Editors' Suggestion

Faster-than-adiabatic cooling of non-neutral plasmas

Jan Bartsch [*] and Alfio Borzì [†]
Institut für Mathematik, Universität Würzburg, Emil-Fischer-Strasse 30, 97074 Würzburg, Germany

Paul-Antoine Hervieux [‡] and Giovanni Manfredi [§]
Université de Strasbourg, CNRS, Institut de Physique et Chimie des Matériaux de Strasbourg, UMR 7504, F-67000 Strasbourg, France

(Received 14 September 2024; revised 23 March 2026; accepted 23 July 2026; published 8 September 2026)

Non-neutral plasmas can be confined for long times in electromagnetic traps. Adiabatic cooling of such plasmas can be achieved by slowly decreasing the effective trap stiffness and letting the plasma occupy a larger volume. Here, we propose an optimal control strategy to cool an antiproton plasma in a time much shorter than that required by an adiabatic process. The optimal cooling protocol is found by minimizing an appropriate cost functional that measures the distance of the system from the desired "cold" state, while keeping the stiffness amplitude within prescribed limits. The proposed optimal protocol is shown to outperform other expansion-based cooling methods in terms of stability of the plasma, duration of the process, and volume occupied by the final plasma.

DOI: 10.1103/yn72-3qrw

I. INTRODUCTION

The study of non-neutral plasmas has gained significant attention in recent decades [1,2], due to the fact that highly precise benchtop experiments can be conducted to investigate their intrinsic properties. Non-neutral plasmas can be confined for considerably longer periods of time than quasineutral plasmas, and they can be cooled to lower temperatures with greater ease. These properties make them ideal candidates for studies in fundamental physics, ranging from turbulence and nonlinear dynamics to antimatter physics [3]. The confinement of non-neutral plasmas is typically achieved through the use of cylindrical traps, which employ an axial magnetic field to ensure radial confinement, while the axial confinement is provided by a set of electrodes [4].

The creation and confinement of non-neutral plasmas consisting of antimatter particles (e.g., antiprotons) represents a significant challenge for the production of substantial quantities of neutral antihydrogen atoms $\bar{\mathrm{H}}$. Significant progress has been made in this field through the efforts of several international collaborations [5–7]. The objective of this recent and promising research is to compare the physical properties of matter and antimatter, e.g., the spectral lines of antihydrogen, or its response to gravitational fields [5,8,9].

For these purposes, it is first necessary to cool down the positrons and antiprotons to sufficiently low energies, so that they can combine to form neutral $\bar{\mathrm{H}}$ atoms. A variety of techniques may be employed to implement antiproton and positron cooling. These include the use of Coulomb collisions with cold electrons [10], evaporative cooling [11], cyclotron radiation [12], and adiabatic cooling [6,13]. Each technique has its own advantages and drawbacks. For instance, during evaporative cooling, the majority of charged particles are lost, resulting in a lower density. Conversely, using adiabatic cooling techniques, Gabrielse *et al.* [6] were able to cool down 3×10^6 antiprotons to approximately 3 K with minimal losses.

Adiabatic cooling is achieved by slowly decreasing the trap frequency $\omega(t)$ of a non-neutral plasma in an approximately one-dimensional (1D) harmonic well, $U_{\mathrm{trap}}(x,t) = m\omega^2(t)x^2/2$ [14]. Adiabaticity requires that the frequency varies little during one oscillation, i.e., $|\dot{\omega}(t)| \ll \omega^2(t)$. If the process is not slow enough, the plasma displays unwanted oscillations around the cold target equilibrium. However, achieving fast cooling is often a crucial experimental requirement, for instance when the charged particles are accumulated by bunches, and each bunch needs to be cooled before the arrival of the next [15].

The purpose of this work is to propose an alternative technique, based on optimal control theory, which brings the plasma from a high-temperature to a low-temperature state in a time much shorter than that required by the adiabaticity condition, while keeping the spurious oscillations to a very low level. Similar control techniques have recently been exploited in atomic, molecular, and solid-state physics, where they are known as "shortcuts to adiabaticity" [16]. Optimal control theory offers a potent mathematical framework to optimize the control parameters (in our case, the trap frequency) in order to achieve a desired outcome (the cooling) [17,18]. The optimal control problem is formulated by defining a cost function that describes the desired outcome and the cost attached to the control [19]. By introducing Lagrange multipliers in the form of adjoint variables, the cost function and system constraints are combined into a single Lagrangian, leading to the

[*]Contact author: jan.bartsch@uni-wuerzburg.de
[†]Contact author: alfio.borzi@uni-wuerzburg.de
[‡]Contact author: paul-antoine.hervieux@ipcms.unistra.fr
[§]Contact author: giovanni.manfredi@ipcms.unistra.fr

 ©2026 American Physical Society

Euler-Lagrange equations that provide the required conditions for optimal control [20,21].

II. PLASMA KINETIC MODEL

We consider a system of charged particles confined in a 1D harmonic potential well, where x represents the axial (longitudinal) direction of the cylindrical trap. Their dynamics is governed by the Vlasov-Poisson equations

$$\frac{\partial f}{\partial t} + v\frac{\partial f}{\partial x} + \left(-\omega^2(t)x + \frac{qE}{m}\right)\frac{\partial f}{\partial v} = \left(\frac{\partial f}{\partial t}\right)_{\text{coll}}, \quad (1)$$

$$\frac{\partial E}{\partial x} = \frac{q\sigma_\perp}{\varepsilon_0}\int_{-\infty}^{\infty} f(x,v,t)\,dv, \quad (2)$$

where $E(x,t)$ is the longitudinal self-consistent electric field, $\sigma_\perp$ is the two-dimensional (2D) particle density in the transverse plane, and $\rho(x,t) = \int f(x,v,t)\,dv$ is the 1D density along x. Here, we focus on proton or antiproton plasmas, with electric charge $q = \pm e$ and mass $m = m_P$. We also define the stiffness of the harmonic trap as $\kappa(t) = m\omega^2(t)$. The transverse trap frequency $\omega_\perp/2\pi = eB_x/(2\pi m) \approx 100$ MHz (generated by an axial magnetic field $B_x \approx 3$ T [6]) is much larger than the longitudinal trap frequency $\omega/2\pi \approx 200$ kHz, which is generated via a series of biased electrodes.

In antiproton cooling experiments, the main effect of collisions is to drive the system toward energy equipartition between the parallel and perpendicular motion [13]. Therefore, we added a Bhatnagar-Gross-Krook (BGK) relaxation term [22] on the right-hand side of Eq. (1):

$$\left(\frac{\partial f}{\partial t}\right)_{\text{coll}} = -\nu_{\text{coll}}\left[f(x,v,t) - \frac{\rho(x,t)\,e^{-m(v-u)^2/2k_BT_\perp}}{(2\pi k_BT_\perp/m)^{1/2}}\right], \quad (3)$$

where $u(x,t) = \int f v dv/\rho$ is the mean velocity and $T_\perp(t)$ is the perpendicular temperature, which obeys the relaxation equation

$$\frac{dT_\perp}{dt} = -\frac{\nu_{\text{coll}}}{2}(T_\perp - T_\parallel), \quad (4)$$

where ν_{coll} is the collision frequency and $k_BT_\parallel(t) = m\int f(0,v,t)v^2dv/\rho(0,t)$ is the parallel temperature at the center of the plasma, taken as a reference value. Note that the parallel temperature profile remains almost uniform along the evolution, as shown in Appendix D 2. A factor of 1/2 appears in Eq. (4) because the perpendicular motion is associated with two degrees of freedom. The total temperature is defined as $T = \frac{1}{3}T_\parallel + \frac{2}{3}T_\perp$. The above BGK term induces the relaxation of the parallel temperature toward the perpendicular one—which acts as an effective external bath—on a timescale of the order of ν_{coll}^{-1}. Over long times, all temperatures converge to the same value $T_\perp = T_\parallel = T$, signaling thermal equilibrium.

Hence, although the expansion is 1D in the longitudinal coordinate, the collisions are intrinsically three dimensional (3D). Space-charge effects in the transverse plane are neglected, as they are largely dominated by longitudinal effects in the direction of the expansion. We emphasize that this model has been shown to quantitatively reproduce the results of adiabatic cooling experiments on antiprotons [13].

If the trap frequency $\omega(t)$ decreases in time, the plasma will undergo an expansion, leading to some cooling effect. In the experiments, the trap stiffness is usually varied as a linear ramp. However, using this procedure, the plasma will experience residual oscillations after the frequency has leveled off at its final value. Only if the ramp is infinitely long in time will the oscillations disappear: this is the case of adiabatic cooling, which has been shown to work in practice, but requires times that are much longer than the inverse of the trap frequency [6].

The purpose of this work is to bring the plasma from an initial thermal equilibrium with temperature T_{init} to a final equilibrium with temperature T_{fin} by decreasing the trap stiffness from κ_{init} to κ_{fin}. The optimization procedure will ensure that this is done in a finite time, much shorter than the time required for adiabatic cooling.

By using normalized variables, it is easy to show that the plasma evolution is characterized by only three dimensionless parameters (see Appendix A): (i) the ratio of the final to the initial temperature $T_{\text{fin}}/T_{\text{init}}$, (ii) the ratio of the final to the initial stiffness $\kappa_{\text{fin}}/\kappa_{\text{init}}$, and (iii) the parameter

$$\eta = \frac{e^2\sigma_\perp N_p}{\varepsilon_0\sqrt{mk_BT_{\text{init}}}\,\omega_{\text{init}}} \propto \frac{\omega_p}{\omega_{\text{init}}}, \quad (5)$$

where $N_p = \int \rho\,dx$ is the number of 1D particles (indeed, charged sheets of perpendicular density $\sigma_\perp$), and ω_p is the plasma frequency. When $\eta \ll 1$, the trap confinement dominates, whereas when $\eta \gg 1$ the Coulomb repulsion dominates. Note that, in the former case (dilute plasma), the equilibrium density is almost Gaussian, whereas for large values of η the density is flat inside the trap. For many experimental situations, the Coulomb repulsion and trap confinement are of the same order of magnitude, so that $\eta \gtrsim 1$.

III. OPTIMAL CONTROL AND SIMULATION METHODS

The optimal control method is formulated by defining a suitable cost function J that represents the desired target of the procedure (here, the low-temperature equilibrium state) and the cost incurred in reaching such a target state:

$$J(f,\kappa) = \frac{1}{2}\int_\Omega |f(z,t_{\text{p}}) - f_{\text{targ}}(z)|^2dz + \frac{\alpha_1}{2}\int_0^{t_{\text{p}}} |\kappa(t) - \kappa_{\text{fin}}|^2dt + \frac{\alpha_2}{2}\int_0^{t_{\text{p}}}\left|\frac{d\kappa(t)}{dt}\right|^2dt, \quad (6)$$

where $z = (x,v)$, Ω is the integration domain in the phase space, t_{p} is the duration of the protocol, and $f_{\text{targ}}(z)$ is the desired target state. By minimizing J, the first term in Eq. (6) enforces the final distribution function to be close to the target, the second term prevents the control $\kappa(t)$ from deviating too strongly from its desired final value, and the third term penalizes large variations in the control. The parameters $\alpha_{1,2} > 0$ act as weights for the control terms: the larger they are, the more stringent the two constraints. Here, we take $\alpha_1 = 10^{-4}$, which is a compromise between computational speed and accuracy in reaching the target state, and $\alpha_2 = 0.1$, a choice that prevents large gradients in the control. The impact of these parameters is further explored in Appendix D 3.

Next, we define the Lagrange functional

$$\mathcal{L}(f,\kappa,q) = J(f,\kappa) + \int_0^{t_p}\int_\Omega (\partial_t f(z,t) + L_\kappa f(z,t))q(z,t)dzdt, \quad (7)$$

where $q(z,t)$ is a Lagrange multiplier and $L_\kappa = v\partial_x - \frac{1}{m}[\kappa(t)x \pm eE]\partial_v$ is the phase-space advection operator. Taking the gradients of $\mathcal{L}(f,\kappa,q)$ with respect to its arguments generates the three fundamental equations that govern the optimization problem. Specifically, $\nabla_q\mathcal{L} = 0$ yields back the Vlasov-Poisson system (1) and (2), $\nabla_f\mathcal{L} = 0$ yields the adjoint equation for $q(z,t)$, and $\nabla_\kappa\mathcal{L} = 0$ yields the gradient (or Euler-Lagrange) equation for the control parameter $\kappa(t)$.

To find the optimal solution, we use an iterative gradient-based optimization method [21,23]. Starting from an initial guess $\kappa_0(t)$ of the control, we iteratively execute the following steps: (i) Solve the Vlasov-Poisson system (1) and (2) to obtain $f(z,t)$; (ii) using this f, solve the adjoint equation for $q(z,t)$ backward in time; (iii) given f and its adjoint q, assemble the gradient equation and update the control using a backtracking scheme in the direction of the negative gradient. This loop is stopped when a termination criterion is reached, for example, if the norm of the gradient is smaller than some given threshold. A detailed description of the optimization strategy is given in Appendix B.

The Vlasov-Poisson equations (1) and (2) are solved with a 1D particle-in-cell (PIC) method adapted to the adjoint equation. The time-stepping is performed using a second-order velocity-Verlet scheme [24]. The collisional term is implemented at each time step by sampling the velocities of $N_{\rm coll} = 2\nu_{\rm coll}\Delta t N_p$ particles from the equilibrium distribution appearing in Eq. (3) (see Appendix C). In the forthcoming simulations, we use $N_p = 10^4$, time step $\omega_{\rm init}\Delta t = 0.25$, and phase-space resolution $\Delta x = 0.5\,\ell$, $\Delta v = 0.5\,v_{\rm th}$, where $v_{\rm th} = \sqrt{k_B T_{\rm init}/m}$ and $\ell = v_{\rm th}/\omega_{\rm init}$.

IV. NUMERICAL RESULTS

We simulate an (anti)proton plasma with the following initial parameters: 3D density $\rho_{\rm 3D} = 4.7\times10^{12}$ m^{-3}, initial temperature $T_{\rm init} = 2.7$ meV $= 31$ K, initial trap frequency $\omega_{\rm init}/2\pi = 200$ kHz. These are standard values for trapped nonneutral plasmas [6,25]. The final target state is characterized by $T_{\rm fin} = 0.1T_{\rm init}$ and $\kappa_{\rm fin} = 0.01\,\kappa_{\rm init}$. With these values, Eq. (5) yields $\eta \approx 2.27$, meaning that the Coulomb repulsion and the trap confinement are of similar strength. The initial and final states are computed by solving self-consistently the stationary Vlasov-Poisson equations (1) and (2) and assuming Maxwell-Boltzmann distributions with temperatures $T_\parallel(0) = T_\perp(0) = T_{\rm init}$ and $T_\parallel(t_p) = T_\perp(t_p) = T_{\rm fin}$.

The unmagnetized collision frequency is roughly $\nu_{\rm coll} = \sqrt{2k_BT/m}\,\rho_{\rm 3D}\,b^2$, where $b = e^2/(2\pi\varepsilon_0 k_B T)$ is the distance of closest approach [13]. This yields, in our case, an initial collision time $\nu_{\rm coll}^{-1} = 0.26$ ms, corresponding to $\nu_{\rm coll}/\omega_{\rm init} = 0.003$. However, as the temperature drops by a factor of approximately 5 (see Fig. 1), the collision frequency increases by a factor of approximatley 10 (because $\nu_{\rm coll} \sim T^{-3/2}$). On the other hand, the confining magnetic field has the opposite effect of *reducing* the collision rate by a factor of approximately 5 for our range of parameters [13,26], thus compensating the increase due to the cooling. On the whole, we use the following effective collision rate, which is constant over the entire process and incorporates both competing effects: $\nu_{\rm coll}/\omega_{\rm init} = 0.006$.

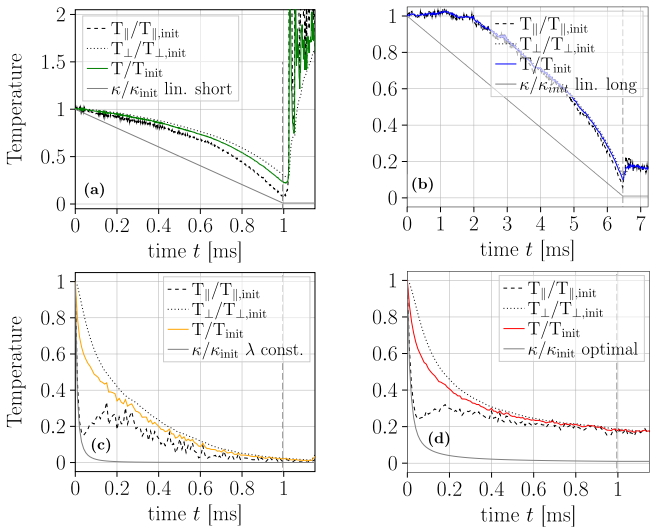


FIG. 1. Evolution of the temperatures at the center of the plasma ($x = 0$), normalized to their initial values, for four protocols: (a) short linear ramp, (b) long linear ramp, (c) constant adiabaticity parameter, and (d) optimal. Note that the long ramp case is about 6.5 times longer than the others. Parallel temperatures are shown as dashed lines, perpendicular temperatures as dotted lines, and total temperatures as solid color lines. The vertical dashed lines indicate the protocol duration t_p.

In the simulations, t_p denotes the total duration of the protocol. After $t = t_p$, the code is run for a short time with constant trap stiffness $\kappa(t_p)$ in order to check the stability of the final state. We compare the cooling efficiency of four different protocols [see Figs. 1(a)–1(d), respectively]: a "short" linear ramp of duration $t_p = 1$ ms, a "long" ramp with $t_p = 6.4$ ms, a protocol with constant adiabatic parameter λ (see below), and our optimal protocol.

The adiabaticity parameter is defined as $\lambda(t) = |\dot\omega(t)|/\omega^2(t)$, with $\lambda(t) \ll 1$ for an adiabatic process (the dot denotes time differentiation). It can be shown that, if $\omega(t)$ decreases faster than $1/t$, then the adiabaticity condition cannot be satisfied for all times [13]. The case $\omega(t) = \omega_{\rm init}/(1+\lambda\,\omega_{\rm init}t)$, which yields a constant λ, thus constitutes the fastest possible decrease of the trap frequency compatible with the adiabaticity condition. This is the case represented in Fig. 1(c) for $\lambda = 0.0525$, a value that was chosen to ensure the corresponding stiffness $\kappa(t)$ is as close as possible to the optimal protocol.

Figure 1 shows the evolution of the plasma temperatures—parallel, perpendicular, and total—for the 4 protocols (the spatial temperature profiles are shown in Appendix D 2). We argue that our proposed protocol [Fig. 1(d)] possesses "optimal" physical properties compared with the other protocols considered in this work. The linear ramp of Fig. 1(a) has the same duration as the optimal case, but clearly does not yield a stable plasma, because large temperature fluctuations

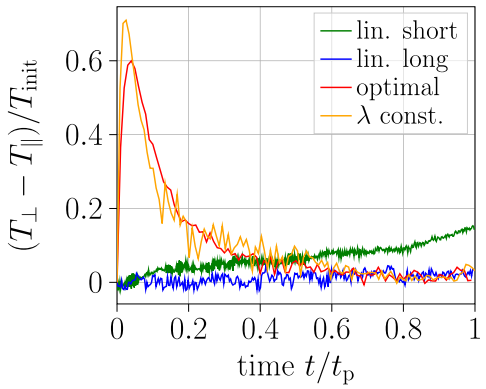


FIG. 2. Evolution of the temperature difference $[T_\perp(t) - T_\parallel(t)]/T_{\rm init}$ for the four protocols shown in Fig. 1 (same color codes), as a function of time normalized to the protocol duration $t_{\rm p}$. The temperature difference is a measure of thermal equilibration and adiabaticity.

are observed after the end of the cooling process. The long linear ramp of Fig. 1(b) is stable and leads to a final temperature similar to the one reached by the optimal protocol, but requires more than a sixfold increase in its duration. The case $\lambda = \text{const}$ [Fig. 1(c)] is a special one. Being adiabatic for all times, it can lead to arbitrarily low temperatures, but at a heavy cost: at the end of the protocol the plasma occupies a much larger region of space. With the present parameters, the final extension of the plasma is approximately 300 cm (see Fig. 3), which is clearly unrealistic for experiments where the trapping device is at most 10–20 cm long. One could try to shorten this protocol, but in that case large fluctuations would appear after $t_{\rm p}$, similar to those observed in Fig. 1(a). More details are given in Appendix D 5.

In summary, the three protocols of Figs. 1(a)–1(c) all fail to produce a final plasma that simultaneously satisfies the

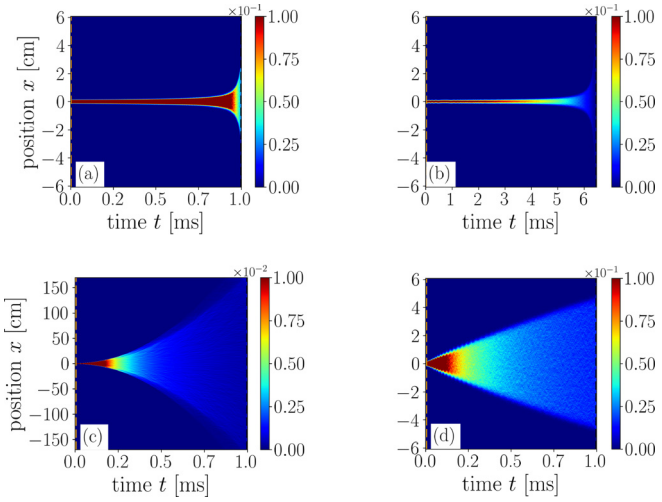


FIG. 3. Space-time evolution of the 1D plasma density $\rho(x,t)$ for the four protocols studied here: (a) short linear ramp, (b) long linear ramp, (c) constant-λ, and (d) optimal. Note the different scale of the vertical axis for panel (c). The color bars represent the values of the plasma density in the normalized units defined in Appendix A.

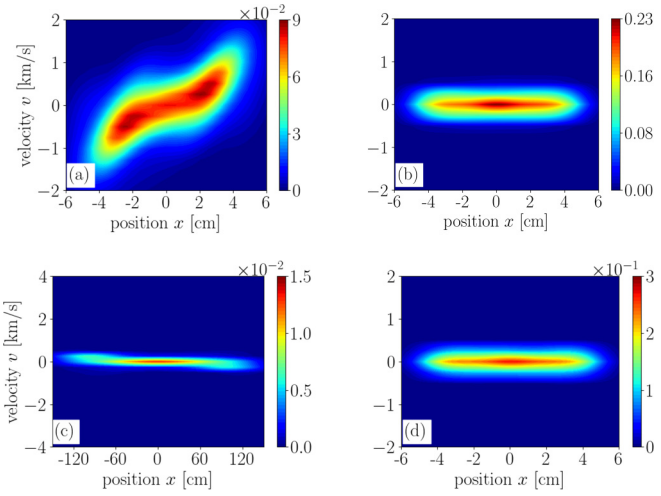


FIG. 4. Phase-space portraits of the distribution functions at the end of the protocol, for the four cases studied here: (a) short linear ramp, (b) long linear ramp, (c) constant-λ, and (d) optimal. Note the different scale of the horizontal (position) axis for panel (c). The color bars represent the values of the distribution functions, in the normalized units defined in Appendix A.

following fundamental requirements: (i) stability of the temperature, (ii) short duration of the protocol, and (iii) moderate volume occupied by the plasma. Only our optimal protocol [Fig. 1(d)] is capable of achieving a stable plasma, with longitudinal length $\lesssim$10 cm, within a time of approximately 1 ms; see Appendix D 1 for a qualitative comparison of the protocols according to the above criteria. The optimal protocol is also robust against random perturbations, as shown in Appendix D 6.

During a perfectly adiabatic (infinitely slow) process, the system is at thermal equilibrium at every instant of time, so that the parallel and perpendicular temperatures are equal. Hence, in Fig. 2, we display the temperature difference as a measure of the adiabaticity of each evolution. As expected, for the slow linear ramp (blue curve) the two temperatures remain close to each other during the whole protocol, signaling quasi-adiabaticity, while for the fast linear ramp (green curve) the temperature difference increases. Interestingly, the curves for the optimal and constant-λ protocols are very similar. In both cases, the plasma is not at thermodynamic equilibrium at intermediate times: indeed, the parallel temperature decreases faster due to the plasma expansion, then grows again under the action of the collisions, finally reaching $T_\perp$, as can be seen in Figs. 1(c) and 1(d). This behavior was already observed in our earlier work [13]. Such nonequilibrium phase appears to be a crucial condition to achieve significant cooling in a short time.

In Fig. 3, we show the spatiotemporal evolution of the 1D plasma density $\rho(x,t)$, evincing the plasma expansion that is at the origin of the cooling. For the linear ramp cases [Figs. 3(a) and 3(b)], the expansion occurs mostly near the end of the process, whereas for the optimal protocol [Fig. 3(d)] the expansion is evenly distributed along its whole duration. The constant-λ case [Fig. 3(c)] is similar to the optimal one, but the plasma expansion is much larger (approximately 300 cm), as anticipated above. The corresponding profiles of the trap and self-consistent potentials are shown in Appendix D 4.

Finally, Fig. 4 shows the phase-space portraits at the end of the protocol, for the same four cases. As expected, the short linear protocol [Fig. 4(a)] is far from equilibrium at $t = t_{\rm p}$, while the constant-λ case [Fig. 4(c)] is very elongated in the x direction, in agreement with the preceding results. The linear long [Fig. 4(b)] and optimal [Fig. 4(d)] protocols appear to produce stable distribution functions, but the optimal protocol does so in a time that is about 6.5 times faster.

Additional numerical results are shown in Appendix D 7 (no collisions, $\nu_{\rm coll} = 0$), Appendix D 8 (no space-charge, $\eta = 0$), and Appendix D 9 (evolution of the emittance).

V. CONCLUSIONS

We have presented an optimal-control method to cool down a non-neutral collisional plasma in a time much shorter than the time required by an adiabatic process. As in the latter, the cooling is obtained by decreasing the trap stiffness $\kappa(t)$ and letting the plasma occupy a larger volume. We compared our optimal protocol to three other standard methods to achieve cooling through expansion: two linear ramp protocols, either short or long (quasi-adiabatic), which is a common procedure used in the experiments, and a protocol with a constant adiabaticity parameter, for which the trap frequency decreases as $1/t$. Our optimal protocol outperformed all others in terms of stability, duration of the process, and volume occupied by the final plasma. This work suggests that the use of optimal-control techniques may lead to more efficient trapping and cooling techniques for non-neutral plasmas. This is a particularly timely issue for the burgeoning field of antimatter plasmas, which need to be confined and cooled before significant losses and annihilation occur.

DATA AVAILABILITY

There are no publicly available research data or software supporting this paper. Requests for further information or data should be sent to the authors.

APPENDIX A: DIMENSIONLESS PARAMETERS AND NORMALIZATION

The time-evolution of the non-neutral plasma is described by the Vlasov-Poisson system (1) and (2). We transform to dimensionless variables (denoted by an overbar), as follows:

$$\bar{t} = \omega_{\rm init} t, \quad \bar{v} = v/v_{\rm th}, \quad \bar{x} = x/\ell,$$
$$\overline{E} = eE\ell/(k_B T_{\rm init}), \quad \overline{\rho} = \rho\ell/N_p, \quad \overline{f} = f v_{\rm th}\ell/N_p,$$

where $\ell = v_{\rm th}/\omega_{\rm init}$ and $v_{\rm th} = \sqrt{k_B T_{\rm init}/m}$ is the initial thermal velocity. Note that, in these normalized variables, it holds $\iint \overline{f}(\bar{x}, \bar{v}, \bar{t})\, d\bar{x}\, d\bar{v} = \int \overline{\rho}\, d\bar{x} = 1$ (integrals are from $-\infty$ to $+\infty$).

Using this normalization, the Vlasov-Poisson system becomes

$$\frac{\partial \overline{f}}{\partial \bar{t}} + \bar{v}\frac{\partial \overline{f}}{\partial \bar{x}} - (\overline{\kappa}(\bar{t})\bar{x} - \overline{E})\frac{\partial \overline{f}}{\partial \bar{v}} = -\bar{\nu}_{\rm coll}[\overline{f}(\bar{x}, \bar{v}, \bar{t}) - \overline{f}^{\rm eq}(\bar{x}, \bar{v}, \bar{t})], \tag{A1}$$

$$\frac{\partial \overline{E}}{\partial \bar{x}} = \eta\, \overline{\rho}(\bar{x}, \bar{t}), \tag{A2}$$

where $\overline{\kappa}(\bar{t}) \equiv \omega^2(t)/\omega_{\rm init}^2 = \kappa(t)/\kappa_{\rm init}$, $\bar{\nu}_{\rm coll} = \nu_{\rm coll}/\omega_{\rm init}$, and

$$\eta = \frac{e^2\sigma_\perp N_p}{m\varepsilon_0 v_{\rm th}\omega_{\rm init}} = \frac{e^2\sigma_\perp N_p}{\varepsilon_0\sqrt{mk_B T_{\rm init}}\,\omega_{\rm init}}. \tag{A3}$$

To rewrite the dimensionless parameter η in a more transparent way, we define the 3D number density as $\rho_{\rm 3D} = \sigma_\perp N_p/L$, where L is the linear extension of the plasma along the longitudinal direction x. With this setting, we have

$$\eta = \frac{e^2\rho_{\rm 3D}}{m\varepsilon_0}\frac{L}{v_{\rm th}\omega_{\rm init}} = \frac{\omega_p}{\omega_{\rm init}}\frac{L}{\lambda_D}, \tag{A4}$$

where we have defined the plasma frequency and the Debye length as follows:

$$\omega_p = \left(\frac{e^2\rho_{\rm 3D}}{m\varepsilon_0}\right)^{1/2}, \quad \lambda_D = v_{\rm th}/\omega_p.$$

Hence, there are only three parameters to be specified for a simulation: $\overline{\kappa}_{\rm fin} = \kappa_{\rm fin}/\kappa_{\rm init}$, $\overline{T}_{\rm fin} = T_{\rm fin}/T_{\rm init}$, and η. One can distinguish two physically different regimes, depending on whether the external harmonic confinement or the Coulomb repulsion dominates: if $\eta \ll 1$, then the trap confinement dominates, whereas if $\eta \gg 1$ the Coulomb repulsion dominates.

APPENDIX B: OPTIMAL CONTROL THEORY

In this section, we consider the dimensionless Vlasov-Poisson (VP) system (A1) and (A2), but drop the overline for ease of notation. Here, the time-dependent continuous function $\kappa(t)$ (stiffness of the trap) represents the control parameter. We prescribe the values of this control at the beginning and end of the cooling protocol. We have

$$\kappa(0) = \kappa_{\rm init}, \quad \kappa(t_{\rm p}) = \kappa_{\rm term}. \tag{B1}$$

The purpose of the control is to drive the evolution of the phase-space distribution function from an initial "hot" equilibrium state $f_0(x, v)$ to a colder equilibrium denoted by $f_{\rm targ}(x, v)$. We also require that the control acts with a small cost, represented by the following cost function:

$$J(f, \kappa) = \frac{1}{2}\iint_\Omega \left|f(x, v, t_{\rm p}) - f_{\rm targ}(x, v)\right|^2 dx dv + \frac{\alpha_1}{2}\int_0^{t_{\rm p}} |\kappa(t) - \kappa_{\rm fin}|^2 dt + \frac{\alpha_2}{2}\int_0^{t_{\rm p}} \left|\frac{d\kappa(t)}{dt}\right|^2 dt,$$

where $\Omega = (-x_{\max}, x_{\max}) \times (-v_{\max}, v_{\max})$ is the integration domain in the phase space [see Eq. (6) in the main paper]. The minimization of the first integral in J forces the phase-space distribution at the final time to reach the desired target state, while the minimization of the last two integrals requires keeping the cost of the control to a minimum. The weighted sum of these terms results in a functional representing a trade-off between these two tasks. Moreover, the choice of the two weights α_1 and α_2 allows one to tune the relative importance of the two integrals representing the total cost. In particular, the second integral of the cost (proportional to α_1) promotes changes of the control that bring it close to the given final value $\kappa_{\rm fin}$, whereas the third integral (proportional to α_2) guarantees that the control is continuous and limits its gradient.

Such an optimal control problem requires the minimization of $J(f,\kappa)$, where $f(x,v,t)$ is governed by the VP system and $\kappa(t)$ satisfies the constraints (B1). We solve this minimization problem through the Lagrange approach, which requires the introduction of a Lagrange multiplier of the same dimensionality as the constraint. We denote this Lagrange multiplier (also called an adjoint variable) with $q(x,v,t)$. Hence, we can write the following Lagrange functional:

$$\mathcal{L}(f,\kappa,q) = J(f,\kappa) + \int_0^{t_p} \iint_\Omega (\partial_t f(x,v,t) + L_\kappa f(x,v,t)) \times q(x,v,t)dxdvdt, \tag{B2}$$

where $L_\kappa = v\,\partial_x - [\kappa(t)x - E]\partial_v$ is the streaming operator. The next step is to compute the Euler-Lagrange equations, which are obtained by taking the functional derivatives of $\mathcal{L}$ with respect to the arguments (f,κ,q) and requiring that they are zero. The result is the optimality system $\nabla_f\mathcal{L} = 0$, $\nabla_\kappa\mathcal{L} = 0$, $\nabla_q\mathcal{L} = 0$. Note that L_κ depends on f since the electric field E is given by

$$E(x,t) = \eta \int_0^x \int_{-\infty}^{\infty} f(s,v,t)dvds.$$

Taking the derivative of $\mathcal{L}$ with respect to the state variable f, we obtain the adjoint equation for q, as follows:

$$\partial_t q(x,v,t) - L_\kappa^* q(x,v,t) = \nu_{coll}\, q(x,v,t) - V(x,t), \tag{B3}$$

$$\partial_x V(x,t) = \eta\, \mathcal{I}(x,t), \tag{B4}$$

$$q(x,v,t_p) = -f(x,v,t_p) + f_{targ}(x,v), \tag{B5}$$

where we introduced the adjoint streaming operator $L_\kappa^* = -L_\kappa$, and the term

$$\mathcal{I}(x,t) = \int_{-\infty}^{\infty} \partial_v f(x,v,t) q(x,v,t)dv. \tag{B6}$$

Note that the adjoint model evolves backward in time, starting with a terminal condition at t_p that is given by a difference of phase-space distributions, as in Eq. (B5). However, this setting hinders the application of the particle method used in our work, because we need to interpret q as a (positive definite) density. Therefore, in order to avoid working with the difference appearing in the terminal condition (B5), we introduce the following two auxiliary adjoint problems:

$$\begin{aligned} -\partial_t q_1(x,v,t) + L_\kappa^* q_1(x,v,t) &= V(x,t) - \nu_{\rm coll}\, q_1(x,v,t), \\ q_1(x,v,t_p) &= f_{\rm targ}(x,v) \end{aligned} \tag{B7}$$

$$\begin{aligned} -\partial_t q_2(x,v,t) + L_\kappa^* q_2(x,v,t) &= -\nu_{\rm coll}\, q_2(x,v,t), \\ q_2(x,v,t_p) &= f(x,v,t_p). \end{aligned} \tag{B8}$$

If subtracted from each other, the two equations give the original adjoint equation, and therefore $q = q_1 - q_2$ fulfills Eq. (B3) with terminal condition (B5).

Next, we discuss the derivative of the Lagrange functional with respect to κ. By requiring that it is zero, we obtain the equation:

$$\alpha_1(\kappa(t) - \kappa_{\rm fin}) - \alpha_2 \frac{d^2}{dt^2}\kappa(t) - \iint_\Omega x\partial_v f(x,v,t) q(x,v,t)dxdv = 0. \tag{B9}$$

This is a second-order differential equation on the time interval $(0,t_p)$, with the boundary conditions $\kappa(0) = \kappa_{\rm init}$ and $\kappa(t_p) = \kappa_{\rm term}$. This equation must be satisfied at optimality. Otherwise, one can prove that the left-hand side of (B9) (the residual) provides the gradient of the functional J with respect to κ. We denote this residual with $\nabla_\kappa J$.

Clearly, the functional derivative of $\mathcal{L}$ with respect to κ refers to the Hilbert space with the usual scalar product of functions defined on $(0,t_p)$. This space is denoted with $L^2(0,t_p)$, and the gradient $\nabla_\kappa J$ is defined on this space. On the other hand, our control function is differentiable, i.e., it belongs to a more regular space.

From the optimization point of view, this means that we cannot use $\nabla_\kappa J$ to update the control function in a gradient-based iterative procedure, since $\kappa^{\rm new} = \kappa^{\rm old} - \nabla_\kappa J(\kappa^{\rm old})$ would result in a $\kappa^{\rm new}$ with lower regularity. Hence, we need to project our gradient onto the space of differentiable functions. Using Riesz's representation theorem, we obtain this projected gradient Ψ as the solution of the following boundary-value problem [27]:

$$\alpha_1\Psi(t) - \alpha_2\frac{d^2}{dt^2}\Psi(t) = \nabla_\kappa J(t), \tag{B10}$$

$$\Psi(0) = \kappa_{\rm init}, \quad \Psi(t_p) = \kappa_{\rm term}. \tag{B11}$$

The projected gradient Ψ is then used to update the control as follows:

$$\kappa^{\rm new}(t) = \kappa^{\rm old}(t) - s\Psi(t),$$

where $s > 0$ is the gradient step size. This procedure ensures that the updated control $\kappa^{\rm new}$ lives in the correct functional space and possesses the required regularity.

APPENDIX C: NUMERICAL METHODS

To solve the Vlasov-Poisson system (1) and (2), we employ a standard particle-in-cell (PIC) method with N_p simulation particles [28,29]. To integrate Newton's laws of motion in the streaming part L_κ, we apply the Velocity-Verlet scheme [24].

To reconstruct the full distribution function $f(x,v,t)$, which is required for the optimization procedure, we use a finite-volume discretization in the phase space [30]. For this purpose, we consider the computational domains $\Omega_x = [-x_{\max}, x_{\max}]$ for position and $\Omega_v = [-v_{\max}, v_{\max}]$ for velocity, where $x_{\max}$ and $v_{\max}$ are chosen sufficiently large such that the distribution function f can be considered zero at the boundaries along the evolution. We split Ω_x and Ω_v in N_x and N_v volumes, respectively. Furthermore, we split the time interval $[0,t_p]$ into N_t subintervals of length Δt.

To implement the collision term, we use a rejection-free method for the BGK operator [31]. Specifically, we calculate the number of particles $N_{\rm coll}$ that undergo, on average, a collision event within one time step, as $N_{\rm coll} = 2\,\nu_{\rm coll}\,\Delta t\,N_p$. Then, we randomly select $N_{\rm coll}$ particles and re-initialize their

TABLE I. Physical parameters expressed in the normalized units defined in Appendix A.

Symbol	value	Symbol	value
κ_{init}	1	T_{init}	1
κ_{term}	0.015	T_{fin}	0.1
$v_{\max}$	10	$x_{\max}$	150
η	2.27	ν_{coll}	0.006

velocities randomly according to the normal distribution,

$$f^{\text{eq}}(x, v, t) = \frac{\rho(x,t)}{\sqrt{2\pi k_B T_\perp/m}} \exp\left(\frac{-m(v-u)^2}{2k_B T_\perp}\right),$$

where $u(x,t) = \int f(x,v,t)\, v\, dv/\rho(x,t)$ is the mean velocity in the x direction. This procedure ensures that (in the absence of any other terms) the temperatures equilibrate exponentially with a rate equal to ν_{coll}.

To solve the adjoint equation in the form of Eqs. (B7) and (B8), we apply an operator splitting technique. For each adjoint equation, we split the time-step process into solving the streaming terms and the source and sink terms separately. For the former, we adapt the standard methods used for the Vlasov equation [32]. For the latter, we add or subtract particles as follows: Given a total amount of available particles N_q to be added or subtracted, in the case of Eq. (B7) we add to q_1 a fraction of N_q particles at each point of the phase space only if V is positive; furthermore, we subtract a number of particles proportional to $\nu_{\text{coll}}\, q_1$. Similarly, in the case of Eq. (B8), we subtract from q_2 a number of particles proportional to $\nu_{\text{coll}}\, q_2$.

To solve the system (B10) and (B11) for the gradient, we employ a finite-difference approximation, and the resulting tridiagonal system is solved using the Thomas algorithm [33]. To solve our optimal control problem, we employ a nonlinear conjugated gradient method using the Fletcher-Reeves formula for the step-direction update [23].

We have run our optimization procedure with the data presented in Tables I and II (all parameters are expressed in the normalized units defined in Appendix A), using a linear-in-time control as the initial guess for the optimization procedure.

APPENDIX D: ADDITIONAL NUMERICAL RESULTS

In this section, we present some additional numerical tests and results that complement those provided in the main text.

TABLE II. Numerical parameters.

Symbol	value	Symbol	value
N_t	5×10^3	Δt	0.25
N_x	600	Δx	0.5
N_v	40	Δv	0.5
N_p	10^4		

TABLE III. Comparison of the different protocols in terms of three criteria: plasma length, protocol duration, and temperature fluctuations. Failure to comply with these criteria is highlighted in red.

	Optimal	Lin. short	Lin. long	Const. λ
Length [cm]	10	10	10	300
Duration [ms]	1	1	6	1
Fluctuations	Low	High	Low	Low

1. Qualitative comparison of the protocols

As explained in the main text, an "ideal" cooling protocol should satisfy a number of conditions: (i) short duration, (ii) low-temperature fluctuations at the end of the protocol, and (iii) moderate length of the plasma. These three requirements are summarized in Table III for the four protocols considered in the present work. It clearly appears that only our optimal protocol satisfies all these criteria simultaneously.

2. Temperature profiles

In Fig. 5, we show the longitudinal temperature profile for the optimal case at different times. Halfway through the protocol, the profile becomes uniform across the length of the plasma. We recall that the perpendicular temperature is assumed to be uniform in the model.

3. Influence of the weights in the cost function

The cost function J to be minimized includes two free parameters α_1 and α_2, see Eq. (6) in the main text. These parameters act as weights for the control terms: the larger they are, the more stringent the two constraints. In the main

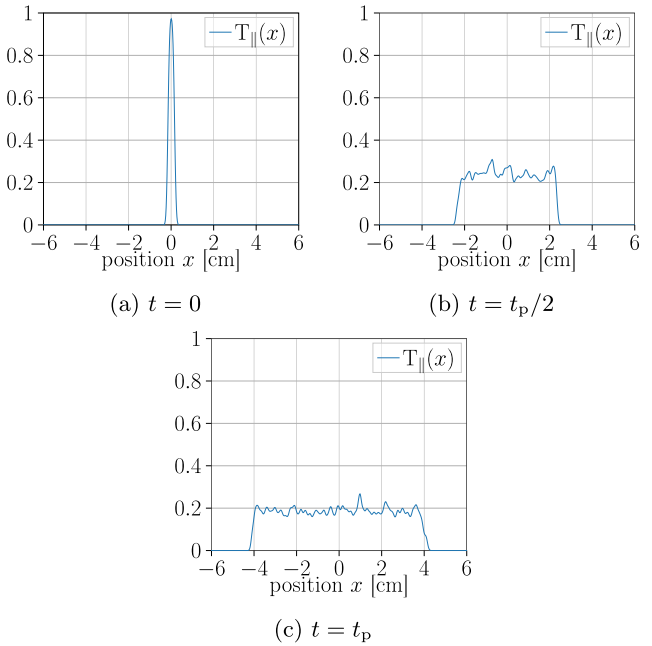
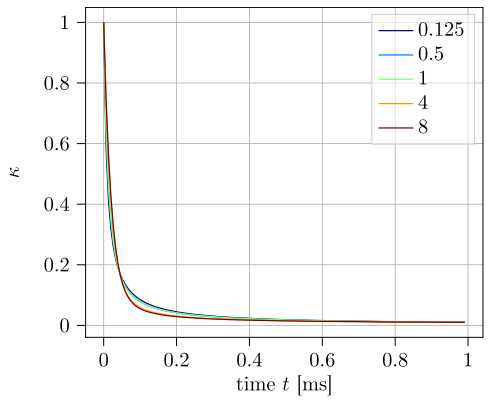
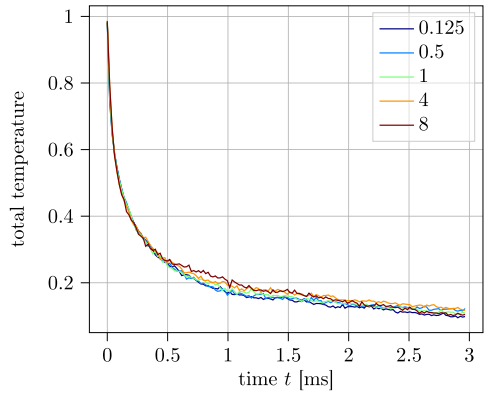


FIG. 5. Longitudinal temperature profiles at different times, for the optimal protocol.



(a) Protocol stiffness



(b) Total temperature

FIG. 6. (a) Protocol stiffness and (b) total temperature, normalized to their initial values, for different weights (α_1, α_2), obtained by multiplying the values used in the main paper, $(\alpha_1, \alpha_2) = (10^{-4}, 0.1)$, by the factors indicated on the figure.

text, we took $\alpha_1 = 10^{-4}$, which is a compromise between computational speed and accuracy in reaching the target state, and $\alpha_2 = 0.1$, a choice that prevents large gradients in the control. Here, we check the stability of the procedure by multiplying both α_1 and α_2 by the same numerical factor in the range [0.125–8]. All other parameters are left unchanged.

The results for the optimal protocols and the total temperature evolutions, shown in Fig. 6, demonstrate that the optimization procedure is rather stable against variations of these parameters.

4. Plasma density and potentials

In Fig. 7 (top panel), we plot the electrostatic, trap, and total potentials at the end of the optimal protocol. We note that the self-consistent electrostatic potential is approximately parabolic, similarly to the trap potential, but with opposite concavity, reflecting the repulsive character of the Coulomb interaction. The parabolic shape is in accordance with the approximately flat plasma density, as seen in Fig. 7 (bottom panel).

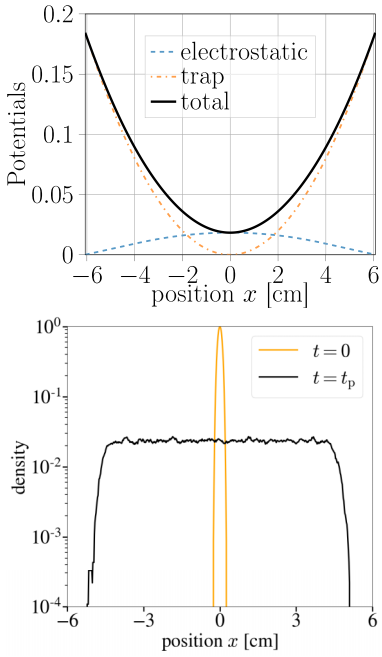


FIG. 7. (top panel) Electrostatic (dashed blue line), trap (dot-dashed orange line), and total (black solid line) potentials at the end of the protocol ($t = t_p$), normalized to $k_B T_{\text{init}}$. (bottom panel) Corresponding plasma density at the initial and final times, normalized so that $\max_x \rho(x, 0) = 1$.

5. Constant-λ protocols

The adiabaticity parameter $\lambda(t)$ is defined as

$$\lambda(t) = \left| \frac{\dot{\omega}(t)}{\omega^2(t)} \right| = \left| \frac{\dot{\kappa}(t)}{2\,\kappa^{\frac{2}{3}}(t)} \right|, \tag{D1}$$

and should be small for an adiabatic protocol. The constant-λ protocol corresponds to the following expression for the trap frequency:

$$\omega(t) = \frac{\omega_{\text{init}}}{1 + \lambda \omega_{\text{init}} t}, \tag{D2}$$

where we take $\lambda = 0.0525$ as in the main text (with this value of λ, this protocol is very close to the optimal one). In Fig. 8, we show the temporal behavior of $\lambda(t)$ for the different cooling protocols considered in the main text.

A crucial property of the constant-λ protocol is that, in contrast with our optimal protocol, the final value of the trap stiffness $\kappa(t_p)$ cannot be fixed *a priori*, but rather depends on the protocol duration t_p, as is clear from Eq. (D2). Hence, the constant-λ protocol cannot control the final plasma extension, which depends directly on the trap stiffness. This is a fundamental limitation, as we show in the next paragraphs.

For a constant-λ protocol duration of $t_p = 1$ ms, similar to that of our optimal protocol, we obtain the final stiffness given in Table IV(a). As discussed in the main text, this protocol yields very low values of the final temperature [see Fig. 9(a) below], but at the prohibitive cost of having a final plasma extension close to 300 cm, as seen in Fig. 10(a), which is experimentally unrealistic.

TABLE IV. Constant-λ protocols with $\lambda = 0.0525$ and different protocol durations t_p. The tables give the final stiffness, temperature, and plasma length.

(a) $t_p = 1$ ms	
$\kappa(t_p)/\kappa_{init}$	0.00035
$T(t_p)/T(0)$	0.035
Plasma length	300 cm
(b) $t_p = 0.2$ ms,	
$\kappa(t_p)/\kappa_{init}$	0.00502
$T(t_p)/T(0)$	0.44
Plasma length	20 cm
(c) $t_p = 0.1$ ms,	
$\kappa(t_p)/\kappa_{init}$	0.0175
$T(t_p)/T(0)$	0.5
Plasma length	3.6 cm

To limit the length of the plasma, one could reduce the protocol duration, taking for instance $t_p = 0.2$ ms or $t_p = 0.1$ ms, as in Tables IV(b) and IV(c). In these cases, the plasma lengths become 20 cm and 7.2 cm, respectively, as shown in Figs. 10(b) and 10(c). However, for these protocols, the final plasma is not stable: its temperature increases significantly after the end of the protocol and displays large oscillations, as seen in Figs. 9(b) and 9(c). The same outcome would have been observed by keeping the same t_p, but reducing the adiabaticity parameter λ.

We conclude that the constant-λ protocols—being unable to control the final length of the plasma—do not constitute a viable alternative to our optimal protocol.

6. Test case with noisy control

To check the robustness of our optimal control strategy, we add some random white noise to the optimal control parameter $\kappa^*(t)$. Specifically, we draw random numbers r_{σ_n} from a normal distribution with zero mean and standard deviation $\sigma_n \in (0, 1)$ and add them to the optimal control to obtain the "noisy" control κ_n:

$$\kappa_n(t) = \kappa^*(t)(1 + r_{\sigma_n}).$$

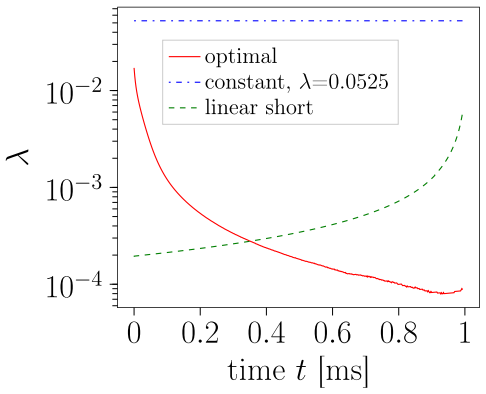
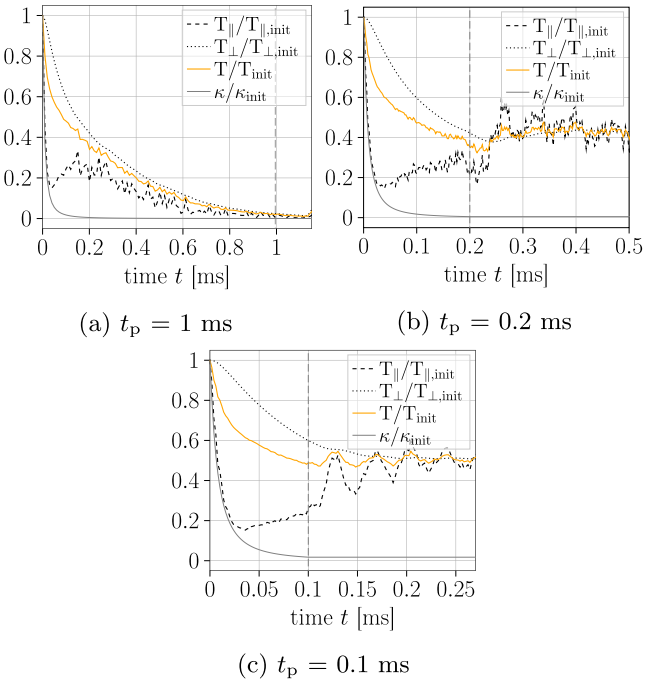


FIG. 8. Evolution of $\lambda(t)$, on a semilogarithmic scale, for the optimal protocol (solid red line), the short linear protocol (dashed green line), and the constant-λ protocol (dash-dotted blue line).

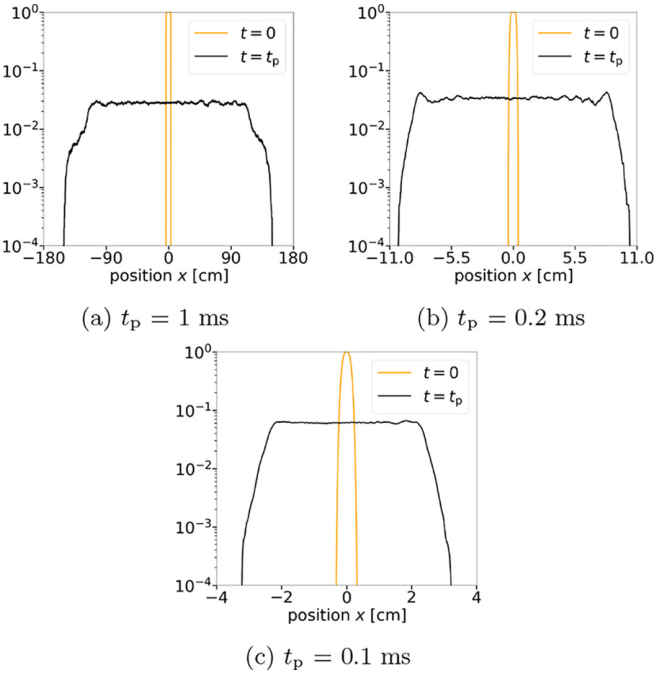


(a) $t_p = 1$ ms (b) $t_p = 0.2$ ms

(c) $t_p = 0.1$ ms

FIG. 9. Evolution of the normalized temperatures and stiffness, for the three protocols of Table IV.

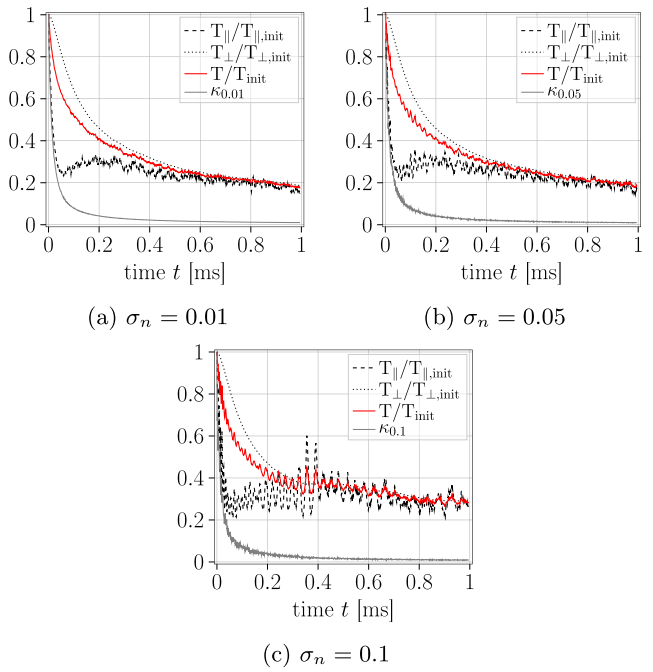
We then consider different levels of noise σ_n (see Fig. 11). For noise levels up to 5%, the control appears to be robust. Moreover, for noise levels below 1%, we see almost no change in the behavior of the temperature compared with the zero-noise case presented in the main paper.



(a) $t_p = 1$ ms (b) $t_p = 0.2$ ms

(c) $t_p = 0.1$ ms

FIG. 10. Plasma densities at $t = 0$ and $t = t_p$, normalized so that $\max_x \rho(x, 0) = 1$, for the three protocols of Table IV.



(a) $\sigma_n = 0.01$ (b) $\sigma_n = 0.05$

(c) $\sigma_n = 0.1$

FIG. 11. Evolution of the temperatures for noisy optimal protocols. The relative amplitude of the noise is (a) $\sigma_n = 0.01$, (b) $\sigma_n = 0.05$, and (c) $\sigma_n = 0.1$.

7. Test case without collisions

We have performed a test with no collisions, i.e., $\nu_{\mathrm{coll}} = 0$ in Eq. (A1). In this case, there is no mechanism coupling the parallel and perpendicular temperatures. The perpendicular temperature remains unchanged, while the parallel one decreases very efficiently, because it is not restrained by the coupling to the perpendicular reservoir. However, such a case is not realistic, because the isotropizing collisions do play a significant role for the type of plasma considered here. The related numerical results for an optimal protocol are plotted in Fig. 12.

8. Test case without self-consistent fields

If we turn off the space charge, i.e., consider the case $\eta = 0$ in Eq. (A2) (all other parameters being unchanged), there is no Coulomb repulsion between the particles. Hence, the distribution function in the phase space remains essentially Gaussian. In this case, the observed temperature decrease, shown in

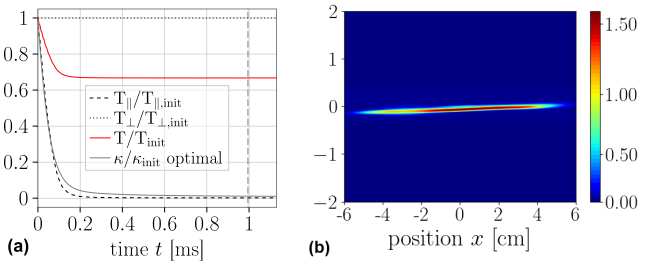


FIG. 12. Test-case with no collisions. (a) Evolution of the temperatures and trap stiffness. (b) Phase space at the end of the protocol.

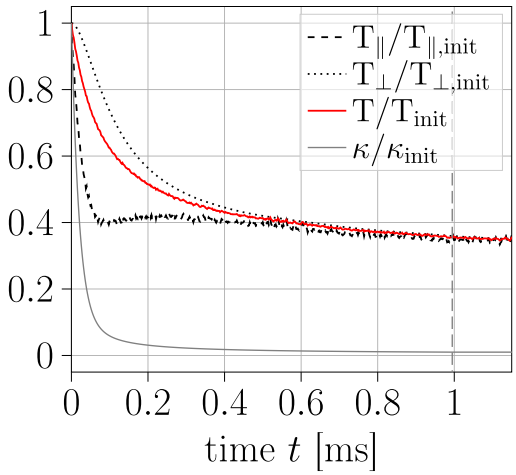
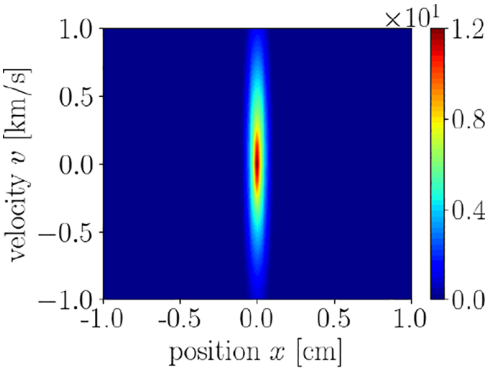
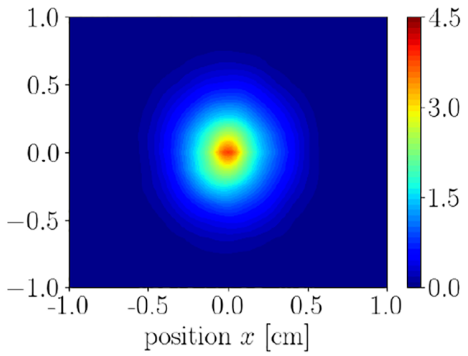


FIG. 13. Test-case without space charge: Evolution of the temperatures and trap stiffness.

Fig. 13, is significantly smaller, with the final temperature reaching only $T_{\mathrm{fin}} \approx 0.35 T_{\mathrm{init}}$. This is probably due to the fact that the self-consistent Coulomb repulsion, which is absent here, favors the plasma expansion and hence the cooling effect. Indeed, the plasma extension is only approximately 1 cm in this case (see Fig. 14), as opposed to approximately 10 cm when the Coulomb repulsion is included.



(a) $t = 0$



(b) $t = t_{\mathrm{p}}$

FIG. 14. Test-case without space charge: (a) initial and (b) final phase-space distributions.

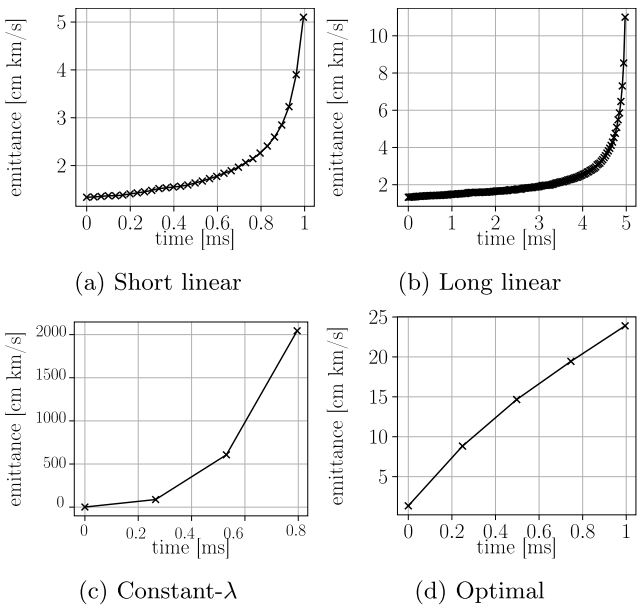


FIG. 15. Emittance for different cooling protocols, as a function of time.

9. Emittance

In the field of charged-particle beams, the emittance is a commonly employed quantity to characterize the area occupied by the beam in the phase space. However, emittance and temperature describe different physical properties despite both reflecting a form of spread in the particle distribution. Temperature quantifies the mean kinetic energy per particle associated with random motion and is well defined only near thermodynamic equilibrium. Emittance, in contrast, measures the phase-space area occupied by a beam and incorporates spatial extent, momentum spread, and their correlations. Emittance characterizes beam quality rather than thermodynamic state, and cannot in general be interpreted as a temperature.

In the present context (cooling of a non-neutral plasma), we are mainly interested in the temperature. However, for completeness, we include in Fig. 15 some results on the emittance. While the temperature decreases for all the protocols considered in the main text, the emittance is found to grow in all cases. The largest increase is observed for those protocols (optimal and constant-λ) that reach the lowest final temperature.

Here, the emittance ε is defined as [see Ref. [34], Eq. (11)]: $\varepsilon = (\sigma_x \sigma_v - \sigma_{xv}^2)^{1/2}$, where

$$\sigma_x = \frac{\sum x^2}{N_p} - \left(\frac{\sum x}{N_p}\right)^2,$$

$$\sigma_v = \frac{\sum v^2}{N_p} - \left(\frac{\sum v}{N_p}\right)^2,$$

$$\sigma_{xv} = \frac{\sum xv}{N_p} - \frac{\sum x \sum v}{N_p^2},$$

and all sums are performed for the N_p particles in the beam, i.e., $\sum x = \sum_{i=1}^{N_p} x_i$.

[1] R. C. Davidson, *Physics of Nonneutral Plasmas* (World Scientific Publishing Company, 2001).

[2] D. H. E. Dubin and T. M. O'Neil, Trapped nonneutral plasmas, liquids, and crystals (the thermal equilibrium states), Rev. Mod. Phys. **71**, 87 (1999).

[3] G. Maero, E. D. Hunter, D. J. Murtagh, and E. V. Stenson, Fundamental physics and other applications using nonneutral plasma, Adv. Phys.:X **9**, 2367438 (2024).

[4] J. Fajans and C. M. Surko, Plasma and trap-based techniques for science with antimatter, Phys. Plasmas **27**, 030601 (2020).

[5] G. B. Andresen *et al.*, Confinement of antihydrogen for 1,000 seconds, Nat. Phys. **7**, 558 (2011).

[6] G. Gabrielse, W. S. Kolthammer, R. McConnell, P. Richerme, R. Kalra, E. Novitski, D. Grzonka, W. Oelert, T. Sefzick, M. Zielinski, D. Fitzakerley, M. C. George, E. A. Hessels, C. H. Storry, M. Weel, A. Müllers, and J. Walz, Adiabatic cooling of antiprotons, Phys. Rev. Lett. **106**, 073002 (2011).

[7] Y. Enomoto, N. Kuroda, K. Michishio, C. H. Kim, H. Higaki, Y. Nagata, Y. Kanai, H. A. Torii, M. Corradini, M. Leali, E. Lodi-Rizzini, V. Mascagna, L. Venturelli, N. Zurlo, K. Fujii, M. Ohtsuka, K. Tanaka, H. Imao, Y. Nagashima, Y. Matsuda, *et al.*, Synthesis of cold antihydrogen in a cusp trap, Phys. Rev. Lett. **105**, 243401 (2010).

[8] P. Perez, D. Banerjee, F. Biraben, D. Brook-Roberge, M. Charlton, P. Cladé, P. Comini, P. Crivelli, O. Dalkarov, P. Debu, *et al.*, The GBAR antimatter gravity experiment, Hyperfine Interact. **233**, 21 (2015).

[9] E. Anderson, C. Baker, W. Bertsche, N. M. Bhatt, G. Bonomi, A. Capra, I. Carli, C. L. Cesar, M. Charlton, A. Christensen, *et al.*, Observation of the effect of gravity on the motion of antimatter, Nature (London) **621**, 716 (2023).

[10] D. W. Fitzakerley, M. C. George, E. A. Hessels, T. D. G. Skinner, C. H. Storry, M. Weel, G. Gabrielse, C. D. Hamley, N. Jones, K. Marable, E. Tardiff, D. Grzonka, W. Oelert, and M. Zielinski (ATRAP Collaboration), Electron-cooled accumulation of 4 × 109 positrons for production and storage of antihydrogen atoms, J. Phys. B: At. Mol. Opt. Phys. **49**, 064001 (2016).

[11] G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, P. D. Bowe, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, J. S. Hangst, W. N. Hardy, R. S. Hayano, M. E. Hayden, A. Humphries, R. Hydomako, S. Jonsell, *et al.* (ALPHA Collaboration), Evaporative cooling of antiprotons to cryogenic temperatures, Phys. Rev. Lett. **105**, 013003 (2010).

[12] C. Amsler, H. Breuker, S. Chesnevskaya, G. Costantini, R. Ferragut, M. Giammarchi, A. Gligorova, G. Gosta, H. Higaki, E. D. Hunter, C. Killian, V. Kletzl, V. Kraxberger, N. Kuroda, A. Lanz, M. Leali, V. Mäckel, G. Maero, C. Malbrunot, V. Mascagna, *et al.*, Reducing the background temperature for cyclotron cooling in a cryogenic Penning–Malmberg trap, Phys. Plasmas **29**, 083303 (2022).

[13] G. Manfredi and P.-A. Hervieux, Adiabatic cooling of trapped non-neutral plasmas, Phys. Rev. Lett. **109**, 255005 (2012).

[14] G. Z. Li, R. Poggiani, G. Testera, and G. Werth, Adiabatic cooling of ions in the Penning trap, Z. Phys. D: At. Mol. Clusters **22**, 375 (1991).

[15] P. Blumer, M. Charlton, M. Chung, P. Cladé, P. Comini, P. Crivelli, O. Dalkarov, P. Debu, L. Dodd, A. Douillet, S. Guellati, P.-A. Hervieux, L. Hilico, A. Husson, P. Indelicato, G. Janka, S. Jonsell, J.-P. Karr, B. Kim, E. Kim, *et al.*, Positron accumulation in the GBAR experiment, Nucl. Instrum. Methods Phys. Res. Sect. A **1040**, 167263 (2022).

[16] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga, Shortcuts to adiabaticity: Concepts, methods, and applications, Rev. Mod. Phys. **91**, 045001 (2019).

[17] J.-L. Lions, *Optimal Control of Systems Governed by Partial Differential Equations*, Die Grundlehren der mathematischen Wissenschaften (Springer-Verlag, New York-Berlin, 1971), Vol. Band 170, pp. xi + 396.

[18] E. D. Sontag, *Mathematical Control Theory*, 2nd ed., Texts in Applied Mathematics (Springer-Verlag, New York, 1998), Vol. 6, p. xvi + 531.

[19] K. Ito and K. Kunisch, *Lagrange Multiplier Approach to Variational Problems and Applications*, Advances in Design and Control (Society for Industrial and Applied Mathematics, Philadelphia, 2008), Vol. 15, p. xviii+341.

[20] F. Tröltzsch, *Optimal Control of Partial Differential Equations: Theory, Methods, and Applications*, Graduate Studies in Mathematics (American Mathematical Society, Providence, RI, 2010), Vol. 112, p. xvi + 399.

[21] A. Borzì and V. Schulz, *Computational Optimization of Systems Governed by Partial Differential Equations*, Computational Science & Engineering (Society for Industrial and Applied Mathematics, Philadelphia, PA, 2012), Vol. 8, p. xx + 282.

[22] P. L. Bhatnagar, E. P. Gross, and M. Krook, A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems, Phys. Rev. **94**, 511 (1954).

[23] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed., Springer Series in Operations Research and Financial Engineering (Springer, New York, 2006), p. xxii + 664.

[24] E. Hairer, C. Lubich, and G. Wanner, Geometric numerical integration illustrated by the Störmer–Verlet method, Acta Numerica **12**, 399 (2003).

[25] G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, P. D. Bowe, E. Butler, P. T. Carpenter, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, J. S. Hangst, W. N. Hardy, M. E. Hayden, A. J. Humphries, J. L. Hurt, R. Hydomako, *et al.* (ALPHA Collaboration), Autoresonant excitation of antiproton plasmas, Phys. Rev. Lett. **106**, 025002 (2011).

[26] M. E. Glinsky, T. M. O'Neil, M. N. Rosenbluth, K. Tsuruta, and S. Ichimaru, Collisional equipartition rate for a magnetized pure electron plasma, Phys. Fluids B **4**, 1156 (1992).

[27] G. von Winckel and A. Borzì, Computational techniques for a quantum control problem with H^1-cost, Inverse Probl. **24**, 034007 (2008).

[28] A. Chertock, A practical guide to deterministic particle methods, in *Handbook of Numerical Methods for Hyperbolic Problems* (Elsevier/North-Holland, Amsterdam, 2017), Vol. 18, pp. 177–202.

[29] G.-H. Cottet and P.-A. Raviart, Particle methods for the one-dimensional Vlasov-Poisson equations, SIAM J. Numer. Anal. **21**, 52 (1984).

[30] T. Barth, R. Herbin, and M. Ohlberger, Finite volume methods foundation and analysis, in *Encyclopedia of Computational Mechanics Second Edition*, edited by E. Stein, R. Borst, and T. J. R. Hughes (Wiley, New York, 2017).

[31] B. Mortier, M. Baelmans, and G. Samaey, A kinetic-diffusion asymptotic-preserving Monte Carlo algorithm for the Boltzmann-BGK model in the diffusive scaling, SIAM J. Sci. Comput. **44**, A720 (2022).

[32] J. Bartsch, P. Knopf, S. Scheurer, and J. Weber, Controlling a vlasov–poisson plasma by a particle-in-cell method based on a monte carlo framework, SIAM J. Control Optim. **62**, 1977 (2024).

[33] S. D. Conte and C. de Boor, *Elementary Numerical Analysis* (McGraw-Hill Book Co., 1980).

[34] K. Floettmann, Some basic features of the beam emittance, Phys. Rev. Spec. Top.–Accel. Beams **6**, 034202 (2003).