CNRS
Centre National de la Recherche Scientifique

INFN
Istituto Nazionale di Fisica Nucleare

Lock acquisition of the Advanced Virgo arm-cavities with reduced force

VIR-0019A-12

Bas Swinkels and Paolo Ruggi

Issue: 1

Date: May 4, 2012

VIRGO * A joint CNRS-INFN Project
Via E. Amaldi, I-56021 S. Stefano a Macerata - Cascina (Pisa)
Secretariat: Telephone (39) 050 752 521 * FAX (39) 050 752 550 * Email W3@virgo.infn.it

Contents

1 Introduction

Gravitational wave detectors like the Virgo interferometer are brought from their uncontrolled state to their working point in a long sequence of actions called the lock-acquisition. One of the critical steps at the beginning of this sequence is the lock-acquisition of the high-finesse arm-cavities. For Advanced Virgo, several parameters will change that will make this lock-acquisition more difficult:

- The finesse of the arm-cavities will increase from 150 today to around 443, which will make resonant optical transients more pronounced.
- The mass of the mirrors will increase from around 21 kg today to around 40 kg, so the same actuation force will only cause half the effect.
- The payloads will be redesigned completely and will probably no longer feature a *recoil mass*. Instead, both the *marionetta* and the mirror will be actuated from *Filter 7*. This might complicate the controllability of the mechanics.
- The strength of the magnets used for the actuation of the mirrors will likely be reduced, which limits the available actuation force. This might be done for several reasons:
  - Magnets with a lower strength will have a lower coupling to environmental magnetic fields.
  - Less powerful magnets are generally smaller. Since the magnets have to be glued to the main mirrors, this lowers the risk that they spoil the high quality factor of the mechanical resonances of the mirrors.

– Smaller magnets will cause less problems with eddy-current damping, which spoil the quality factor of the pendulum mode [1]. This effect was however never observed in Virgo. It can be solved by using a dielectric recoil mass.

– A lock acquisition that uses less force will cause less disturbances in the payload. This might be important for the new design of the payload, which will probably feature a high coupling from the actuation on the mirror to the *tx* degree-of-freedom of Filter 7. It is also important to reduce any shocks that could excite the *violin modes*. Their quality factor should be very high, so it might take days to calm down in case of excitation.

One possibility would be to first lock the cavity with an auxiliary laser at a different wavelength, after which the lock is handed over to the main laser [2, 3]. In this note, such systems will not be discussed. Instead, we will focus on the limits of the current method of locking the cavities and will see what can be gained by improving the locking algorithm only.

2 Theory

2.1 Threshold-velocity

Due to various physical limits, it is not possible to lock a cavity with a linear controller when the mirror is passing a resonance with arbitrary large speed. Some of these limits can be calculated [4]. These formulas are not exact, but they should give a good idea how the threshold velocity scales with various parameters. The time it takes to cross a resonance t_{res} is

$$t_{\mathrm{res}} = \frac{\lambda}{2v\mathcal{F}}, \tag{1}$$

where λ is the wavelength, v the velocity of the mirror and $\mathcal{F}$ is the finesse of the cavity. The first limit v_{max1} follows from the response time of the feedback loop, which is roughly $t_{\mathrm{loop}} = 1/2\pi B$. Equating this to Eq. 1 and solving for the velocity yields

$$v_{\mathrm{max1}} = \frac{\pi\lambda B}{\mathcal{F}}, \tag{2}$$

with B is the bandwidth of the control loop in Hz. If the resonance is passed with a larger velocity, the loop will not be fast enough to acquire the lock.

A second limit comes from the impulse of the mirror mv, which can be stopped by applying the maximum available force F_{max} for a time $t_{\mathrm{pulse}} = mv/F_{\mathrm{max}}$ If the required time is longer than the time to cross the resonance, the cavity cannot be locked. This gives a threshold velocity

$$v_{\mathrm{max2}} = \sqrt{\frac{F_{\mathrm{max}}\lambda}{2\mathcal{F}m}}. \tag{3}$$

A final requirement comes from the resonant behavior of the cavity, which makes that the field inside the cavity takes a finite time to build up. Requiring that the time to pass resonance is longer than the *storage time* $t_{\mathrm{st}} = 2\mathcal{F}L/(\pi c)$ yields

$$v_{\mathrm{max3}} = \frac{\lambda\pi c}{4\mathcal{F}^2 L}, \tag{4}$$

with c the speed of light and L the cavity length. Passing the resonance with a higher speed will lead to *ringing effects* in the signals in reflection and transmission of the cavity, which could spoil the signal that is used to lock the cavity [5].

Filling in some typical values for the Virgo, Virgo+ and Advanced Virgo interferometers, the threshold velocities can be calculated, see Table 1. Lacking a design of the new payload, it is assumed that the maximum available force per actuator will not change. Up to now, the lock was always performed using a single coil-pair (up & down coils of the end-mirror). It is assumed that for Advanced Virgo, all 4 coil-pairs per cavity will be used

(both up-down and left-right coils of both the input and end-mirrors). Furthermore, it is assumed that due to upgrades in the control system, a factor two in bandwidth can be gained.

Table 1: Calculated maximum velocities

Property	Symbol	Unit	Virgo	Virgo+	Adv. Virgo
mirror mass	m	kg	20	20	40
arm cavity finesse	$\mathcal{F}$	-	50	150	445
loop bandwidth	B	Hz	100	100	200
maximum force per coil-pair	$F_{\max}$	mN	2.9	2.9	2.9
number of coil pairs	-	-	1	1	4
threshold-velocity due to finite bandwidth	$v_{\max 1}$	um/s	6.7	2.2	1.5
threshold-velocity due to finite force	$v_{\max 2}$	um/s	1.2	0.7	0.6
threshold-velocity due to ringing effects	$v_{\max 3}$	um/s	33.4	3.7	0.4

As shown in the table, the threshold velocity due to the bandwidth will not decrease a lot for Advanced Virgo, since we assumed that the effect of the increasing finesse can be partly compensated with an increase of the control bandwidth. The threshold velocity due to the maximum force can be partly recovered by using more than 1 coil-pair. This limit will however get more critical if it is decided to reduce the strength of the magnets. The decrease of the threshold velocity due to the ringing effect seems to be the most most critical for Advanced Virgo, especially since it scales inversely with the square of the finesse.

2.2 Signal linearization

The signal used to lock a cavity is usually a Pound-Drever-Hall (PDH) signal, which is obtained by phase-modulating the laser field and demodulating the signals in reflection or transmission of the cavity. This signal is linear over a range of the cavity length that is roughly equal to its linewidth, so $\lambda/(2\mathcal{F})$. It is possible to increase this range by normalizing the PDH signal with the transmitted power of the cavity [4]. When the cavity is almost out-of-resonance, this method requires dividing by a small number that is affected by shot or electronic noise. The increase of the linear range is therefor only a small factor (~ 3). This works well to somewhat relax the requirements due to the bandwidth and the maximum force. It does however not help for the threshold due to the ringing effects, since both the PDH-signal and the transmission are affected in a different way.

2.3 Guided lock (single impulse)

The typical speed of a free-swinging mirror, after being *damped* by the *local control* system, is on the order of 1 um/s. According to the values of Table 1, it should thus be possible to lock the cavities for Virgo and Virgo+ without much problems. This is confirmed by the experience that it until now, it was always easy to lock the arm-cavities. The change of parameters for Advanced Virgo, however, might just be enough to change this from an easy to a hard (or even impossible) task.

It thus becomes important to reduce the velocity of the mirrors before the actual lock acquisition is attempted. A solution to this problem, sometimes called *guided locking*, has been studied many years ago at LIGO and TAMA [6, 7]. The method consists in estimating the velocity when the uncontrolled cavity crosses a resonance, after which a single, long impulse is applied to the actuators. The length of the impulse is chosen so that it is enough to completely stop the mirror and accelerate it back towards resonance with reduced velocity. The controller then waits for the mirror to drift back to another crossing of the resonance. If the new velocity is low enough, a feedback loop with a standard linear controller can be engaged. If the velocity is still too high, another pulse can be applied. How to calculate the pulse length based on the measured velocity is explained in Appendix A.1.

2.4 Guided lock (double impulse)

During the time it takes for the mirror to return back to resonance, it can be disturbed by the restoring force of the pendulum and the residual seismic noise. To limit this effect, this time should be kept as short as possible. We therefore investigated a slightly modified algorithm. The idea is to apply both a positive and a negative pulse, so that nominally, the mirror arrives at resonance with velocity close to zero. Instead of simply waiting for it to drift back to resonance after the first pulse, the pulse is extended to accelerate it further. To finally arrive at low speed, a second pulse with opposite sign is applied to brake the mirror just at the right moment. How to calculate the length of the two pulses based on the measured velocity is explained in Appendix A.2. The double-pulse method does indeed seem slightly faster, but it is not sure if this justifies the increase in complexity of the algorithm.

3 Velocity estimation during a resonance crossing

As described in section 2.3, the guided-lock algorithm needs an accurate measurement of the velocity of the mirror when it passes through a resonance. Ideally, such a measurement gives a value that is linear to the velocity, but any function that gives a value that depends in a smooth way on the velocity would be sufficient, since this could be corrected with e.g. a calibration table. Other requirements to the measurement are that it should be fast to calculate (within a few samples of passing the resonance), it should be robust against noisy signals and it should be able to measure velocities that are larger than the velocity to which the system can be pre-damped using the local controls.

One complex method of measuring the velocity based on the ringing effect was given by Matone [8]. This method requires finding the minima and maxima of the ringing signal, after which a polynomial fit has to be done, which is non-trivial to implement.

An alternative method uses only the derivative of the linearized error signal described in section 2.2. At a velocity that is low enough to not cause ringing effects, this error signal is linear to the position, so its derivative is proportional to the velocity. It was found that the same method can be used to accurately measure higher speeds, if a simple correction by a power-law is applied [7]. The measured value v_{meas} is thus calculated as

$$v_{\mathrm{meas}} = C \left(\frac{\partial}{\partial t} \frac{S_{\mathrm{PDH}}}{S_{\mathrm{DC}}} \right)^{P}, \tag{5}$$

where S_{PDH} and S_{DC} are the PDH and transmission signals from the photodiodes, C is a calibration constant and P is a power. This calculation can be implemented easily in real-time by using a high-pass filter.

At high speeds, the calculated velocity is only valid during a few samples around the crossing of the resonance. Care must thus be taken to extract the velocity in a robust way. Trail-and-error, both in experiment and simulation showed that the best moment to *sample* this velocity is at the moment that the PDH-signal crosses zero. A threshold on the transmitted power of the cavity should be used to discriminate a zero-crossing due to the resonance crossing of the carrier from a zero-crossing due to a sideband or due to noise. As an alternative sample time, one could choose the maximum of the transmitted power. This proved to be less reliable, probably since it occurs slightly later in time and is thus closer to the moment the ringing effect starts to spoil the signal.

3.1 Experimental calibration with Virgo+

To demonstrate the velocity estimation using the normalized error signal, some data of a free swinging arm-cavity of Virgo+ was analyzed. To get data at higher speeds than normal, one of the mirrors was excited by a kick to make it oscillate at its main pendulum mode.

If the mirror oscillates by more than a few wavelengths, the velocity can be easily calibrated. Every time the cavity passes a resonance (as seen by the transmission of the cavity), the length of the cavity should be an

exact integer number of half-wavelengths. By looking at the derivative of the PDH error signal, the sign of the velocity can be obtained, so the mode number of the resonance can be obtained by counting up and down [12]. The position of the mirror when it is between two resonances can then be estimated using a cubic-spline interpolation, from which the velocity can be obtained using simple derivation, see Fig. 1.

For every resonance crossing, the velocity is calculated using the algorithm described in the previous section. The obtained values are plotted versus the calibrated velocity obtained with the interpolation method. As shown in Fig. 2, the obtained relation is slightly non-linear. This can be corrected by applying a simple power law, as suggested in [7]. Some trial-and-error shows that $P \approx 1.1$ works best. After correction, the signal is linear up to 10 um/s, with a standard deviation of the relative error of around 2 percent.

3.2 Simulation using E2E for Advanced Virgo

Performing an accurate measurement of the velocity will become more difficult for Advanced Virgo, since the ringing effect will be bigger due the increase of the finesse. To see if such a measurement is still possible in the future can be checked using a simulation. A simple Fabry-Perot cavity was simulated using E2E [13], with all the optical parameters set as in Advanced Virgo. The length of the cavity was changed at constant velocity, see Fig. 3 for an example of the simulated signals during a resonance crossing.

The simulation was performed for a range of velocities up to 10 um/s. In each case, the velocity was calculated using the algorithm of section 3 and then plotted versus the true velocity, see Fig. 4. As with the measurements on Virgo+, the obtained relation is slightly non-linear, but this can again be corrected by raising the obtained value to the power 1.1.

Note that the simulation did not include any electronic, ADC or shot-noise. Only the discretization in the time domain is simulated at 10 kHz. The small steps seen in the calibration curve might be due to the sampling point jumping by 1 sample. Around 9 um/s, the algorithm breaks down, probably due to the finite sample rate. The maximum speed at which the speed can be measured will probably be much lower if the real noise is considered.

4 Non-linear lock acquisition

4.1 Experiment with Virgo+ using a DSP

A first experimental test of a non-linear locking algorithm similar to guided locking was implemented using one of the DSPs, just before the shutdown of Virgo+. Global Control was modified to simply forward the signals `Pr_B5_ACq` and `Pr_B7_DC` to the North-End DSP. The algorithm in the DSP was modified to calculate the normalized error signal and take its derivative, which should be proportional to the speed. A correction to the mirror was then sent to the mirror proportional to this value, causing more or less a viscous damping. Since the goal of this test was to see with how small forces we could still lock, the correction was clipped to low values. Due to this clipping, the correction was saturating, so the resulting correction resembles a bit the rectangular pulses used in the guided lock algorithm. The correction is switched off when the transmission of the cavity drops below a certain, very low threshold. If the velocity is low enough, a standard linear lock is engaged.

Even though this method is not an exact implementation of the guided-lock algorithm, it was able to reliably lock the cavity with a limit on the actuator voltage as low as $V_{\max} = 0.02\,\mathrm{V}$ (corresponding to a force of the actuator $F_{\max} = 0.073\,\mathrm{mN}$). The maximum value for a single coil-pair is $V_{\max} = 0.8\,\mathrm{V}$ ($F_{\max} = 2.9\,\mathrm{mN}$), so 40 times less than is used normally with the up-down coils of the end-mirror. Since up to 4 coil-pairs can be used for the lock (both up-down and left-right for both the input and end-mirror) the total force used for the lock was thus a factor 160 lower than the maximum available force. Fig. 5 shows an example of such a lock acquisition, see [9, 10] for more details.

One problem that was discovered was that with such low force, the cavity usually unlocks after a few seconds due to low-frequency disturbances which saturate the actuator, as seen at the end of Fig. 5. This has been

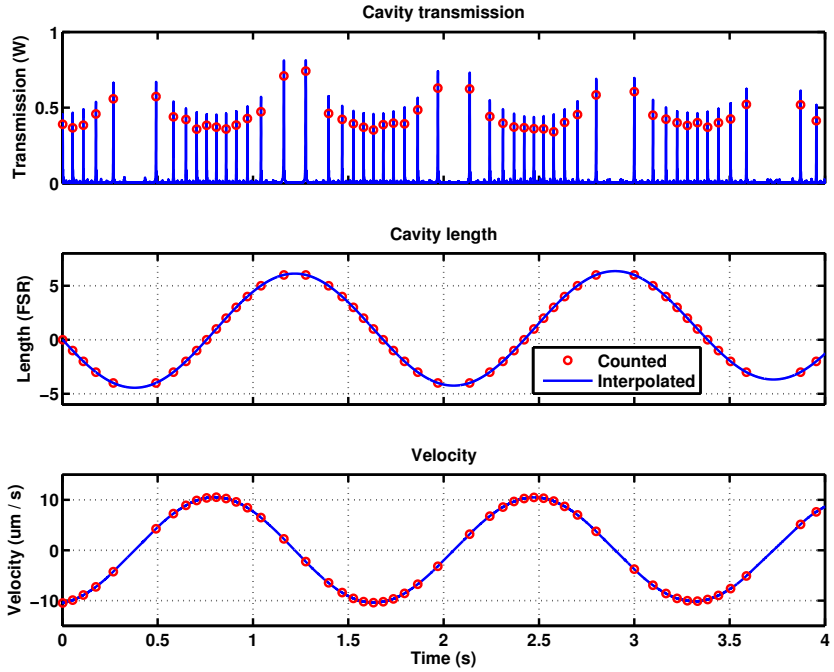

Figure 1: Calibration of velocity by interpolating signals from a free-swinging cavity.

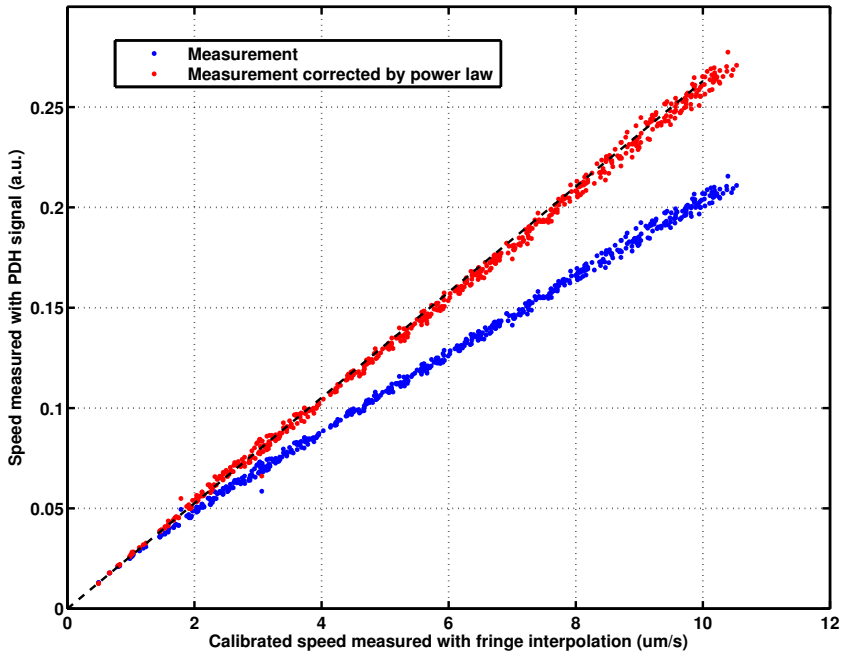

Figure 2: Speed measured with the Pound-Drever-Hall signal vs calibrated speed obtained by fringe-interpolation, using data shown in Fig. 1. The slight non-linearity can be corrected by raising the measured value to the power 1.1.

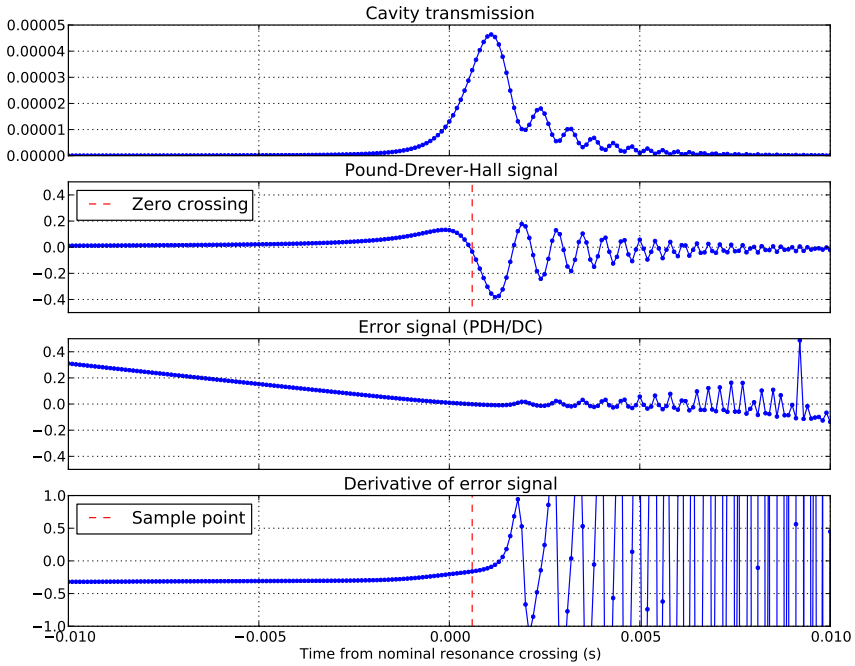

Figure 3: Simulation of an arm cavity of Advanced Virgo that passes a resonance with a velocity of 5 um/s. Note that no electronic and ADC noise is simulated. In reality, the error signal will be very noise when the transmission of the cavity is less than 1/10th of the maximum value and a usable signal for the speed measurement is only available for a few samples. The red dotted line indicates the moment the velocity is calculated.

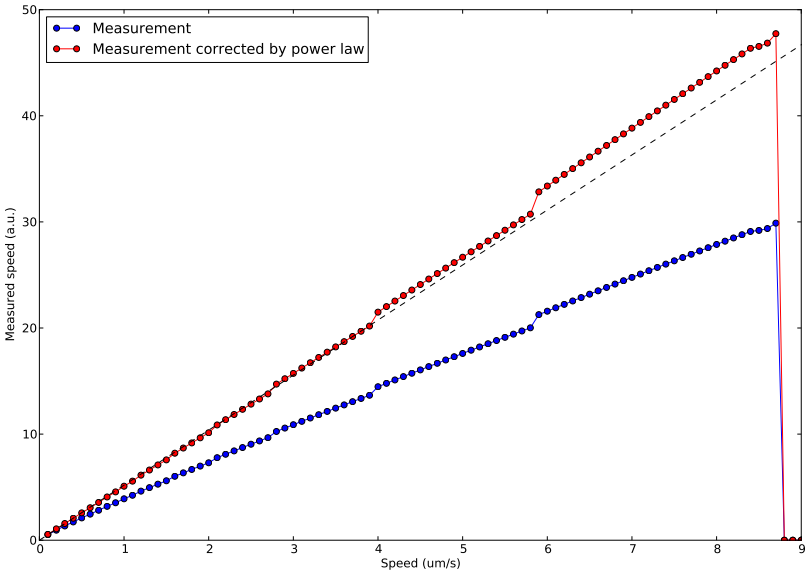

Figure 4: Measured speed vs real speed for Advanced Virgo, using data simulated with E2E as shown in Fig. 3. The slight non-linearity can be corrected by raising the measured values to the power 1.1.

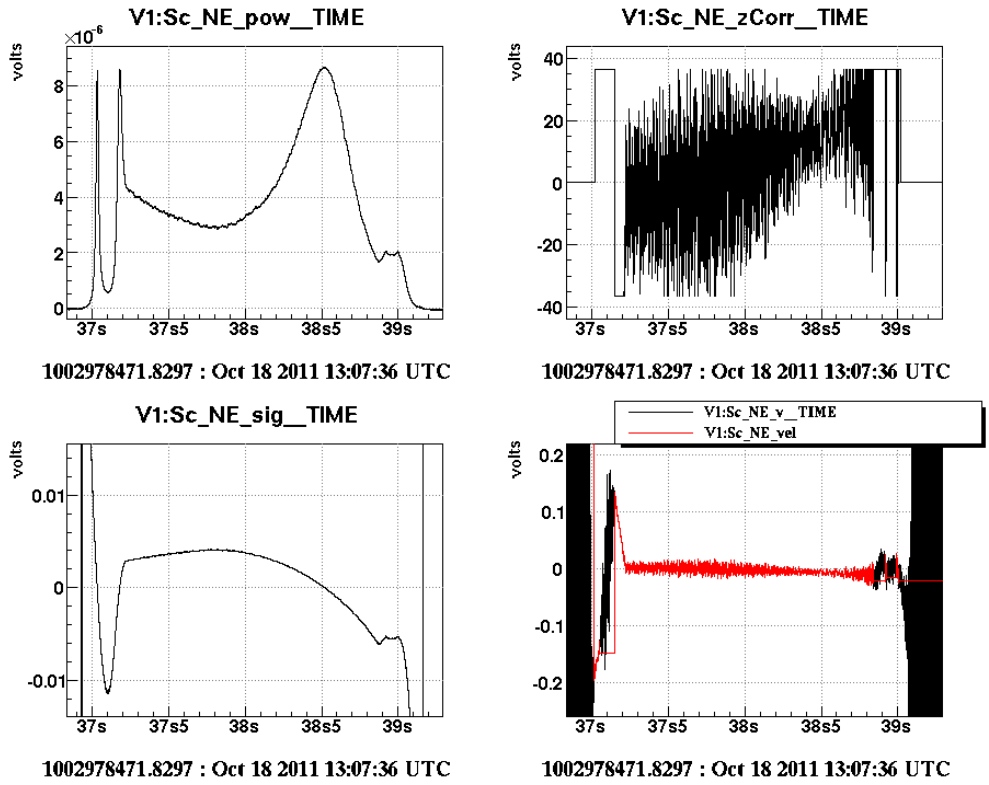

Figure 5: Demonstration of lock with the DSP. After 2 seconds, the lock is lost due to an artificial saturation of the actuator. The zCorr signal has been rescaled, it was limited to 0.01 V in this experiment.

solved by an almost immediate *reallocation* of the low-frequency part of the correction to the *marionetta*, see Fig. 6. The marionetta, which is located one stage higher in the suspension chain, has a much larger dynamic range at low frequencies. In Virgo, the reallocation was only done at later steps of the lock acquisition.

4.2 Experiment with Virgo+ using Global Control

In a second experimental test, a guided-lock algorithm was implemented in Global Control, using the 2-pulse method described in section 2.4. As with the test with the DSP, the error signal is calculated as `Pr_B5_ACq/Pr_B7_DC` and the velocity as the derivative of this. Two pulses with opposite sign are then sent to the North-End mirror, with the length of the pulses proportional to the measured velocity. After sending the two pulses, the algorithm waits until the transmission of the cavity is above a certain threshold and then engages a linear feedback loop. If the resonance is not found within a certain timeout period, the algorithm starts all over. The logic of the algorithm is implemented in *C* as a simple state machine, see Appendix B. Further details are given in [11].

For the algorithm to work well, both the velocity measurement and the force of the actuator should be calibrated. In practice, since the length of the pulses is calculated using the product of these two calibration values, only a single value has to be tuned until the algorithm could lock the cavity. This value was easily found with a bit of trial-and-error, see Fig. 7 for an example of a successful lock acquisition. Once the cavity could be locked reliably, the maximum force was reduced (which means that the length of the pulses will become longer), until it was impossible to lock.

The minimum force with which the lock could be easily acquired was with a limit on the actuator voltage of $V_{\max} = 0.05\,\mathrm{V}$ (corresponding to a force of the actuator $F_{\max} = 0.18\,\mathrm{mN}$). The maximum value for a single

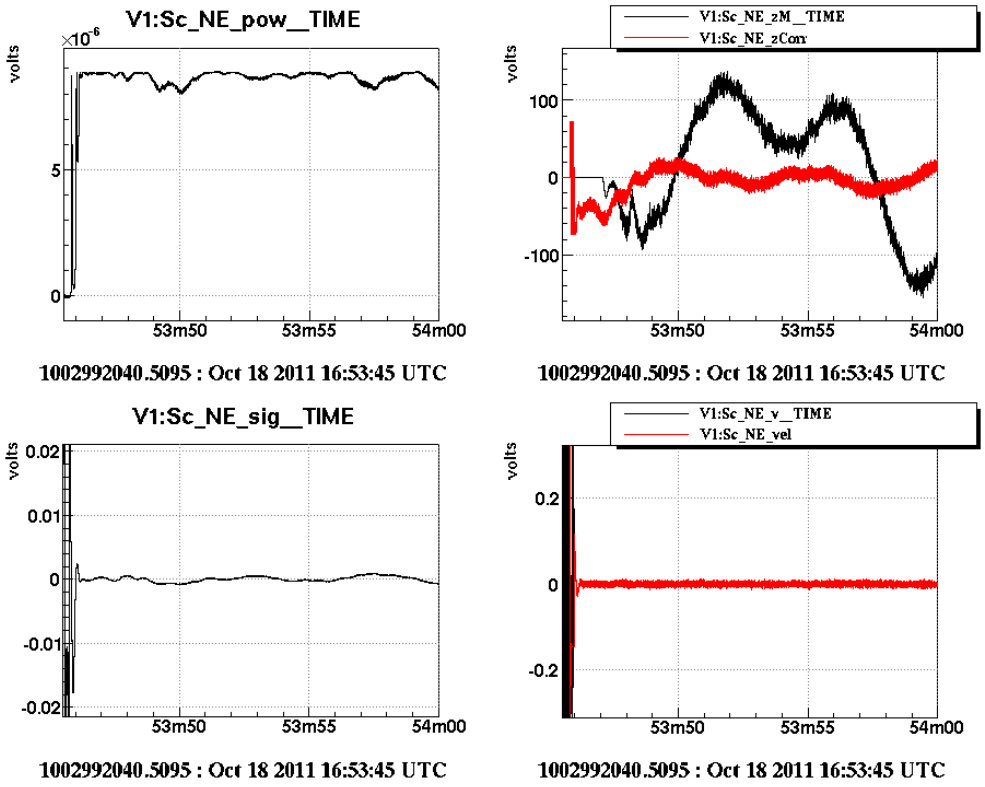

Figure 6: Demonstration of a lock with the DSP, in which the low-frequency part of the correction is reallocated to the marionetta after 1 second. The zCorr signal has been rescaled, it was limited to 0.02 V in this experiment.

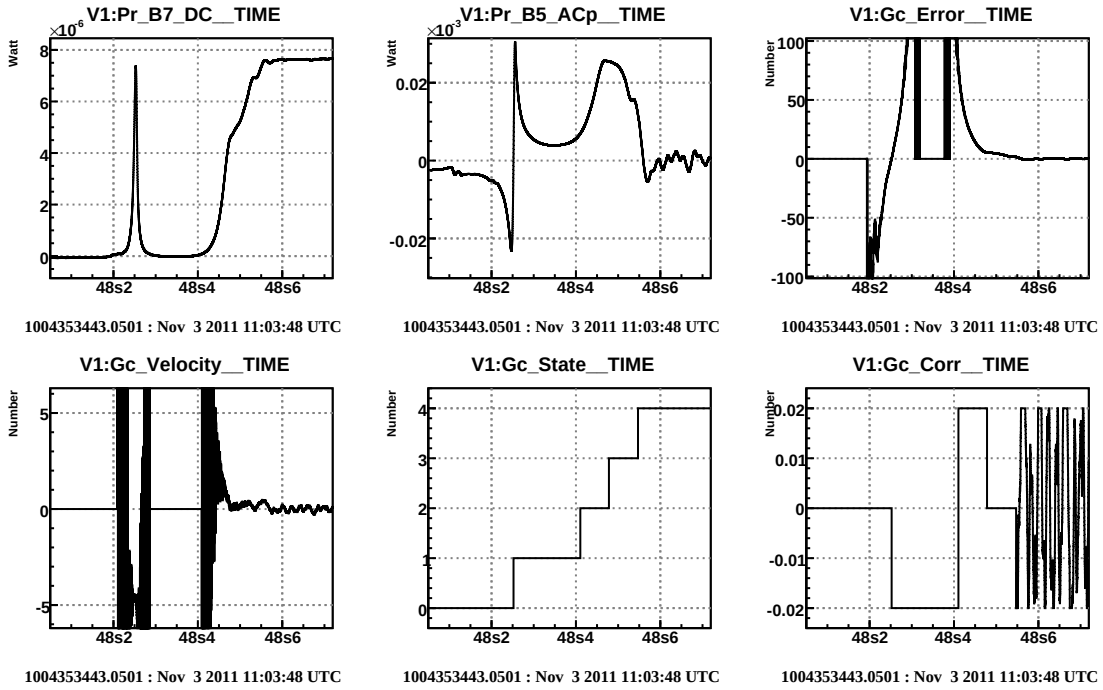

Figure 7: Demonstration of lock with Global Control. Using the length of the pulse and the known force of the actuator, the velocity of the first resonance crossing can be estimated to be around $0.3\,\mu$m/s.

coil-pair is $V_{\max} = 0.8\,\mathrm{V}$ ($F_{\max} = 2.9\,\mathrm{mN}$), so 16 times less than normally used. It is possible to use 4 coil-pairs per cavity, which means a factor of 64 lower than possible. Note that not as much time was spent fine-tuning the algorithm and no reallocation to the marionetta was used, as were done for the experiment with the DSP.

4.3 Simulations using E2E for Advanced Virgo

The experimental tests described in the last two sections were performed using the Virgo+ interferometer, so the obtained results are not directly applicable to Advanced Virgo. To make any quantitative statement about Advanced Virgo, we must rely on simulations. A model of a simple cavity with simplified suspensions was therefore implemented in E2E [13]. Since E2E allows the inclusion of blocks written in *C*, the two-pulse algorithm used in Global Control could be used without modification. Various tests were done with both the 1-pulse and 2-pulse algorithm. See Fig. 8 for an example of a simulated lock using the 1-pulse algorithm.

The mechanics of the mirror were simulated using a simulation block that is based on a model of the Initial LIGO suspensions. Its parameters were modified to resemble a mirror of Advanced Virgo, but this is still far from realistic. For the moment, we can thus not make any quantitative conclusions. This simulation must just be seen as a proof-of-principle. Once more details about the mechanics of the future payloads are available, we should be able to obtain quantitative results in a short time.

5 Conclusion

In this note, we discussed how the lock of a high-finesse cavity can be acquired with minimal force. We discussed some theoretical limits to the maximum velocity of a mirror in case a linear feedback is used. These limits were

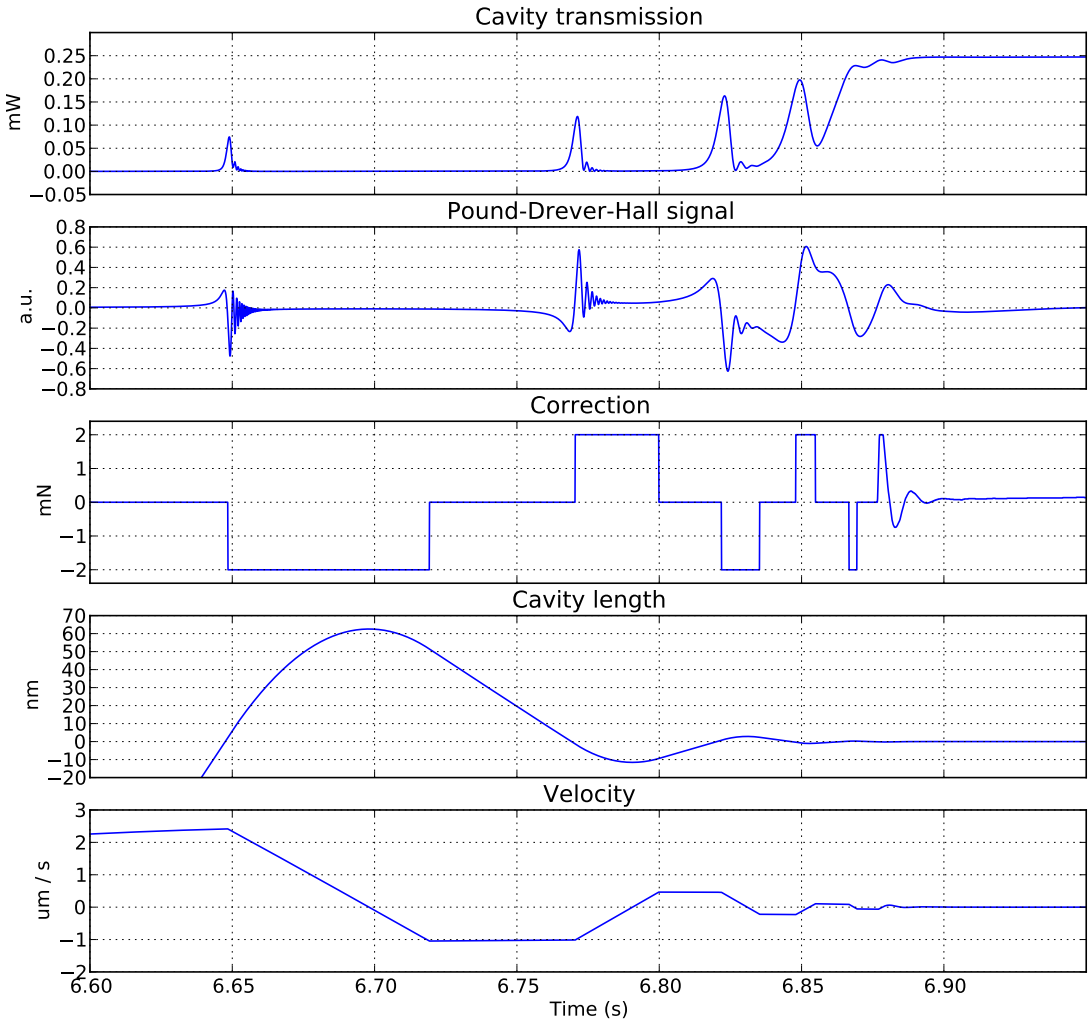

Figure 8: Example of a lock acquisition of an Advanced Virgo arm-cavity with the 1-pulse algorithm, as simulated with E2E. Note that the cavity length and velocity are properties that are only available in simulation and cannot be observed directly in reality.

not violated for Virgo and Virgo+, which is confirmed by our experience that it was always easy to lock. Some of these limits will probably be exceeded for Advanced Virgo, so it might become difficult or even impossible to lock with a conventional feedback loop.

One solution to this problem is *guided locking*, in which the velocity of the free-swinging mirror is measured during a resonance crossing. A long pulse is then sent using the actuators to bring it back at resonance back with reduced velocity. When the velocity is reduced enough, the lock can be acquired using a standard linear controller. This method has been tested by LIGO and TAMA in the past. As described in this note, we have experimentally tested a similar method using the Virgo+ interferometer and have implemented the algorithm in simulation with the parameters of Advanced Virgo, both with great success. The guided-lock method is thus extensively tested and easy to implement, so it can be used for Advanced Virgo without any problems.

In two different experiments, we showed that we could still acquire the lock while the maximum actuator force was artificially limited by a factor 40 of what is normally used. We were only using one out of the four available actuators, so we can claim that, in good weather conditions, we were operating with a safety margin of 160.

One crucial part of the method is to accurately measure the velocity when passing a resonance. We described experimental tests with Virgo+, which showed that we can accurately calibrate the measurement up to velocities of $10\,\mu m/s$. Simulations using parameters of Advanced Virgo showed that the same measurement should still work up to velocities that can be reached by pre-damping the mirrors with the local controls. What is still missing in the simulation is an accurate model of the noise of the photodiodes, which might limit the maximum velocity that can be measured and which might introduce larger errors in the measurement. Once more details of the detection system are available, we will have to verify in simulation that this does not affect the robustness of the guided lock algorithm.

The standard scheme in guided locking uses a single pulse. To reduce the time to lock, we investigated a modification of the algorithm that uses two pulses of opposite sign. A detailed theoretical analysis showed that the possible gain is small, so it is possibly not worth the effort. Anyhow, such variations are easy enough to implement and test, so a possible optimization of the algorithm can probably be done during the commissioning of the arm cavities.

For Advanced Virgo in power-recycled configuration, using the guided-lock algorithm is probably enough to guarantee a successful lock-acquisition without the use of an auxiliary laser system. Such a system might still be needed when we change to a double-recycled configuration.

5.1 Requirement on locking force

One of the parameters that should come out of this study is the maximum required locking force, which is needed for the design of the coils and the strength of the magnets glued to the mirror. We can only answer this question when an accurate model of the new payload is available. This model should include both the transfer-functions of the mirror and the marionetta, since it might be required to immediately reallocate the low-frequency part of the correction to the marionetta.

To determine the required force, requirements on the velocity due to the high finesse given in section 2 are no longer very relevant, since they can be fulfilled by using the guided-lock algorithm. That does not mean that we could lower the force of the actuators arbitrarily. One possible criterion might come from the total time to lock in relation to the relevant time-constants of the disturbances of the system, which is probably dominated by the 0.6 Hz pendulum mode and some low-frequency modes of the whole suspension chain. To not violate the assumption that the mirror is a free-moving object, we can thus require that the total time to lock is a large factor (> 10) smaller, so lets say $t_{\mathrm{total}} < 1/(0.6 \cdot 10) \approx 0.1\,\mathrm{sec}$. (Note that it would be OK if it takes 1 minute of failed lock attempts, followed by a single successful lock within 0.1 second). As shown in section A.3, the total time to reach a resonance with a 1-pulse algorithm is $t_{\mathrm{total}} = 4v_0/a = 4v_0 m/F_{\mathrm{max}}$. Solving for the force and using a typical velocity of $1\,\mu$m/s gives

$$F_{\mathrm{min}} = \frac{4v_0 m}{t_{\mathrm{max}}} = \frac{4 \cdot 10^{-6} \cdot 40}{0.1} = 1.6\,\mathrm{mN}. \tag{6}$$

This is probably an overestimate, since you could wait a bit for a very slow resonance crossing. The calculated value has to be compared with the currently available force of $4 \cdot 2.9 = 11.6\,\mathrm{mN}$.

In experiments with Virgo+, a very large safety margin on the force was found. Some amount of safety margin should be kept to account for bad weather conditions or problems like broken actuators. It is expected that we could reduce the maximum force by a small factor (say 5-10) without major problems. This will have to be confirmed using future simulations.

A Calculation of pulse lengths for the guided lock algorithm

In this section, we will calculate the length of the impulses needed to bring a mirror back to a resonance with a given velocity, which is needed in the guided-lock algorithm. We will also calculate the total time required to bring the mirror to a resonance with zero velocity, which is the parameter which we want to minimize.

It is assumed that the velocity of the mirror during a resonance crossing v_0 can be accurately measured using the method described in section 3. Furthermore, it is assumed that the mirror is a free-moving body, so that the force of the pendulum and the seismic disturbances can be ignored. This is valid if the time to stop the mirror is much smaller than the pendulum period. The mirror is only actuated at maximum force, which causes an acceleration of $a = F_{\mathrm{max}}/m$. All times, velocities and distances are assumed to be positive.

A.1 Single impulse

In the original guided-lock algorithm [6, 7], a single rectangular pulse is applied after the mirror crosses a resonance. To make the mirror come back with velocity $v_{\mathrm{end}} = \epsilon v_0$ (with $0 \leq \epsilon \leq 1$), a pulse of length t_{pulse} should be applied

$$t_{\mathrm{pulse}} = \frac{v_0 + v_{\mathrm{end}}}{a} = \frac{v_0}{a}(1+\epsilon). \tag{7}$$

At the end of this pulse, the mirror is away from the resonance by a distance x_{pulse}

$$x_{\mathrm{pulse}} = \frac{v_0^2}{2a} - \frac{v_{\mathrm{end}}^2}{2a} = \frac{v_0^2(1-\epsilon^2)}{2a}. \tag{8}$$

After this, one has to wait a time t_{drift} until the mirror drifts back to resonance

$$t_{\mathrm{drift}} = \frac{x_{\mathrm{pulse}}}{v_{\mathrm{end}}} = \frac{v_0^2(1-\epsilon^2)}{2a\epsilon v_0} = \frac{v_0}{a}\left(\frac{1}{2\epsilon} - \frac{\epsilon}{2}\right). \tag{9}$$

The total time to return to resonance t_{return} is thus

$$t_{\mathrm{return}} = t_{\mathrm{pulse}} + t_{\mathrm{drift}} = \frac{v_0}{a}\left(1 + \frac{\epsilon}{2} + \frac{1}{2\epsilon}\right). \tag{10}$$

Note that if a small return velocity is chosen, the time to return back to resonance will become very long.

A.2 Double impulse

As shown in the previous section, when using only a single pulse and choosing a low return velocity, it can take a long time before the mirror drifts back towards resonance. This time can be shortened by, after the desired velocity is reached, continuing to accelerate towards the resonance at maximum force. After a while, the actuation is changed to braking at maximum force, to arrive at resonance with the desired velocity.

This method can be divided in 3 periods: in the first part the mirror is slowed down to the desired target velocity, this is exactly equal to t_{pulse} of Eq. 7. In the second part, which lasts a time t_{accel}, the mirror is

accelerated further towards resonance, while in the last part that takes t_{brake}, the mirror is slowed down for the second time to the target velocity. Since the accelerating and braking happens with the same force, it follows that $t_{\text{accel}} = t_{\text{brake}}$. During these last two periods, the mirror moves over a distance

$$x_{\text{accel+brake}} = t_{\text{accel}}\left(v_{\text{end}} + at_{\text{accel}}/2\right) + t_{\text{brake}}\left(v_{\text{end}} + at_{\text{brake}}/2\right) = 2t_{\text{brake}}v_{\text{end}} + t_{\text{brake}}^2 a \tag{11}$$

Equating this to Eq. 8 and solving for a positive solution of t_{brake} yields

$$t_{\text{accel}} = t_{\text{brake}} = \frac{\sqrt{\frac{v_{\text{end}}^2}{2} + \frac{v_0^2}{2}} - v_{\text{end}}}{a} = \frac{v_0}{a}\left(\sqrt{\frac{\epsilon^2+1}{2}} - \epsilon\right). \tag{12}$$

The first two periods are combined into one pulse of the same sign, which lasts

$$t_{\text{pulse1}} = t_{\text{pulse}} + t_{\text{accel}} = \frac{v_0}{a}\left(\sqrt{\frac{\epsilon^2+1}{2}} + 1\right), \tag{13}$$

followed by a second pulse of opposite sign that lasts $t_{\text{pulse2}} = t_{\text{brake}}$. The total time to return to resonance t_{return} is finally

$$t_{\text{return}} = t_{\text{pulse1}} + t_{\text{pulse2}} = \frac{v_0}{a}\left(\sqrt{2\epsilon^2+2} + 1 - \epsilon\right). \tag{14}$$

The difference between the two 1-pulse and 2-pulse algorithms is clearly illustrated in Fig. 9, in which the return speed was chosen as $\epsilon = 0.2$ in both cases.

A.3 Comparison

During the time that the mirror returns to resonance, the motion of the mirror will be disturbed by the restoring force of the pendulum and by the residual seismic noise. It could thus return with a velocity that is totally different than expected or even change direction and move towards the next resonance. To reduce this effect as much as possible, we want to bring back the mirror to resonance as fast as possible.

To compare the single and double-pulse methods, the time to return to resonance is plotted as a function of ϵ, see Fig. 10a. When the return velocity is chosen as zero, the time for the 1-pulse method goes to infinity, while the double-pulse method reaches a minimum of $t_{\text{return}} = (1+\sqrt{2})V_0/a$. It can be proven that this is the fastest possible way to get back to resonance, given a limited force. The single pulse method reaches a minimum of $t_{\text{return}} = 2V_0/a$ for $\epsilon = 1$ (so the mirror returns with the original speed).

The time it takes to return to resonance is actually not very relevant, what is really important is the total time to bring the mirror at resonance with speed zero so that the linear lock can be engaged. When the mirror returns with reduced speed at resonance, the algorithm can be repeated so that the speed converges towards zero. The total time t_{total} can be calculated as

$$t_{\text{total}} = (1 + \epsilon + \epsilon^2 + \epsilon^3 + \ldots)t_{\text{return}} = \frac{t_{\text{return}}}{1-\epsilon}, \tag{15}$$

which is plotted in Fig. 10b. As before, the double pulse algorithm gives a minimum time of $t_{\text{total}} = (1+\sqrt{2})v_0/a$ for $\epsilon = 0$, while the single pulse method now shows a minimum of $t_{\text{total}} = 4v_0/a$ for $\epsilon = 1/3$.

The 2-pulse algorithm thus seems faster than the single pulse algorithm by a factor 1.65. In practice, due to uncertainties in measuring the velocity v_0 and due to disturbances, it is probably not desirable to give the 2-pulse algorithm a target velocity of zero. Because of these uncertainties, it is possible that the mirror stops slightly before resonance and then starts moving away from the resonance, so it could take a long time before it returns back to the same or the next resonance. To improve the robustness of the lock, the mirror should therefore be brought back with a lower but nonzero velocity, which increases the total time and thus further reduces the advantage with respect to the 1-pulse algorithm.

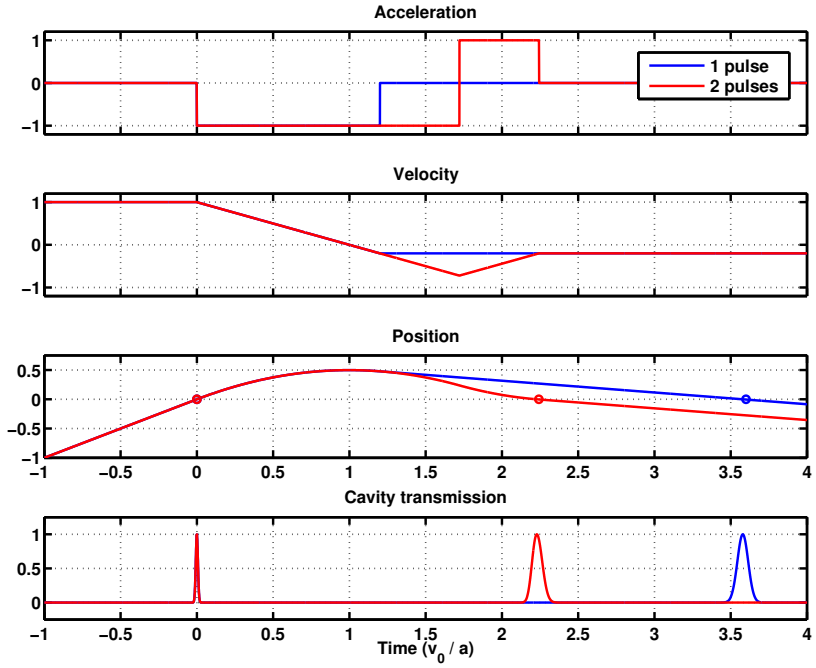

Figure 9: Illustration of acceleration (a), velocity (b), position (c) and transmission (d), for the 1 and 2 pulse algorithms. The return speed is chosen as $\epsilon = 0.2$.

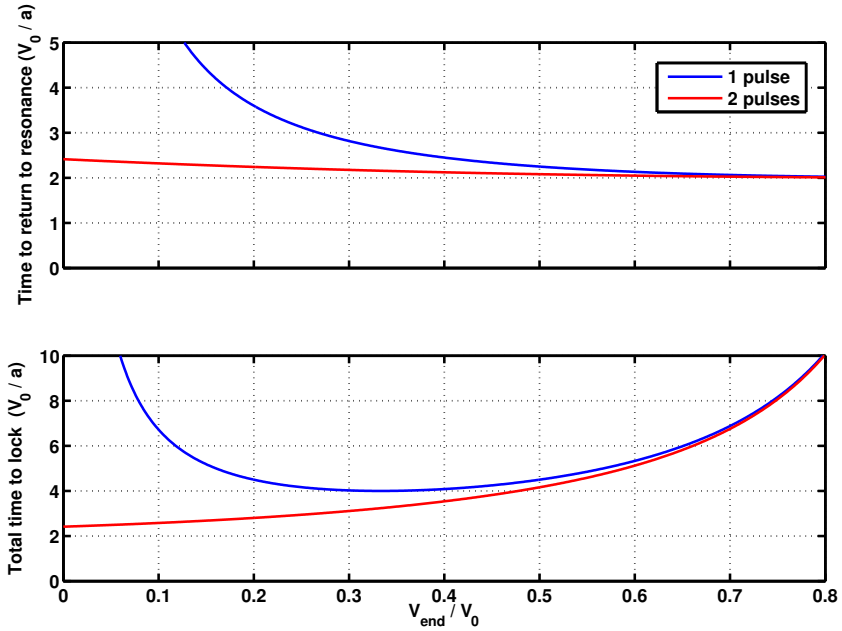

Figure 10: (a) Time to return to resonance and (b) total time to lock vs return speed, for 1 and 2 pulse algorithms

B Guided-lock algorithm in Global Control

Below is an excerpt of the code of GcLocking that implements the state-machine for the 2-pulse algorithm. The complete algorithm is implemented in `CavityTestSensing` that is part of `GcLocking` version `v2r14p1`. All capitalized variables are parameters that can be changed *on-fly*. The only inputs to the algorithm are the transmission of the cavity (`signals[D7_DC]`) and the Pound-Drever-Hall signal in reflection(`signals[D5_P]`), while the only output is the correction sent to the north-end mirror (`corr`).

```
dc = signals[D7_DC] / DC_MAX;  //calculate error signals
error_last = error;
if (dc > 1e-4) { //prevent divide by zero and insanely large values
  error = signals[D5_P] / signals[D7_DC];
} else {
  error = 0;
}
if (dc > TH_CALC) { //to not pollute filter
  velocity = derivator.compute_output(error) * CALI;
  derivator.update_registers();
} else {
  velocity = 0;
  derivator.reset_history();
}

switch(state) {        //state machine for north-cavity lock
case 0:                //uncontrolled, waiting for first resonance
  corr = 0;
  counter = 0;
  if (dc > TH_NONLIN && (error_last > 0 ^ error > 0)) { //zero-crossing of error signal
    V_save = velocity;  //TODO: protect for too large values
    counter = int(C1 * fabs(V_save));
    state = 1;
  }
break;

case 1:                // impulse 1
  corr = V_save > 0 ? CORR_MAX : -CORR_MAX;
  counter -= 1;
  if (counter <= 0) {
    counter = int(C2 * fabs(V_save));
    state = 2;
  }
break;

case 2:                // impulse 2
  //corr = - CORR_MAX * sign(V_save);
  corr = V_save > 0 ? -CORR_MAX : CORR_MAX;
  counter -= 1;
  if (counter <= 0) {
    counter = TIMEOUT;
    state = 3;
  }
break;
```

```
case 3:                 //wait for return to resonance
  corr = 0;
  counter -= 1;
  if (dc > TH_LINEAR) {// or check velocity again?
    state = 4;
  } else if (counter <= 0) { //timeout
    state = 0;
  }
break;

case 4:                 //linear lock
  corr = GAIN * controller.compute_output(error);
  controller.update_registers();
  corr = corr >  CORR_MAX ?  CORR_MAX : corr;     //clip
  corr = corr < -CORR_MAX ? -CORR_MAX : corr;
  counter = 0;
  if (dc < TH_UNLOCK) { //unlock
    state = 0;
    controller.reset_history();
  }
break;
} //switch(state)
```

References

[1] S. Frasca, E. Majorana, P. Puppo, P. Rapagnani and F. Ricci, *"Electromagnetic coupling dissipation between mirrors and reaction masses in Virgo"*, Physics Letters A, vol. 252, pp. 11-16 (1999).

[2] ESPCI-LAL-LAPP-LMA, *"CALVA - R&D report"*, Virgo technical note VIR-0299A-10 (2010).

[3] K. Izumi et. al., *"Multi-color Cavity Metrology"*, to be published (2012).

[4] L. Barsotti, *"The control of the Virgo interferometer for gravitational wave detection"*, PhD thesis, Università di Pisa (2006).

[5] M. Rakhmanov, *"Dynamics of laser interferometric gravitational wave detectors"*, PhD thesis, California Institute of Technology (2000).

[6] J. Camp et al, *"Guided lock acquisition in a suspended Fabry-Perot cavity"*, Optics Letters, vol. 20 (24), pp. 2463-2465 (1995).

[7] K. Izumi, *"Lock acquisition of Fabry-Perot cavity with guided lock"*, presentation at GWADW conference, Elba (2008).

[8] L. Matone, *"Étude du contrôle global de l'interféromètre central de Virgo"*, PhD thesis, Université Paris XI (1999).

[9] P. Ruggi, G. Vajente, D. Passuello and E. Calloni, *"Low force mirror damping"*, Virgo logbook entry #30431 (2011).

[10] P. Ruggi, *"Low force cavity lock"*, Virgo logbook entry #30433 (2011).

[11] B. Swinkels and G. Vajente, *"Non-linear lock acquisition with reduced force"*, Virgo logbook entry #30472 (2011).

[12] B. Swinkels and R. Day, *"Analysis of free-swinging cavity"*, Virgo logbook entry #20752 (2008).

[13] H. Yamamoto et al., *End to End* time-domain optical simulator http://www.ligo.caltech.edu/~e2e/