symmetry in the vdP-waveform for the y(t), y’(t) and y’’(t) signals that is characteristic to the square wave too. Comparably, the y(t), y’(t) and y’’(t) signal spectra will only have odd harmonics (see Figure 3).

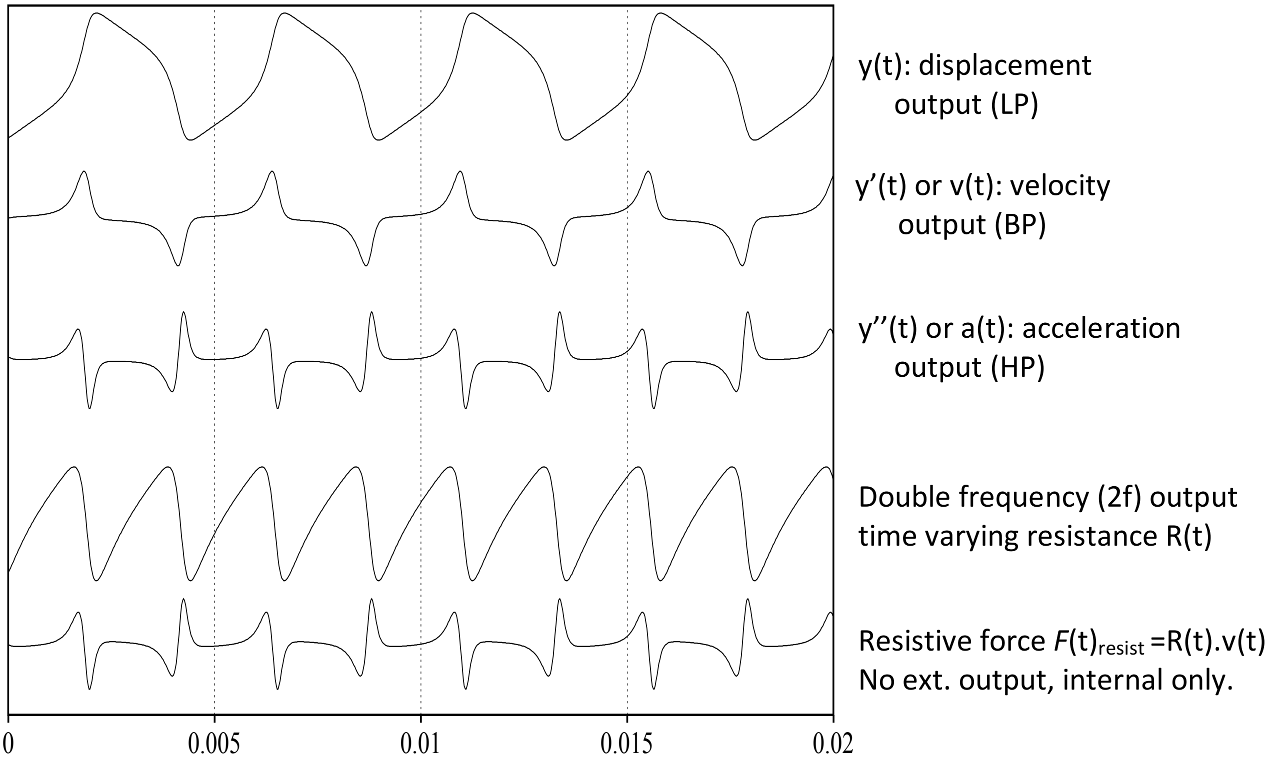

Figure 1. Signals at different points in the vdPol circuit for μ=10.

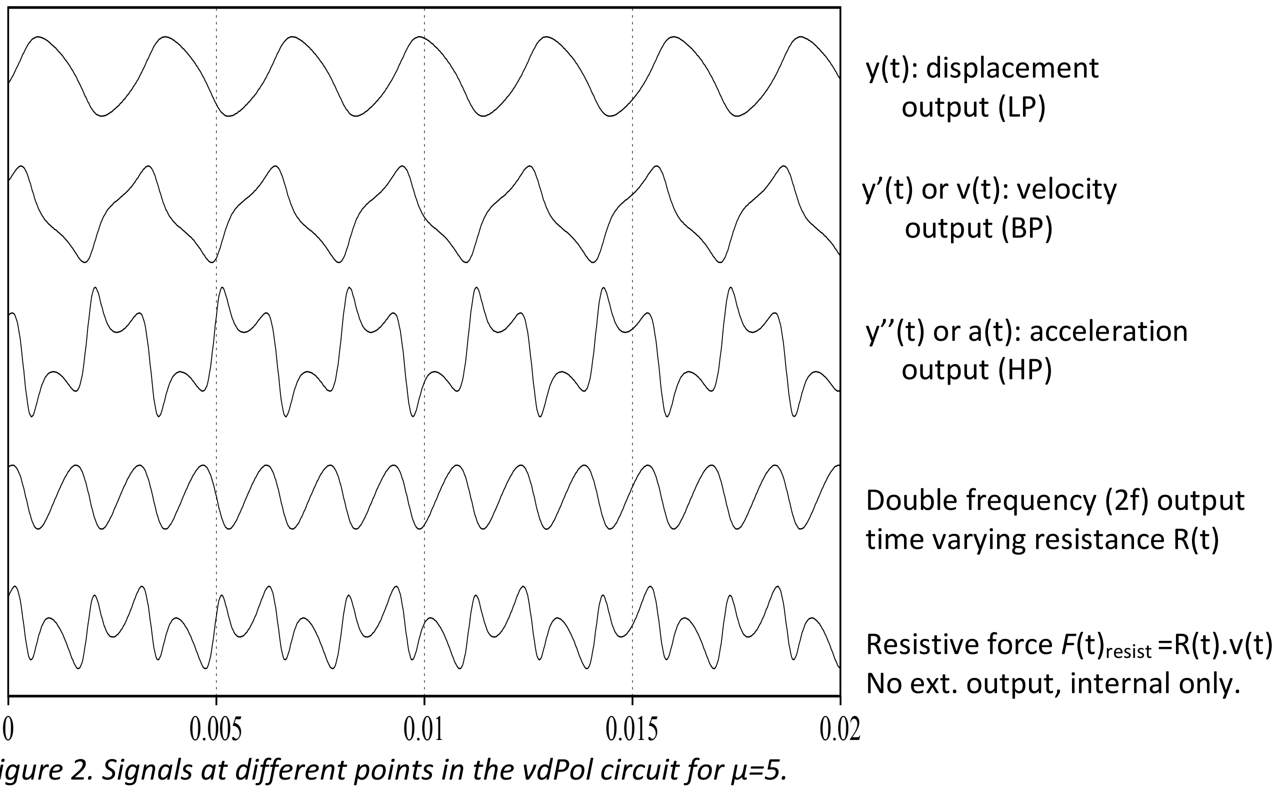

Figure 2. Signals at different points in the vdPol circuit for μ=5.

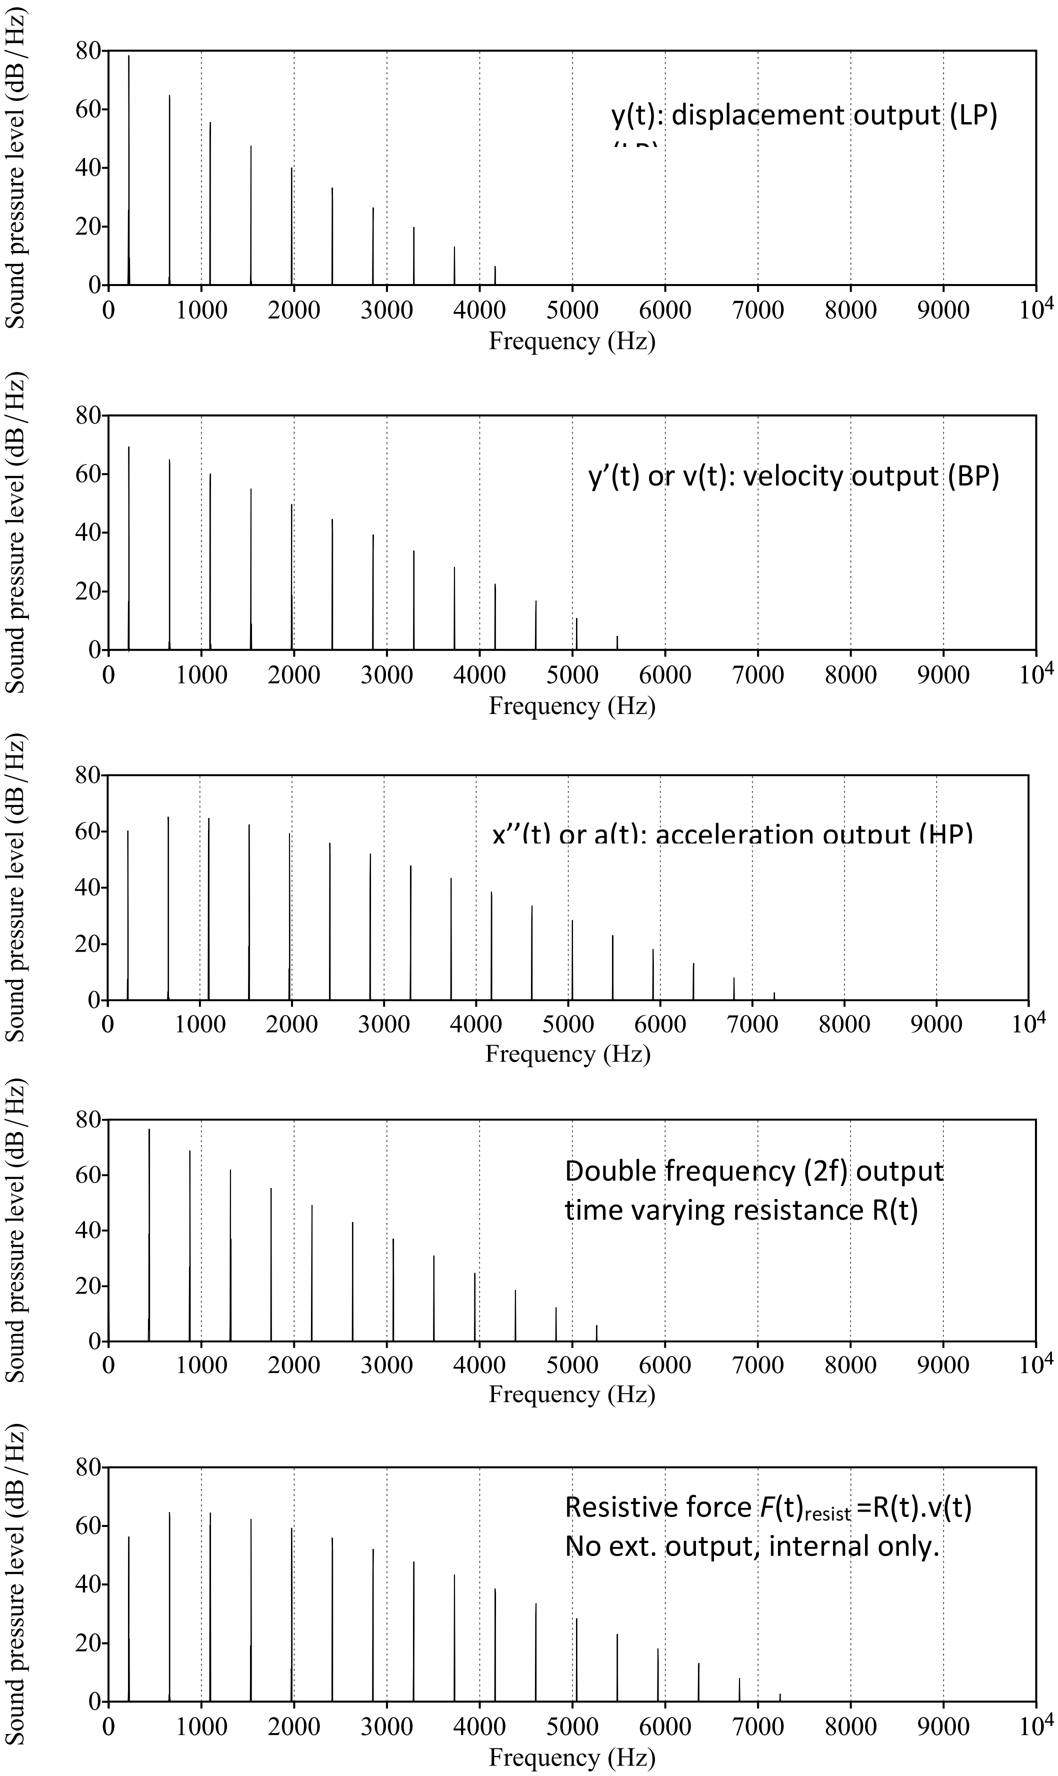

Figure 3. Spectra for the signals at different points in the vdPol circuit for μ=10.

Modulation

When the vdP oscillator is driven by an external oscillator (connected to the external force input) or when for instance two vdP-oscillators are coupled, the circuit may show a behavior that is completely different from results obtained with other modules that are available in the analog studio. The output will be neither a band filtered version of the input, nor an amplitude modulation, but something in between. Depending on the frequency (mis)match, the output will either be:

state (A) where the VdP entirely locks its internal oscillation to the external driving frequency. This mechanism works whatever the exact shape of the incoming waveform, while still the characteristic vdP waveform remains preserved, or:

state (B) where the vdP circuit oscillates completely independent from the external source, leading to an amplitude modulated mix at the output.

The moment of locking or unlocking to one of the two regimes is only partly predictable.[1]

A stronger driving force at the input will increase the chance that the system gets frequency locked, but increasing the $\boldsymbol{\mu}$ -the meaning or physical interpretation of this μ-control parameter will be further explained later- will make the system again more resistant and critical on the minimal frequency distance needed for this locked 'slave' behaviour. Once frequency (Phase) locked, the vdP-circuit has a tendency to stay locked, even when the input frequency changes to a value that it would not sync-to in a free, unlocked, situation (this is called a hysteresis effect). Once caught, the system can be dragged away from its preferred frequency, but once unlocked again it falls back to its original preference.

Relaxation oscillator

The vdP circuit is a **relaxation oscillator**. What this qualification implies in practice is best exposed and exploited with a setting where: (A) the µ-control is given a high value (favouring fast accelerations-decelerations to occur), (B) the oscillatory frequency 'f' is brought down to a very low frequency (low spring stiffness) and (C) the knob (osc/filter mix) is turned in the ocs-direction to expose maximal vdP behaviour. In this setting, the vdP-oscillatory cycle will show a typical vibration pattern where a rapid jump in polarity (the relaxation moment), will be followed by an elongated recovery phase during which the output gradual falls down to amplitude level +/- 1; from where again a relaxation jump to the opposite polarity is induced, and so alternating on. There are thus two relaxation instances (jump up and jump down) within a period. In fact, the electronic circuit that provokes the jump is tapped and the absolute values of the sensed electronic 'kicks' are presented as the double frequency output (see Figs 1,2,3). To understand how to exploit this relaxation behaviour it is better to first understand the implemented principle.

The connected physical model

The typical vdP-waveform that emerges from the circuit reflects the properties of a very specific and clever designed self-balancing (+/-) feedback mechanism that Balthasar van der Pol invented in the 1930-ies. This mechanism and much of the behaviour described above becomes easier to understand, and thereby more efficiently controlled or abused, by grasping the physics of the mechanical model that this electrical system simulates.

---

[1] The general name for this frequency locking effect is 'Arnold Tongues' and the fact this might not be a familiar term to you perhaps demonstrates the specialty of this type of oscillator.

Physical parameters as voltages

Many concepts of a mechanical system have an electrical analogue. For instance, the concept of mechanical resistance is completely comparable to the concept of electrical resistance. A mechanical displacement could be conceptually matched to a change in voltage and similarly a mechanical velocity could be conceptually matched to an amount of charge flowing per time unit through an electronic part (a current as a measure of the speed of voltage change). However, a current running through a resistor will again produce a proportional voltage over this resistor. This voltage difference perfectly shadows the current and thus a voltage level may conceptually represent a velocity too. So, in the actual electrical implementation of the physical model, the matching of mechanical parameters to an electrical analogue is not a simple one-to-one translation. The art of designing an electronic model of a physical system, is how to scale each "parameter" of the mechanical model as a representative electrical voltage appearing somewhere in the circuit. For instance, within the electronic design of the state-variable filter that simulates a mass-spring system, there are typically three designated points in the schematic that represent: (1) the virtual displacement (called the LP output), (2) the virtual velocity (called the BP output), and (3) the virtual acceleration/force (called the HP output). All three physical parameters are time varying signals that can be simultaneously probed within the electrical circuit. Apart from probing, time-varying signals can be added at designated points in the circuit to change or modulate the modelled physical property. For instance, the voltage control input in "audio/CV" adds a voltage to the acceleration/force point to simulate an external driving of the virtual mass of the system. In the vdP-circuit, there is a designated point where the sensed voltage represents the mechanical resistance in the modelled physical system. The earlier mentioned voltage controlled µ-parameter allows you to modulate (rescale) this simulated resistance.

The mechanical model

The next graph sketches schematically the physical system that is modelled by the electronic vdP-circuit:

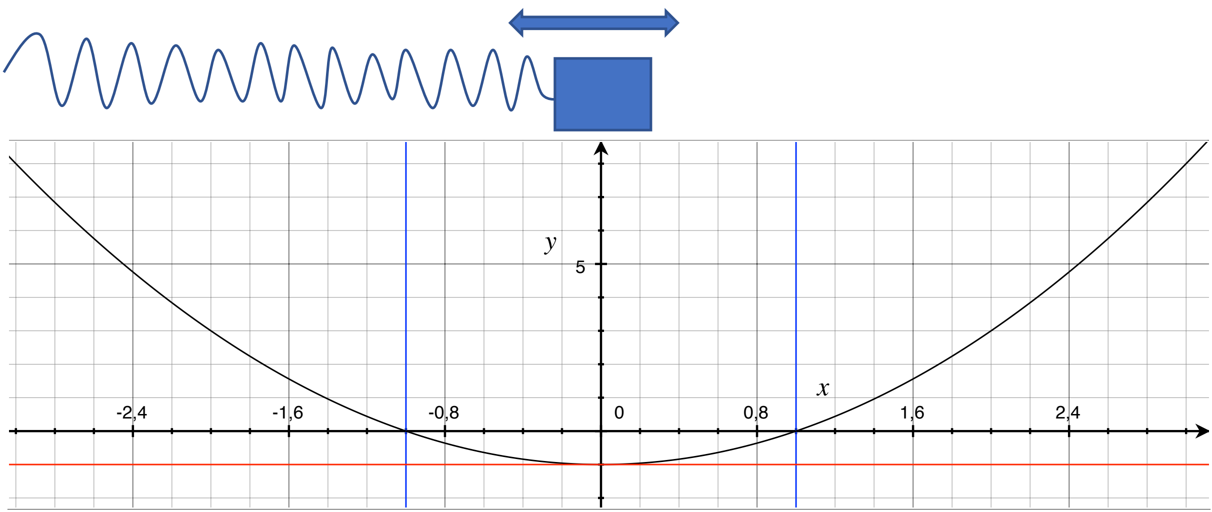

Figure 4. The x-axis measures the horizontal displacement of the virtual mass that is connected to a virtual spring that is at equilibrium at x=0. The (x^2-1) function on the y-axis specifies the virtual resistance this mass feels locally when it moves through a virtual medium with a viscosity gradient that varies with the displacement.

Potentials (voltages) in an electric circuit will not change if there are no currents flowing through the electronic parts. Similarly, the above modelled physical system will not gain or lose potential energy if the mass does not move. When the mass in Fig. 4. would be displaced horizontally from its equilibrium position (at x=0) by a short push from an external force, the system will respond with a familiar sinusoidal movement due to the working of the spring. During this cycle, there is a periodic exchange between kinetic energy (mass velocity) and potential energy (stretching of the spring). The kinetic energy is at its maximum when the mass due to its inertia crosses the zero point. In the next phase of the cycle, this kinetic energy is fully exchanged for potential energy when the mass loses velocity while stretching the spring to its highest amplitude. Note, that at points of maximum amplitude deflection (zero velocity) there will never be any frictional energy exchange with the environment. A trivial, but not unimportant insight, is that only by movement a system can lose or gain kinetic energy. In the process where kinetic energy is lost to-, or gained from the surroundings, having more velocity '*v*' and/or finding more resistance '*R*', both lead to a larger frictional force $F_{\text{friction}}=v.R$. It is the application of this force over the travelled distance (Force*displacement=Work) that determines the energy dissipation to the environment. A basic insight is that by dynamically controlling the amount of resistance (while not actively controlling the velocity), the rate at which a system may lose (or gain) kinetic energy (velocity) can be modulated, and thus a self-sustaining oscillatory mechanism can be constructed.

Variable resistance

The heart of the vdPol-oscillator is centred around the concept of having a resistance that varies as a function displacement. The (x^2-1) function sketched in the above figure is the actual dependence implemented in the electronic vdP-circuit. The curve illustrates that the resistance of the medium through which the mass moves is not uniform. The resistance increases with a squared factor when moving either to the left or right extreme (you can interpret this as the surrounding fluid rapidly becomes more viscous further away from the equilibrium position). A varying resistance means that for the same velocity value, depending on where the mass is moving along the x-axis, the system may lose (or gain) different portions of kinetic energy while moving. Note that the sketched resistance gradient in Fig. 4. is not representing a force field or potential that will automatically set the mass into motion. When there is no movement (zero velocity), there will be no resistive force acting that could displace the mass or guarantee an energy exchange with the environment. This is why there is a spring added to the vdP-model, as this mechanical element will guarantee a basic initial oscillatory movement where the spring periodically drives the mass back and forth along the x-axis. With the resistive force factor scaled down to zero, (by zeroing the μ control parameter that vertically scales this function), the vdP-oscillator will behave as a classical mass-spring system that does not gain or lose energy, and will respond as a band filter with a corresponding sinusoidal motion. By upscaling the resistance effect of the above gradient function (with the μ control), the initial sinusoidal motion will get modulated as a function of the current x-displacement and as a result the sinusoidal periodic motion gradually deforms to the typical vdP-waveform. The cycle duration will change too as a result of the different resistances felt along the cyclic path.

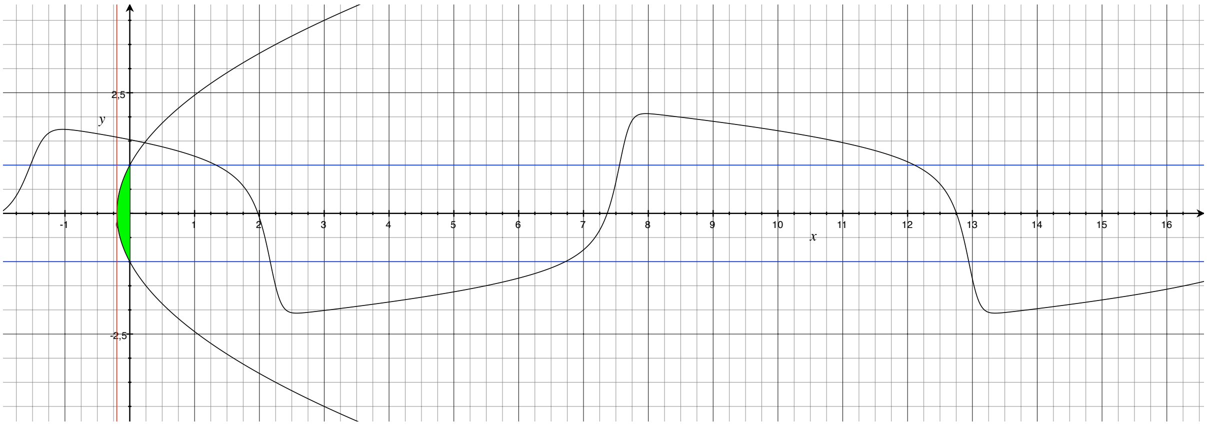

Figure 5. Displacement (now vertically) waveform as a function of time (horizontal axis).

Traversing the negative resistance zone

Let us inspect the path of the mass movement and sketch its physical state as a function of time (horizontally) in detail (see Fig. 5). Note, that the first horizontally oriented displacement axis from fig. 4, is now rotated 90 degrees and displacement is projected vertically as y-axis. We start at the moment just following a positive zero-crossing jump. The virtual moving mass that is modelled by the electrical system, will start to lose virtual velocity once it comes above absolute amplitude '+1'. Once outside the central region between the two blue lines, the resistance is always positive, which means damping, or energy loss. The damping factor increases with more extreme displacement amplitudes. The virtual material gets only stickier and more resistant the larger the amplitude, and thus any movement rapidly stops. The only force that can release the clogged mass from the modelled sticky medium is the spring that will gradually pull the mass back to the midline. The way out of this viscous material may at first be a slow process, but there is hope… as with lower displacement amplitudes, the virtual surrounding gets less and less viscous. Actually, the resistance decreases with a quadratic factor when the amplitude decreases to zero, and so the pull of the modelled spring becomes more effective, and the velocity increases rapidly (the mass accelerates). When crossing displacement amplitude value '+1' a curious phenomenon happens; the virtual resistance of the surrounding material becomes negative! A negative resistance is not a familiar physical phenomenon, but we are in a virtual environment where weird things may happen. This virtual anti-resistance theoretically implies that there will be a force acting in the direction the object is already moving, where a larger negative resistance implies more force per velocity unit. As a result of this anti-resisting force, the velocity will increase on itself and between the blue lines the mass accelerates exponentially. Note, that at displacement=0 the negative resistance is at its minimum. Thus, at the point where the spring force will be zero, there is maximum acceleration due to the anti-resisting force finding a maximum. Effectively, the mass will launch itself over the zero crossing. After the crossing, the negative resistance gradually gets less extreme, but still remains negative. Although the anti-resisting force ceases, the accelerating force is not yet changing sign. The mass will still gain velocity, and thus kinetic energy, as long as the force is in the direction it is travelling. Once outside the section bounded by the two horizontal blue lines at amplitudes -1 & +1, the resistance curve will again be positive (damping), and deceleration sets in. Due to its high velocity, the mass will rapidly enter the high viscosity zone and lose most of the kinetic energy that it gained along the zero-crossing path almost instantly. That the decelerating force increases with a quadratic factor helps a lot decelerate fast. The mass is now at the other extreme deflection point, and with nearly no velocity, again in a state where only the spring force can slowly pull it out of this sticky environment and what follows is again a first slow approach in the direction of the negative resistance zone in the middle, to get boosted again during the next the zero crossing etc. The same process is repeated in the opposite direction, and so on. Note, that any extra velocity gained along the way will not help to make the process repeat faster, as more velocity implies a deeper launching into very sticky surroundings.

Interpretation of the μ-parameter

The μ-control parameter scales the resistance curve. The effect of increasing μ is a faster acceleration/deceleration within the cyclic process: with faster transients the sound gets sharper. The higher velocities shorten the zero-crossing time, and thus could lead to an increase in pitch. Contrary, the time spend due to effective launching in sticky regions increases, which again lengthens the periods and thus reduces the pitch. How this effect will balance out on the effective period duration will also depend on other variables and system dependencies, which in practice makes it hard to predict or generalize how this μ control may affect the observed pitch of the oscillator.

For instance, the intermediate speeding up due to the high μ, increases the velocity values, and thus the amplitudes of the voltages in the BP output that represent this velocity, that thus should increase too. However, the electrical circuit is designed with an automatic gain control. Thus, the amplitude of this BP signal no longer shadows the amplitude of the physical velocity, but its frequency content will still represent it.

Use and abuse of the vdP circuit.

Generating bursts of pulses

The physical interpretation of the oscillating mechanism given above, can help to understand how to exploit the relaxation mechanism and put the vdP-circuit in a burst-like oscillating behaviour. To reach this state:

(1) The resonance frequency knob '*f*' [or the control voltage for *f*(cv)] needs to be brought down its lowest value. You can interpret this control parameter action as a weakening of the internally modelled spring. The less stiff a spring, the less force there will be to pull the mass back to its equilibrium position, and the lower the frequency at which the mass oscillates due to spring forces.

(2) Set the μ-parameter [or the control voltage for μ(cv)] to a high value. The μ-parameter scales the range of the resistance curve, which means faster accelerations/decelerations within the cyclic process.

(3) Add some noise to the external force input.

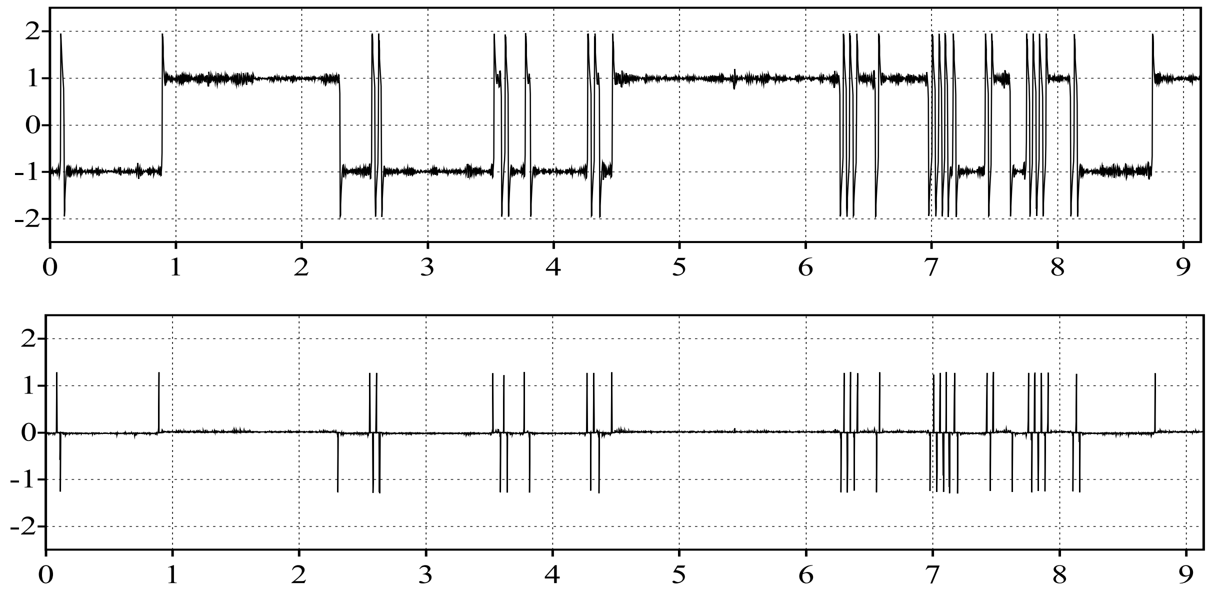

Figure 6. Displacement and velocity time signals for μ=60 & negligible spring stiffness. Some noise is added to the force input to randomly dislocate the mass from its docking position.

With a small displacement from the equilibrium position by some external noise signal at the force input, there will not be that much counter force by the weak spring pulling the mass back to the equilibrium position. However, the little bit of velocity gained at the disruption, will due to the high negative (anti-)resistance around zero displacement, immediately produce a lot of force in the direction the mass is going. The mass will accelerate rapidly in the direction it was accidentally going. Once above displacement amplitude 1, a rapid deceleration starts when the mass lands in very sticky environment at one of the extremes. The spring will be able to slowly pull the mass out of this environment and the vdP-circuit could start to get in a continuous state of self-oscillation. However, the continuous noise that is been added to the force input, will have a curious effect when the mass crosses amplitude +/- 1 with a low velocity. This is the point where the resistance effect changes its polarity from + to -. Although the mass could shorty enter in the negative (anti-resisting) range, the random sign changes of the external noise force may still be enough to flip the direction of the velocity and thus make the mass accelerate back up again to the amplitude +/- 1 line. In practice, this little noise added increases the chance to get stuck at either the +1 or -1 amplitude value (see Figure 6). It could happen that a sudden higher spike within the noise in the direction of the zero-crossing will be just enough to get the mass accelerate on itself and jump to the other side. Once gaining speed and thus anti-resisting force, the external noise spikes forcing in the opposite direction will lose their effect. Moreover, when the system makes this randomly triggered relaxation jump to the other side, there can be so much energy gained along the way, that on the way back from the sticky environment, while moving in the direction amplitude +/-1, there is still enough kinetic energy left to temporary overrule the noise that could again reflect the mass in the opposite direction at the +1/-1 line. As a result, there will be burst of periodic vdP-oscillations, that could by accident be stopped when mass is around amplitude +1 or -1.

With too little disrupting noise, the circuit will go on oscillating uninterruptedly. With just enough noise amplitude, the chance that the mass will be caught in rapid random oscillations around the +/-1 displacement points will increase. With even more noise, such a catch gets even more secured. With again more noise, the chance to leave the zone increases again. So, there are many variables to explore here, as what happens when you give a little voltage offset to the noise. Have fun with the vdP!