

Article

Not peer-reviewed version

Adaptive Time-Domain Simulation of Optical Cavities with Arbitrary Dynamics

[A. Svizzeretto](#), [J. Casanueva Diaz](#), B. L. Swinkels, [M. Bawaj](#)*

Posted Date: 15 May 2026

doi: 10.20944/preprints202605.0969.v1

Keywords: Fabry-Perot cavity simulation; Pound-Drever-Hall sensing; reinforcement learning control



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Adaptive Time-Domain Simulation of Optical Cavities with Arbitrary Dynamics

A. Svizzeretto ^{1,2} , J. Casanueva Diaz ³ , B. L. Swinkels ⁴  and M. Bawaj ^{1,2} *

¹ Università di Perugia, I-06123 Perugia, Italy

² INFN, Sezione di Perugia, I-06123 Perugia, Italy

³ European Gravitational Observatory (EGO), I-56021 Cascina, Pisa, Italy

⁴ Nikhef, 1098 XG Amsterdam, The Netherlands

* Correspondence: mateusz.bawaj@unipg.it

Abstract

We present a fast time-domain simulator for optical cavities capable of reproducing non-linear dynamical regimes arising from ring-down effect during resonance crossings at high mirror velocities. The model is based on a recursive formulation of the intracavity electric field as a sum over round trips, preserving the cavity memory while maintaining high computational efficiency. The simulator is designed to achieve three main goals. First, the boundary conditions of the cavity can be modified at each simulation step, allowing arbitrary time-dependent variations of both mirror positions and input electric field. Second, the sampling frequency can be flexibly chosen by the user, however, it is internally adjusted before effectively executing the simulation to remain consistent with the cavity round-trip structure. Finally, high computational efficiency was obtained by avoiding the repeated evaluation of the full electric field history. The framework is validated through comparison with experimental data from the Virgo interferometer during a mechanical excitation experiment, showing good agreement in non-adiabatic regimes. Due to its efficiency and flexibility, the simulator provides a versatile tool for time-domain studies of optical resonators and future applications in real-time control and reinforcement-learning-based lock acquisition.

Keywords: Fabry–Perot cavity simulation; Pound–Drever–Hall sensing; reinforcement learning control

1. Introduction

Time-domain simulation of optical cavities is required whenever the assumptions underlying frequency-domain or adiabatic models break down, such as in the presence of fast transients, short optical pulses, or rapidly varying boundary conditions. In these regimes, the intra-cavity field evolution depends explicitly on the finite round-trip time and the storage of optical energy [1].

This is particularly relevant in precision interferometry and gravitational-wave detectors, where the interaction between delayed intra-cavity fields and moving mirrors leads to non-linear, history-dependent dynamics that cannot be captured by stationary transfer functions [2,3].

Beyond interferometric gravitational-wave detectors, similar requirements arise in pulsed cavity ring-down spectroscopy, fast cavity tuning experiments, and active stabilization of high-finesse resonators, where transient dynamics and non-adiabatic effects determine system performance [4]. In all these contexts, such simulations provide direct access to transient regimes and provide a natural interface for modern data-driven control strategies.

Legacy interferometer simulation frameworks like e2e and SIESTA provide detailed optical modelling capabilities [5,6]. However, their use in real-time and reinforcement-learning applications is limited by their complexity and software design. The simulator presented here is designed to address transient regimes with minimal approximations for a wide range of cavity parameters. It explicitly propagates the intra-cavity field as a sum over round trips, retaining the full temporal memory of the cavity [7,8]. Crucially, both the input optical field and the mirror positions can be

modified at every simulation step. This allows arbitrary laser amplitude and phase modulation, as long as their characteristic frequencies remain below the simulator sampling frequency, as well as time-dependent cavity length variations under the same condition.

This level of flexibility is essential for our primary target application: integration into a Gymnasium-compatible reinforcement-learning environment for training real-time artificial intelligence (AI) agents dedicated to lock acquisition of suspended optical cavities currently under development [9,10]. Lock acquisition is inherently a non-stationary control problem, dominated by transients, delays, and non-linear field-mirror coupling [11]. Its stepwise formulation naturally interfaces with discrete-time control algorithms, enabling realistic training conditions and supporting an effective sim-to-real transfer [12].

In addition to the theoretical and numerical developments presented here, the simulator is released as open-source software [?] under a separate license to ensure transparency and reproducibility.

The remainder of the paper is organized as follows. Sec. 2 introduces the numerical framework and the recursive formulation used to propagate the intracavity field, together with the implementation details of the simulation strategy. In particular we discuss the adaptive selection of the sampling frequency and the different integration regimes.

Sec. 3 presents the validation of the simulator through a comparison with ring down experimental data from the arm cavity of the Virgo interferometer cavity, while Sec. ?? extends the discussion to the broader applicability of the time-domain formulation to optics and photonics in general, as well as to machine learning (ML)-aided control strategies.

Finally, Sec. 5 summarizes the main results and outlines potential applications of the simulator, particularly in the context of real-time control and reinforcement-learning-based lock acquisition strategies.

Additional supporting material is provided in the Appendix, including the full comparison between simulated and experimental data from Virgo auxiliary channel over a 10 s time window, a reproduction of the pulsed cavity ring-down regimes reported by [4], and the complete benchmark table used to characterize the computational performance of the simulator.

2. Algorithm

The algorithm of the simulator is capable of reproducing non-linear dynamical effects in optical cavities caused by the ring-down phenomenon. The temporal evolution of the electric cavity field is described in [7] by the recursive equation:

$$E(t) = t_a \sum_{n=0}^{N-1} \left[(r_a r_b)^n e^{-2ikS_n(t)} E_{\text{in}}(t - 2nT) \right] + (r_a r_b)^N e^{-2ikS_N(t)} E(t - 2NT), \quad (1)$$

where, t_a is the amplitude transmissivity of the input mirror, r_a and r_b are the amplitude reflectivities of the input and end mirrors, respectively, $k = 2\pi/\lambda$ is the optical wave-number associated with the laser wavelength, and $T = L/c$ is the one-way propagation time inside the cavity of length L . $E_{\text{in}}(t)$ denotes the input optical field, while $E(t)$ is the intracavity electric field. The integer n labels the successive cavity round trips and N is the total number of round trips retained in the calculation. Eventually, the quantity $S_n(t)$ represents the accumulated optical path over the previous n round trips and accounts for the mirror motion and cavity length variations,

$$S_n(t) = \sum_{p=0}^{n-1} d(t - 2pT), \quad (2)$$

specifically, $d(t) = L + (x_b(t - T) - x_a(t))$ where $x_b(t)$ and $x_a(t)$ are the end and input mirror positions. Equation (1) can be used to obtain the optical power $P(t) = |E(t)|^2$ and the Pound-Drever-Hall (PDH) error signal [13], designed for optical cavity control at resonance. In the approximation where

the modulation sidebands are fully reflected by the cavity and act as a stable reference field, the demodulated PDH signal can be expressed as:

$$V_{\text{PDH}}(t) = -\text{Im}\{e^{i\gamma} E_{\text{in}}(t)^* E(t)\}, \quad (3)$$

where $V_{\text{PDH}}(t)$ is the PDH error signal, and γ is the demodulation phase. $E_{\text{in}}(t)$ is the complex conjugate of the input optical field, while $E(t)$ is the intracavity electric field.

These signals are the optical observables of the cavity state and become strongly non-linear under specific conditions, thereby making standard control technique for locking ineffective. In particular, ring down effect occurs when the mirror velocity becomes sufficiently large. Thresholds to this speed, referred to as critical velocities, can be defined [14,15]. In the present treatment, we consider, as a limit to define the non-linear regime, only the threshold obtained by equating the resonance crossing time to the cavity storage time, which represents the average time a photon remains stored in the cavity. This condition corresponds to the critical velocity defined as follows:

$$v_{\text{cr}} \approx \frac{\lambda \pi c}{4L\mathcal{F}^2}, \quad (4)$$

where, λ is the laser wavelength, c is the speed of light, L is the cavity length, and \mathcal{F} is the cavity finesse.

The simulator exploits the structure of Equation (1) to enable the selection of a sampling frequency over a wide range, as well as the independent motion of both the input and end mirrors. The input mirror displacement is modelled as a phase shift of the input laser field, as adopted in the algorithm.

Because the cavity field equation is recursive, each new sample depends only on an internal simulator state rather than on a full re-evaluation of the entire optical history. The cavity memory is therefore represented by a finite buffer of past input fields together with a single stored values of fields and mirrors position. This mechanism supports incremental, state-space-like updates instead of repeated global re-integration. Geometric electric field decay and phase factors are precomputed once during initialization and reused at every step, so the per-step computational load is dominated by multiply-accumulate operations.

The simulator allows the user to select the sampling frequency, but its final value is approximated according to an internal optimization rule designed to preserve consistency with the cavity round-trip time. We will refer to the sampling frequency desired by the user as $f_{\text{calc}}^{\text{desired}}$. The effective sampling frequency f_{calc} is computed through a rule-based decision tree that compares $f_{\text{calc}}^{\text{desired}}$ with $f_{2T} = 1/2T$, the round trip frequency. The main algorithm pseudocode is shown in Alg. 3. Alg. 1 and Alg. 2 clarify rounding strategies.

Algorithm 1 number_of_subhistories(η_{2T})

Require: $\eta_{2T} > 0$

- 1: **if** $\eta_{2T} < 1$ **then**
 - 2: **return** round($1/\eta_{2T}$)
 - 3: **else**
 - 4: **return** 1
 - 5: **end if**
-

In particular, Alg. 2 implements a custom quantization rule that maps the continuous variable η_{2T} to an integer order k used by the simulator. It computes $k_0 = \lfloor \eta_{2T} \rfloor$, evaluates the switching threshold

$$b = \frac{2k_0(k_0 + 1)}{2k_0 + 1},$$

and returns $k = k_0$ if $\eta_{2T} < b$, otherwise $k = k_0 + 1$. Alg. 1 and Alg. 2 are both used to round the dimensionless round-trip normalized time-step defined as $\eta_{2T} = (f_{\text{calc}}^{\text{desired}} 2T)^{-1}$.

Algorithm 2 round_for_inverse_curve(η_{2T})

Require: $\eta_{2T} > 0$

- 1: $k_0 \leftarrow \lfloor \eta_{2T} \rfloor$
- 2: $b \leftarrow \frac{2k_0(k_0 + 1)}{2k_0 + 1}$
- 3: **if** $\eta_{2T} < b$ **then**
- 4: $k \leftarrow k_0$
- 5: **else**
- 6: $k \leftarrow k_0 + 1$
- 7: **end if**

Within the full algorithm, this function is the core discretization step. Its role is to choose the integer round-trip count N that best approximates the requested sampling condition, thereby controlling the trade-off between physical fidelity and numerical feasibility of the final computed sampling frequency $f_{\text{calc}} = f_{2T}/N$. We will come back to this function further in the description.

Alg. 3 checks if $f_{\text{calc}}^{\text{desired}} > f_{2T}$, meaning the requested time-step is shorter than a photon round-trip time. In this case, the algorithm interleaves multiple sub-histories, each shifted in phase to increase the effective sampling, and adjusts f_{calc} to an integer multiple of the round-trip frequency.

For $f_{\text{calc}}^{\text{desired}} < f_{2T}$, we distinguish between two regimes. The first case occurs when the desired sampling frequency becomes too low compared to the cavity storage time; the corresponding integration window may exceed the effective memory of the cavity. In this case, the number of round-trips included in the summation is capped to a maximum value $N_{\text{max}} = 5N_{\text{eff}}$, where the effective photon lifetime expressed in round-trips is defined as

$$N_{\text{eff}} = \frac{1}{|\ln(r_a r_b)|}.$$

The desired sampling frequency is then preserved, while the integration window is allowed to contain a fractional contribution of the last round-trip, ensuring that negligible long-past contributions are not propagated. This strategy guarantees temporal consistency while balancing numerical efficiency against the physical memory of the cavity.

The second case corresponds to intermediate frequencies, where an integer number of round-trips is selected using an inverse-curve rounding rule that finds optimal N , which minimizes the error with respect to the desired sampling frequency. In this case, the effective sampling frequency is set to $f_{\text{calc}} = f_{2T}/N$.

To quantify the impact of the internal frequency-selection strategy, we evaluate the accuracy of the effective sampling frequency with respect to the user-requested value. The result is shown in Figure 1. The plot shows the relative accuracy

$$A_f = 1 - \frac{|f_{\text{calc}} - f_{\text{calc}}^{\text{desired}}|}{f_{\text{calc}}^{\text{desired}}},$$

highlighting the minimal achievable deviation introduced by the approximation to values compatible with the cavity round-trip time.

The simulator workflow can be summarized in a few simple stages. First, the cavity is defined by specifying its physical parameters: mirror reflectivities, transmissivities, and cavity length, together with the desired sampling frequency $f_{\text{calc}}^{\text{desired}}$. During the initialization phase, the internal frequency-selection rule determines the effective sampling frequency f_{calc} , the corresponding time step Θ , and the number of round-trips included in each integration step. The effective sampling frequency is then returned to the user and must be used as the temporal grid of any external simulation environment. The simulation proceeds iteratively, updating the intracavity field at each step through the recursive field equation.

Algorithm 3 `sampling_frequency_selection`($f_{\text{calc}}^{\text{desired}}, f_{2T}, N_{\text{eff}}$)

```

1:  $N \leftarrow 1$ 
2:  $n_{\text{sub}} \leftarrow 1$ 
3: partial_Theta  $\leftarrow$  False
4:  $_{N_{\text{eff\_factor}}} \leftarrow 5$  ▷ field amplitude has decayed to  $\approx 0.67\%$  of initial value
5:  $\eta_{2T} \leftarrow f_{2T} / f_{\text{calc}}^{\text{desired}}$  ▷ dimensionless round-trip normalized time-step
6: if  $f_{\text{calc}}^{\text{desired}} > f_{2T}$  then
7:    $n_{\text{sub}} \leftarrow \text{number\_of\_subhistories}(\eta_{2T})$ 
8:    $f_{\text{calc}} \leftarrow f_{2T} \cdot n_{\text{sub}}$ 
9: else
10:   $N_{\text{max}} \leftarrow _{N_{\text{eff\_factor}}} \cdot N_{\text{eff}}$ 
11:  if  $f_{\text{calc}}^{\text{desired}} < \frac{f_{2T}}{N_{\text{max}}}$  then
12:     $N \leftarrow N_{\text{max}}$ 
13:     $f_{\text{calc}} \leftarrow f_{\text{calc}}^{\text{desired}}$ 
14:    partial_Theta  $\leftarrow$  True
15:  else
16:     $N \leftarrow \text{round\_for\_inverse\_curve}(\eta_{2T})[0]$ 
17:     $f_{\text{calc}} \leftarrow f_{2T} / N$ 
18:  end if
19: end if
20:  $\Theta \leftarrow 1 / f_{\text{calc}}$ 
21: return ( $f_{\text{calc}}, \Theta, N, \text{partial\_Theta}$ )

```

3. Validation and Benchmarks

The simulator performance is evaluated through comparison with experimental data acquired from the north arm cavity of the Virgo interferometer [16,17] during mode matching measurement using the external excitation method, which consists of applying corrections to excite mechanically the end mirror in the longitudinal degree of freedom and analysing the second order mode during free swinging of the test mass. This mechanical excitation produces non-adiabatic crossings of the cavity resonances, exposing ringing effect for different crossing velocities. These represents optimal conditions for validating the simulation framework.

The selected dataset corresponds to a 10 s time window of the DC transmission power signal recorded at the cavity auxiliary B7 photodetector. Together with the transmission power, PDH error signal data from the B4 photodetector in reflection has been collected to compare with simulated data obtained from Equation (3).

The experimental data are shown in Figure 2 where it is also depicted the result of its processing. First, the first sample timestamp was set to zero. Then the transmission was detrended by subtracting its median value. Finally, resonance peak identification was performed with an algorithm based on a prominence-based peak detection: only peaks above 99th percentile of the detrended signal were retained. Each detected peak corresponds to a resonance condition $2L(t_n) = N_n \lambda$, with $\lambda = 1064$ nm. Successive resonances therefore correspond to cavity length variations of $\Delta L = \lambda/2$.

An integer resonance index n was assigned to each detected peak to reconstruct the cavity-length variation. The direction of the first crossing was determined from the sign of the local derivative of the PDH signal, after which the index was increased or decreased by one at each subsequent resonance. Turning points were identified based on a threshold for the transmission, which will be higher when the mirror velocity becomes small before changing direction. At these points, the sign of the index increment was reversed. The absolute cavity length for each n was reconstructed and expressed in terms of unity of Free Spectral Range (FSR). The discrete length samples were interpolated using a cubic spline to obtain a continuous estimate $L(t)$ over the oscillation window. A cubic-spline interpolation provides a smooth reconstruction of the resonance positions, but it is not equivalent to a sinusoidal model. Unlike a sinusoidal fit, the spline does not impose a global physical model on the




Figure 1. Accuracy of the effective calculation frequency f_{calc} with respect to the user-defined frequency $f_{\text{calc}}^{\text{desired}}$, as a function of the normalized ratio $f_{\text{calc}}^{\text{desired}}/f_{2T}$.

mirror motion and may therefore introduce non-physical distortions in the derived velocity. For this reason, spline-based reconstruction is useful as a flexible interpolation tool, but can affect the velocity estimation.

The instantaneous mirror velocity was computed from the time derivative of the reconstructed cavity length $v(t) = dL(t)/dt$.

The computation of the simulated data was performed by initializing the simulation with cavity parameters consistent with arm cavity of the Virgo interferometer [18]: mirror reflectivities $R_a = 0.986$, $R_b = 0.99999$, input transmission $T_a = 0.01377$, and nominal length $L = 3000$ m. The input laser field was set to a mean value of $|E_{\text{in}}| = 12.7$, obtained by the least squares method fit described below. The calculation frequency was matched to the sampling frequency of the raw data which is 10 kHz, ensuring temporal resolution well below the resonance crossing time. It is important to specify that in this case the actual calculation frequency approximated by the simulation strategy was ≈ 9.99 kHz, with an approximation accuracy of 99.9%. With this values only one history of the optical field evolution was evaluated and the regime was the aggregated round-trips one, considering the full Equation (1).

The reconstructed velocity was used as input to the simulation, defining a time-dependent cavity length variation. The simulated cavity power was obtained from the intra-cavity electric field through the output mirror transmissivity.

For detailed analysis, one resonance has been chosen, corresponding to $\approx 4.9 \mu\text{m s}^{-1}$, *i.e.*, a factor of ≈ 7.6 critical velocities. The compared data are showed in Figure 3

In addition to the DC transmission, the simulated and reconstructed PDH error signals exhibit consistent zero-crossing slope and phase structure, confirming that the simulator correctly captures the dynamics induced by high-velocity mirror motion.

In the appendix we present an additional plot for validation of the simulator which consists of a comparison of the whole experimental data time window (10 s) with some observations about integration errors.

In order to find the laser input power, Equation (1) can be approximated by a sum of adiabatic contribution plus a Doppler term:

$$E(t) \approx D_0 \exp\left(-\frac{t}{\tau} - i\frac{kv}{2T}t^2\right) + \frac{t_a E_{\text{in}}}{1 - r_a r_b e^{-2ikvt}}, \quad (5)$$

where, D_0 is the complex amplitude of the transient contribution, τ is the cavity storage time, v is the mirror velocity, and E_{in} is the constant incident field. The prefactor D_0 is given by:

$$D_0 = t_a E_{in} \sqrt{\frac{i\pi}{2kvT}} \exp\left(\frac{iT}{2kv\tau^2}\right). \quad (6)$$

This approximation holds in the length-scan transient regime for $v \geq v_{cr}$. Its derivation assumes a continuum limit of the discrete round-trip dynamics $n \gg 1$ and timescales $t \gg T$. The immediate resonance region is neglected [7]; therefore, the least-squares fit is performed only on the tail of the ring-down.

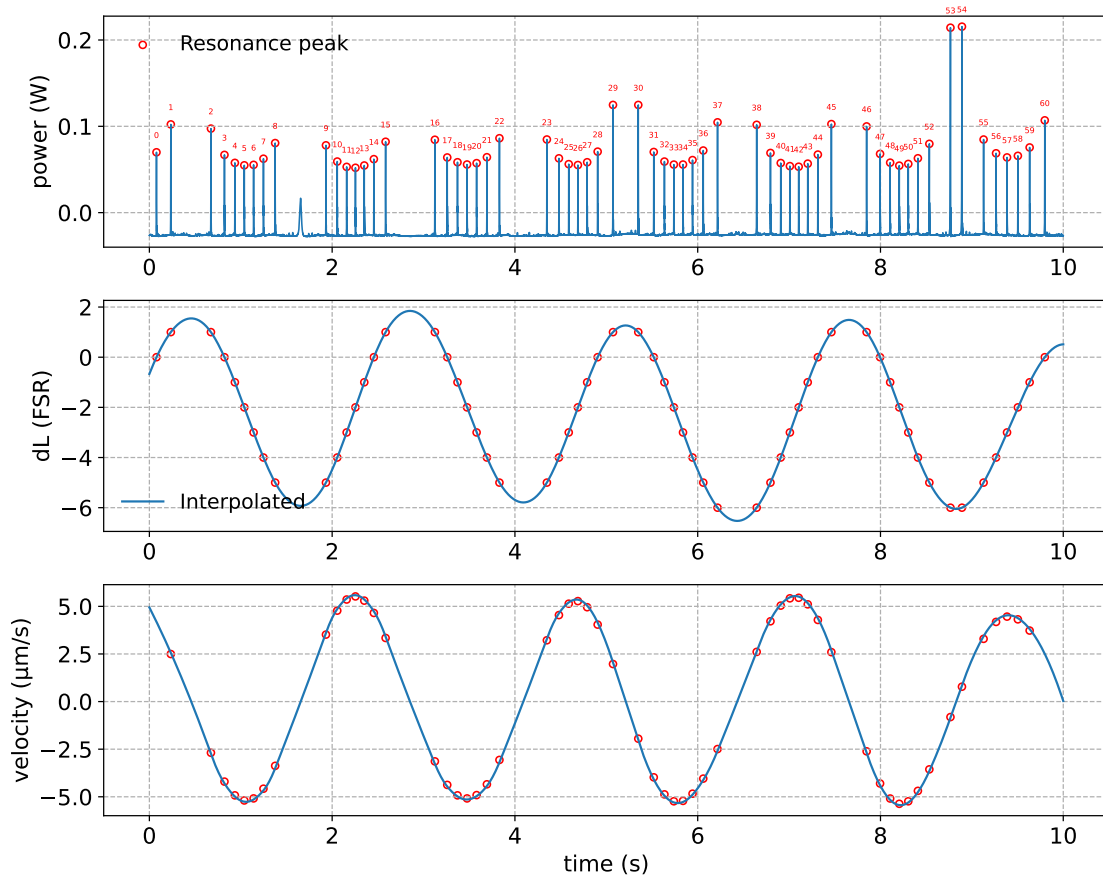


Figure 2. Cavity length and velocity reconstruction during mirror oscillation following a mechanical excitation. All panels share the same horizontal axis, representing time. The upper panel shows the recorded transmission peaks in the experimental data; peaks used in the analysis are indicated by circles. The middle panel depicts the cavity length in FSR units and the corresponding polynomial fit. The bottom panel shows the cavity mirror velocity derived from the cavity length. The first point and the last point of this pane has been cut due to edge effect caused by numerical derivative calculation.

In addition to the physical validation against experimental data, the computational performance of the simulator has been systematically assessed through a dedicated benchmark study. The benchmark on almost 400 cases compares the pure Python and Numba-accelerated backends on a test set of 99 cavity configurations spanning different lengths, reflectivities, and calculation frequencies. In all cases, the execution time scales approximately linearly with the number of round trips N , reflecting the dominant cost of the summation over stored field history. The pure Python implementation shows a clear increase in runtime with growing N , reaching several hundred microseconds for high-finesse, short cavities. In contrast, the Numba backend exhibits a much weaker dependence on N , reaching 16 μs for the most demanding cavity tested, with $N = 300$.




Figure 3. Comparison between experimental data and simulation for a single cavity resonance crossing. The top panel shows the experimental and simulated DC transmission, while the bottom panel displays the corresponding normalized PDH error signal. The simulation is driven by the reconstructed mirror motion and reproduces both the amplitude and temporal structure of the resonance, using a critical velocity scaling factor of 7.6, which corresponds to $4.9 \mu\text{m s}^{-1}$

As a result, the speed-up factor increases systematically with N , ranging from 4.67 for $N \approx 1$ up to 17.61 for $N = 300$, demonstrating that just-in-time (JIT) compilation is necessary for efficient simulations in the high-finesse regime. Detailed results are shown in Table A1.

4. Relevance to Optics and Photonics

Beyond Fabry-Perot (FP) cavities, the time-domain simulator developed in this work may be of relevance to a broader class and physical systems characterized by energy storage, ringing effects, and finite propagation time, all of which exhibit memory-driven dynamics analogous to those of optical cavities. Within this framework, the same delay-based formulation naturally extends to integrated optical resonators, such as ring resonators [19], which can be treated as a unidirectional cavities where the intra-cavity field is recursively re-injected after each round-trip with attenuation and phase accumulation. Ring resonators constitute key building blocks of integrated photonic, including frequency-comb micro-resonators, Kerr non-linear devices, optical buffers, all-optical switching architectures, and quantum photonic circuits.

As an additional validation of the generality of our time-domain approach, we reproduce the characteristic ring-down dynamics obtained in [4], namely the transition between pulsed and smooth exponential decay regimes as a function of pulse duration and detuning. Figure A2 in the Appendix A presents simulated ring-down signals for different normalized pulse durations and detunings, showing qualitative agreement with the temporal feature described in [4], including the emergence of transient peaks and intensity modulation arising from the partial temporal overlap of successive photon round trips.

Another promising application of the simulator lies in the optimization of locking schemes for extremely high-finesse cavities ($\mathcal{F} \sim 10\,000$). In this regime, resolving the temporal evolution of the intra-cavity field is crucial to understand transient dynamics and control performance. Machine learning techniques, particularly reinforcement learning (RL), can exploit this information to identify optimal control strategies, avoiding direct interaction with experimental hardware, which can be both time-consuming and potentially harmful [9,10].

Finally, such simulations can support the design and optimization of gravitational-wave interferometers such as Virgo and Einstein Telescope, improving their design and further upgrades [20–22]. In particular, in a situation where phenomena like transient dynamics, mirror motion, radiation pressure effects, play a key role in maintaining stable operation.

5. Conclusions and Applications

In this work we presented the logic and the validation of a fast time domain optical simulator, capable of reproducing the dynamical evolution of the intra-cavity field in optical resonators. The simulator allows the user to flexibly define cavity parameters, boundary conditions, and mirror motion.

Computational efficiency is achieved through a simple yet effective simulation strategy, in which the user-defined sampling frequency is approximated to the nearest value compatible with internal optimization constraints. This adjustment ensures consistency between the discrete time step and the cavity round-trip time, preserving causal propagation while avoiding unnecessary evaluations of redundant round-trip histories. Efficiency is further improved thanks to the use of JIT compilation.

The resulting framework is designed to provide a computationally efficient and physically consistent simulation tool, spanning applications from large-scale interferometric systems to integrated photonic devices. Computational efficiency is particularly important for predictive modelling, but it becomes even more critical when considering real-time control development and design of RL-based lock acquisition protocols. Lock acquisition in suspended optical cavities is intrinsically a non-stationary control problem, in which the cavity state is observed through signals that deviate from the linear regime assumed by classical controllers.

The feature of a step-by-step evolution of the intracavity field and its associated observables (DC transmission, PDH error signal), under arbitrary time-dependent mirror motion and input field modulation, naturally matches the interaction paradigm required by modern RL algorithms, where an agent observes the system state at discrete time steps, applies an action (e.g., actuation on mirror position or laser frequency), and receives a reward based on the effectiveness of its actions in reaching the desired goal.

The development of framework empowered with this simulator enables large-scale training campaigns, in which thousands of interaction steps are required to learn robust policies. Moreover, addressing one of the biggest challenge in RL applications, the so called sim-to-real transfer, requires high-fidelity simulation of the physical system. Such fidelity is mandatory to train RL agents under realistic dynamics, thereby ensuring a safe and more reliable deployment on the real optical setups.

Funding: Research is supported by Italian Ministry of University and Research (MUR) through the program "Dipartimenti di Eccellenza 2023-2027" (Grant SUPER-C), Italy. This work was partially supported by the European Union's Horizon Europe programme under the Marie Skłodowska-Curie Actions Staff Exchanges project GRAVITY (G.A. 101236384).

Acknowledgments: The authors acknowledge the use of technical data and interferometer parameters from the Virgo Collaboration for validation and comparison purposes. The authors thank the collaboration members involved in detector commissioning and characterization activities for making these data available.

Appendix A. Auxiliary Materials for Validation

Figure A1 shows the comparison between the measured DC transmission during the mode-matching measurements and the corresponding simulated data for the full 10s time window. It is important to note that the simulator is not specifically designed for long-time propagations, as the numerical implementations relies on discrete time steps and finite round-trip summations. Consequently, small integration errors, together with finite numerical precisions, accumulate over time and may lead to a gradual degradation of accuracy.

Despite these limitations, the agreement is remarkably good: the simulated response accurately reproduces both the sequence and relative amplitudes of the resonance crossings.

Part of the residual discrepancy also arises from the resonance identification procedure described in Sec. 3. Resonance positions were estimated by selecting transmission peaks above the 99 % percentile of the signal and assigning them as the exact resonance crossing times. This assumption becomes less accurate as the mirror velocity increases. In the non-adiabatic regime, the first power peak does not, in general, temporarily coincide with the exact resonance position. This temporal offset propagates into the reconstructed cavity length and, consequently, into the velocity estimation.

Therefore, the observed deviations between simulation and experiment should be interpreted as the combined effect of numerical integration errors and systematic uncertainties in the resonance localization procedure.




Figure A1. Comparison between the measured DC transmission of the North arm cavity following a mechanical excitation of the suspended mirror and the corresponding simulation. The simulated response reproduces the sequence of resonance crossings induced by the mirror oscillation, using the reconstructed speed evolution as input.

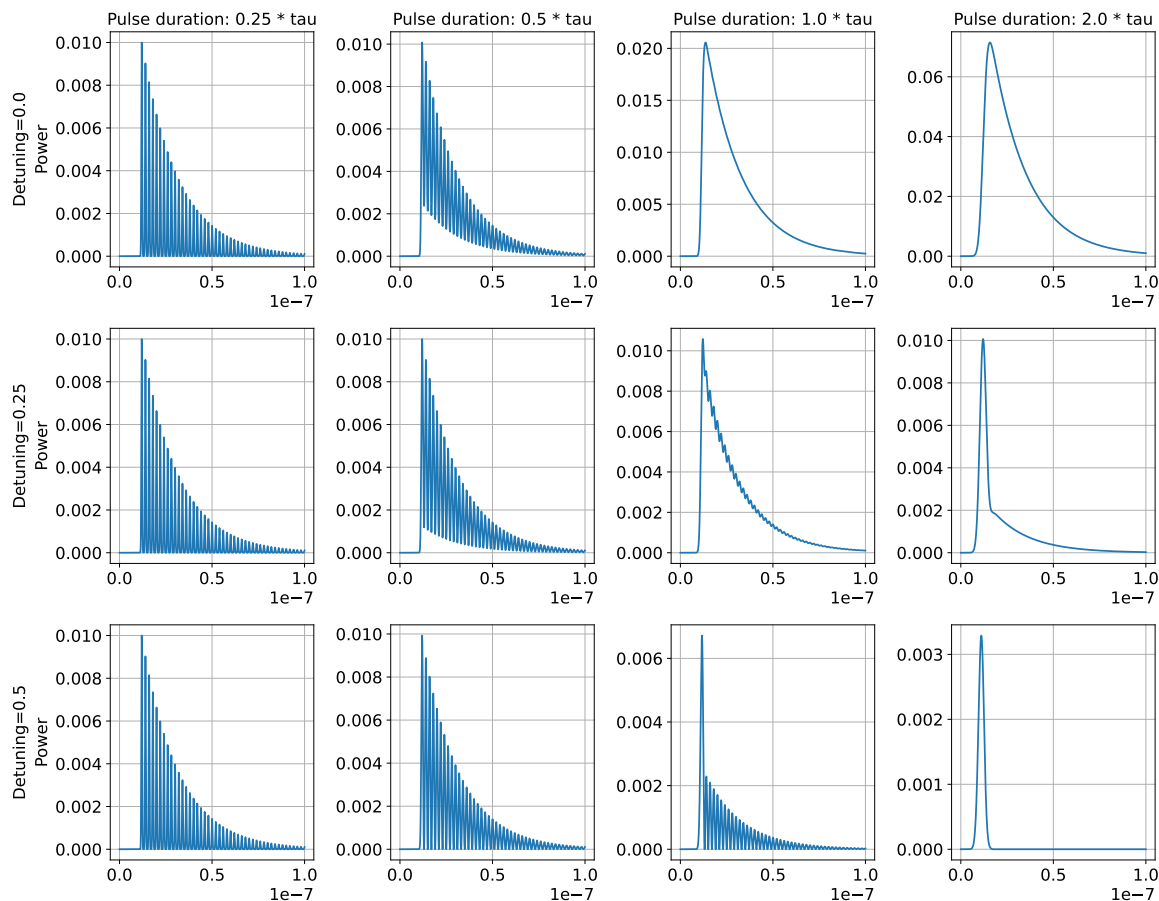


Figure A2. Reproduction of Figure 2 from [4]: simulated cavity ring-down signals under pulsed laser excitation for different ratios of pulse duration to cavity round-trip time. The simulator reproduces all reported regimes: well-separated replicas, partial temporal overlap with interference modulation, and fully overlapped excitation leading to a smooth exponential decay. In this reproduction, a sampling frequency of 2 THz was used, and 2×10^5 simulation points were generated for each panel.

Our simulator reproduces results from [4] using the same physical parameters and without ad hoc fitting. This agreement validates the time-domain modelling of round-trip summation, phase accumulation, and cavity decay, and demonstrates that the published result emerges as a limiting case of the more general simulation framework.

Table A1. Table with benchmark results on the full set of test cavities.

N	Count	Speedup factor (mean \pm std)	Memory ratio (mean)	Pure time (μ s)	Numba time (μ s)	Pure mem (kB)	Numba mem (kB)
1	140	4.67 \pm 0.07	6.48	26.18	5.61	6.98	1.08
2	26	5.09 \pm 0.26	6.60	29.03	5.74	7.27	1.10
3	32	5.29 \pm 0.11	6.49	29.86	5.65	7.30	1.12
4	6	5.38 \pm 0.07	6.37	30.64	5.70	7.32	1.15
5	14	5.60 \pm 0.35	6.27	32.10	5.74	7.34	1.17
6	20	5.64 \pm 0.07	6.18	32.45	5.76	7.39	1.20
7	11	5.74 \pm 0.06	6.06	33.15	5.78	7.39	1.22
8	3	5.76 \pm 0.09	5.97	34.08	5.91	7.41	1.24
9	2	5.94 \pm 0.03	5.88	34.52	5.82	7.44	1.27
10	1	6.11	5.79	35.98	5.89	7.46	1.29

Table A1. Cont.

N	Count	Speedup factor (mean±std)	Memory ratio (mean)	Pure time (μ s)	Numba time (μ s)	Pure mem (kB)	Numba mem (kB)
11	12	6.16 ± 0.10	5.70	36.38	5.91	7.48	1.31
12	3	6.37 ± 0.03	5.62	37.55	5.89	7.51	1.34
13	7	6.46 ± 0.15	5.54	38.23	5.92	7.53	1.36
14	6	6.56 ± 0.06	5.46	39.04	5.95	7.55	1.38
16	2	6.79 ± 0.06	5.32	40.89	6.03	7.60	1.43
17	3	6.87 ± 0.10	5.25	41.32	6.01	7.62	1.45
18	6	6.98 ± 0.07	5.18	42.52	6.09	7.65	1.48
19	1	7.06	5.11	43.33	6.14	7.67	1.50
22	1	7.47	4.93	45.95	6.15	7.74	1.57
23	2	7.32 ± 0.01	4.87	46.21	6.31	7.77	1.59
24	1	7.56	4.82	80.24	10.61	7.79	1.62
25	2	7.68 ± 0.11	4.76	48.72	6.34	7.81	1.64
26	1	7.87	4.71	49.27	6.26	7.84	1.66
27	4	7.74 ± 0.07	4.66	49.29	6.37	7.86	1.69
28	9	7.89 ± 0.05	4.61	51.05	6.47	7.88	1.71
30	12	8.09 ± 0.07	4.51	52.44	6.48	7.93	1.76
31	8	8.17 ± 0.11	4.46	53.25	6.52	7.95	1.78
38	1	8.77	4.17	58.81	6.71	8.12	1.95
46	1	9.59	3.89	65.85	6.87	8.30	2.13
47	1	9.55	3.86	65.93	6.90	8.33	2.16
50	1	9.76	3.77	69.42	7.11	8.40	2.23
53	1	10.05	3.69	71.55	7.12	8.47	2.30
55	5	10.21 ± 0.09	3.63	74.01	7.25	8.52	2.34
57	1	10.34	3.58	74.04	7.16	8.56	2.39
60	12	10.35 ± 0.11	3.51	77.21	7.46	8.63	2.46
62	2	10.77 ± 0.02	3.46	79.23	7.36	8.68	2.51
63	5	10.73 ± 0.14	3.44	79.96	7.45	8.70	2.53
78	1	11.73	3.14	93.75	7.99	9.05	2.88
96	1	12.57	2.88	111.52	8.87	9.52	3.30
102	1	13.09	2.80	115.24	8.80	9.66	3.45
104	1	12.94	2.78	116.89	9.04	9.71	3.49
108	1	13.36	2.73	120.86	9.05	9.80	3.59
112	1	13.68	2.69	125.01	9.14	9.90	3.68
118	1	13.72	2.63	128.56	9.37	10.04	3.82
132	1	14.27	2.50	139.91	9.81	10.37	4.15
136	1	14.50	2.47	145.02	10.00	10.46	4.24
137	1	14.41	2.45	140.53	9.75	10.44	4.27
139	1	14.81	2.43	144.25	9.74	10.48	4.31
140	4	14.78 ± 0.09	2.42	144.96	9.81	10.51	4.34
166	1	15.62	2.26	169.71	10.86	11.16	4.95
176	1	15.77	2.20	176.53	11.19	11.40	5.18
212	1	16.77	2.03	206.86	12.34	12.24	6.02
256	1	17.37	1.88	248.35	14.29	13.27	7.05
276	4	17.56 ± 0.09	1.82	262.30	14.93	13.70	7.52
300	8	17.61 ± 0.10	1.76	282.15	16.02	14.26	8.09

References

1. Lawrence, M.J.; Willke, B.; Husman, M.E.; Gustafson, E.K.; Byer, R.L. Dynamic response of a Fabry–Perot interferometer. *Journal of the Optical Society of America B* **1999**, *16*, 523. <https://doi.org/10.1364/josab.16.000523>.
2. Bhawal, B. Real-time simulation of interferometric gravitational wave detectors involving moving mirrors. *Journal of the Optical Society of America A* **1998**, *15*, 120. <https://doi.org/10.1364/josaa.15.000120>.
3. Beausoleil, R.G.; Sigg, D. Spatiotemporal model of the LIGO interferometer. *Journal of the Optical Society of America A* **1999**, *16*, 2990. <https://doi.org/10.1364/josaa.16.002990>.

4. Yong Lee, J.; Won Hahn, H.W.L. Time Domain Study on Cavity Ring-Down Signals from a Fabry-Pérot Cavity under Pulsed Laser Excitations. *Japanese Journal of Applied Physics* **1999**, *38*, 6287. <https://doi.org/10.1143/jjap.38.6287>.
5. Bhawal, B. LIGO end-to-end simulation program. In Proceedings of the AIP Conference Proceedings. AIP, 2000, Vol. 523, p. 469–470. <https://doi.org/10.1063/1.1291916>.
6. Caron, B. SIESTA, a time domain, general purpose simulation program for the VIRGO experiment. *Astroparticle Physics* **1999**, *10*, 369–386. [https://doi.org/10.1016/s0927-6505\(98\)00059-0](https://doi.org/10.1016/s0927-6505(98)00059-0).
7. Rakhmanov, M. Dynamics of Laser Interferometric Gravitational Wave Detectors. PhD thesis, California Institute of Technology, Pasadena, California, 2000. <https://resolver.caltech.edu/CaltechETD:etd-05242004-154845>.
8. Svizzeretto, A. Optical cavity locking in non linear regime: A Reinforcement Learning strategy for future applications in Gravitational Wave Detectors. Master's thesis, Università di Perugia, Perugia, 2024.
9. Svizzeretto, A.; Bawaj, M. Partial Observability and Domain Randomization in RL-Based Strategy for Optical Cavity Locking Optimization. In Proceedings of the Proceedings of Fifth MODE Workshop on Differentiable Programming for Experiment Design — PoS(MODE2025). Sissa Medialab, 2026, MODE2025, p. 035. <https://doi.org/10.22323/1.491.0035>.
10. Bawaj, M.; Svizzeretto, A. Applying reinforcement learning to optical cavity locking tasks: considerations on actor-critic architectures and real-time hardware implementation, 2026, [[arXiv:physics.ins-det/2509.14884](https://arxiv.org/abs/physics.ins-det/2509.14884)].
11. Hassen, S.Z.S.; Petersen, I.R. Frequency Locking of an Optical Cavity Using a Time-Varying Kalman Filtering Approach. *IEEE Transactions on Control Systems Technology* **2014**, *22*, 1143–1150. <https://doi.org/10.1109/tcst.2013.2266693>.
12. Zhao, W.; Queraltà, J.P.; Westerlund, T. Sim-to-Real Transfer in Deep Reinforcement Learning for Robotics: a Survey. In Proceedings of the 2020 IEEE Symposium Series on Computational Intelligence (SSCI), 2020, pp. 737–744. <https://doi.org/10.1109/SSCI47803.2020.9308468>.
13. Black, E.D. An introduction to Pound–Drever–Hall laser frequency stabilization. *American Journal of Physics* **2001**, *69*, 79–87. <https://doi.org/10.1119/1.1286663>.
14. Barsotti, L. The control of the Virgo interferometer for gravitational wave detection. PhD thesis, Università degli Studi di Pisa, Pisa, 2006. https://gwic.ligo.org/assets/docs/theses/Barsotti_Thesis.pdf.
15. Swinkels, B.; Ruggi, P. Lock acquisition of the Advanced Virgo arm-cavities with reduced force, 2012. <https://tds.virgo-gw.eu/ql/?c=8821>.
16. Acernese, F.; Alshourbagy, M.; Amico, P.; Aoudia, S.; Antonucci, F. Measurement of the optical parameters of the Virgo interferometer. *Applied Optics* **2007**, *46*, 3466. <https://doi.org/10.1364/ao.46.003466>.
17. Acernese, F.; Agathos, M.; Ain, A.; Albanesi, S.; Allocca, A.; Amato, A.; Andrade, T.; Andres, N.; Andrés-Caracsona, M.; Andrić, T.; et al. Virgo detector characterization and data quality: results from the O3 run. *Classical and Quantum Gravity* **2023**, *40*, 185006. <https://doi.org/10.1088/1361-6382/acd92d>.
18. Capocasa, E. Optical and noise studies for Advanced Virgo and filter cavities for quantum noise reduction in gravitational-wave interferometric detectors. PhD thesis, Université Paris Diderot, Paris. https://theses.hal.science/tel-03141038/file/capocasa_thesis_2017.pdf.
19. Little, B.; Chu, S.; Haus, H.; Foresi, J.; Laine, J.P. Microring resonator channel dropping filters. *Journal of Lightwave Technology* **1997**, *15*, 998–1005. <https://doi.org/10.1109/50.588673>.
20. Dorigo, T.; Giammanco, A.; Vischia, P.; Aehle, M.; Bawaj, M.; Boldyrev, A.; de Castro Manzano, P.; Derkach, D.; Donini, J.; Edelen, A.; et al. Toward the end-to-end optimization of particle physics instruments with differentiable programming. *Reviews in Physics* **2023**, *10*, 100085. <https://doi.org/10.1016/j.revip.2023.100085>.
21. Punturo, M.; Abernathy, M.; Acernese, F.; Allen, B.; Andersson, N.; Arun, K.; Barone, F.; Barr, B.; Barsuglia, M.; Beker, M.; et al. The third generation of gravitational wave observatories and their science reach. *Classical and Quantum Gravity* **2010**, *27*, 084007. <https://doi.org/10.1088/0264-9381/27/8/084007>.
22. Hild, S.; Abernathy, M.; Acernese, F.; Amaro-Seoane, P.; Andersson, N.; Arun, K.; Barone, F.; Barr, B.; Barsuglia, M.; Beker, M.; et al. Sensitivity studies for third-generation gravitational wave observatories. *Classical and Quantum Gravity* **2011**, *28*, 094013. <https://doi.org/10.1088/0264-9381/28/9/094013>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.