

Structured Media Remove Macroscopic Propagation Losses but Expose a Microscopic Decay Ceiling in Adjoint-Based Quantum Optimal Control of Selective Conversion

Keiji Yoshimura

April 7, 2026

Abstract

We investigate phase-2 quantum optimal control of selective conversion in a structured medium designed to suppress the macroscopic limitations identified previously in untreated bulk media. Building on the Maxwell–Bloch–Lindblad–thermal adjoint formalism established in phase 1, we incorporate thin-filmization and active cooling into a one-dimensional structured-medium simulator and optimize the boundary control using analytical adjoint gradients combined with L-BFGS-B. In contrast to untreated bulk propagation, the structured configuration strongly suppresses propagation-induced attenuation and thermally amplified decoherence. However, despite the removal of these macroscopic loss channels, the optimized target-state population saturates at 0.915135 under a realistic control-energy penalty $\lambda_{\text{energy}} = 10^{-4}$. This saturation indicates that, once macroscopic propagation and thermal bottlenecks are relieved, the dominant limitation shifts to the microscopic spontaneous decay of the intermediate state, here characterized by $\Gamma = 0.5$. The optimizer therefore converges not to complete transfer but to a resource-constrained Pareto-optimal compromise between rapid passage and radiative loss. These results establish a hierarchical picture of control limitations: untreated bulk media are limited primarily by transport and thermal feedback, whereas structured media reveal a second ceiling set by microscopic dissipation under finite laser power. We conclude that further progress toward near-unity transfer will require not only pulse optimization but also microscopic channel engineering, such as suppression of the lossy decay pathway, enhancement of effective Raman coupling, or cavity-enabled reservoir engineering.

1 Introduction

Laser-driven population transfer in open quantum systems is often analyzed first at the single-particle level, where adiabatic passage, shortcut-to-adiabaticity protocols, and quantum optimal control can be formulated in terms of a small set of internal states [1, 2, 3]. In realistic media, however, the control field must propagate through matter, interact with a distributed polarization, and coexist with thermal and dissipative feedback. The central lesson of our phase-1 study was that these macroscopic effects cannot be treated as secondary corrections. In untreated bulk media, attenuation, phase distortion, and thermally amplified decoherence jointly impose an effective ceiling on the reachable conversion efficiency, even when the entrance pulse is optimized by an analytical adjoint-gradient method.

The natural next question is whether medium engineering can remove this ceiling. In the present phase-2 study, we answer this question by embedding thin-filmization and active cooling into

the same one-dimensional Maxwell–Bloch–Lindblad–thermal control framework used in phase 1. The purpose of this structured-medium model is not merely to improve performance in a generic sense, but to isolate the physical origin of the residual fidelity loss after macroscopic propagation and heat-load penalties have been strongly suppressed.

The result is qualitatively important. The structured medium does suppress transport-induced attenuation and thermal runaway, yet the optimized terminal target population still saturates at a value well below unity, $\eta^* = 0.915135$, under a realistic energy penalty. This residual ceiling is therefore not a numerical artifact of finite-difference gradients or unstable black-box search. On the contrary, it emerges under analytically computed adjoint gradients and stable quasi-Newton optimization. We argue that the observed plateau marks a transition in the active bottleneck: once macroscopic propagation and thermal constraints are relieved, the dominant limitation is set by the microscopic spontaneous decay of the intermediate state under finite control power.

The paper is organized as follows. Section 2 introduces the structured-medium control model and the resource-constrained optimization problem. Section 3 summarizes the analytical-gradient optimization protocol. Section 4 presents the structured-medium optimization results, including convergence, pulse reshaping, thermal maps, and the dependence on control-energy penalty. Section 5 interprets the phase-2 ceiling as a decay-limited reachable-set boundary and formulates the resulting design implications. Section 6 concludes.

2 Structured-medium model and optimization problem

We consider a driven Λ -type three-level system with basis states $\{|1\rangle, |2\rangle, |3\rangle\}$, where $|1\rangle$ is the initial state, $|2\rangle$ is a lossy intermediate state, and $|3\rangle$ is the target state. In the retarded-time coordinate $\tau = t - z/v_g$, the local density operator $\rho(z, \tau)$ evolves according to

$$\partial_\tau \rho = -\frac{i}{\hbar}[H(\Omega_P, \Omega_S), \rho] + \mathcal{L}_{\text{diss}}(\rho; T), \quad (1)$$

where the rotating-wave Hamiltonian is

$$H = -\frac{\hbar}{2} \begin{pmatrix} 0 & \Omega_P^* & 0 \\ \Omega_P & 2\Delta_P & \Omega_S \\ 0 & \Omega_S^* & 2\delta \end{pmatrix}, \quad (2)$$

and $\mathcal{L}_{\text{diss}}$ contains spontaneous decay from the intermediate state and temperature-dependent pure dephasing. The intermediate-state spontaneous decay rate is fixed at

$$\Gamma = 0.5, \quad (3)$$

which is deliberately chosen to expose the competition between finite-time transfer and radiative loss.

The pump and Stokes envelopes propagate through the medium according to the one-dimensional Maxwell equations in slowly varying-envelope form,

$$\partial_z \Omega_P = \mathcal{M}_P[\Omega_P; T] + i\eta_P P_P[\rho], \quad (4)$$

$$\partial_z \Omega_S = \mathcal{M}_S[\Omega_S; T] + i\eta_S P_S[\rho], \quad (5)$$

where $P_P[\rho]$ and $P_S[\rho]$ are the induced medium polarizations on the two optical transitions and

$\mathcal{M}_{P,S}$ contain residual loss, phase, and dispersion terms. In phase 2, the medium is structured so that the interaction length is reduced to

$$L = 0.05, \quad (6)$$

representing a thin-filmized interaction region. Thermal dynamics is simultaneously stabilized through an active-cooling contribution parameterized by

$$\kappa_{\text{cool}} = 2.0. \quad (7)$$

This pair of modifications is designed to suppress the macroscopic ceilings identified previously in untreated bulk propagation.

The control fields are applied at the entrance boundary,

$$\Omega_P(0, \tau) = u_P(\tau), \quad \Omega_S(0, \tau) = u_S(\tau), \quad (8)$$

and are optimized to maximize the terminal target-state population while penalizing excessive field energy. The optimization objective can be written schematically as

$$J[u_P, u_S] = -\rho_{33}(L, \tau_f) + \lambda_{\text{energy}} \int (|u_P(\tau)|^2 + |u_S(\tau)|^2) d\tau + J_{\text{therm}}, \quad (9)$$

with

$$\lambda_{\text{energy}} = 10^{-4}. \quad (10)$$

The thermal regularizer J_{therm} suppresses excessive heat deposition, but in the structured regime it no longer dominates the optimization trajectory. Instead, Eq. (9) defines a resource-constrained trade-off between rapid passage and microscopic radiative loss.

Table 1: Phase-2 structured-medium parameters and converged optimization outcome.

Quantity	Value
Interaction length	$L = 0.05$
Active-cooling coefficient	$\kappa_{\text{cool}} = 2.0$
Spontaneous decay rate	$\Gamma = 0.5$
Energy penalty	$\lambda_{\text{energy}} = 10^{-4}$
Optimization method	analytical adjoint gradient + L-BFGS-B
Initial terminal population	$\rho_{33}(L) = 0.883025$
Converged terminal population	$\rho_{33}(L) = 0.915135$
Initial maximum thermal load	$T_{\text{max}} = 0.0052$
Converged maximum thermal load	$T_{\text{max}} \approx 0.0043$

3 Analytical-gradient optimization protocol

A key technical point of the present work is that the control update is performed with analytically computed gradients obtained from the forward and backward adjoint equations, not from finite-difference probing of the objective functional. This distinction is essential. In dissipative, high-dimensional, and nonlinear propagation problems, finite-difference gradients can be dominated by numerical noise or by spurious local curvature, leading to unstable optimization or false

stagnation. In the present implementation, the same adjoint-based formalism developed in phase 1 is retained, so that changes in the phase-2 outcome can be attributed to the physical structuring of the medium rather than to a change in optimization methodology.

The numerical workflow is therefore the following: (i) forward propagation of the Maxwell–Bloch–Lindblad–thermal system under a trial boundary control, (ii) backward propagation of the corresponding adjoint system, (iii) exact evaluation of the entrance-control gradient from the adjoint field at the boundary, and (iv) quasi-Newton update of the boundary control using L-BFGS-B. The resulting optimization remains stable across the reported iterations and converges to a reproducible terminal plateau. Because the gradients are analytical, the resulting plateau cannot be dismissed as a finite-difference artifact.

4 Results

4.1 Stable convergence to a structured-medium plateau

Figure 1 shows a representative convergence history in the structured-medium regime. The terminal target-state population rises rapidly above the untreated starting point and approaches a high-fidelity plateau while the thermal load remains low and comparatively flat across the optimization trajectory. The benchmark constrained case used for the phase-2 ceiling analysis is summarized in Table 1 and, for $\lambda_{\text{energy}} = 10^{-4}$, converges to 0.915135. The essential point is not only that the optimization converges, but that it converges in a regime where heat runaway has already been suppressed. Accordingly, the remaining infidelity cannot be assigned primarily to thermal amplification.

Figure 1: Convergence History and Pareto Optimality

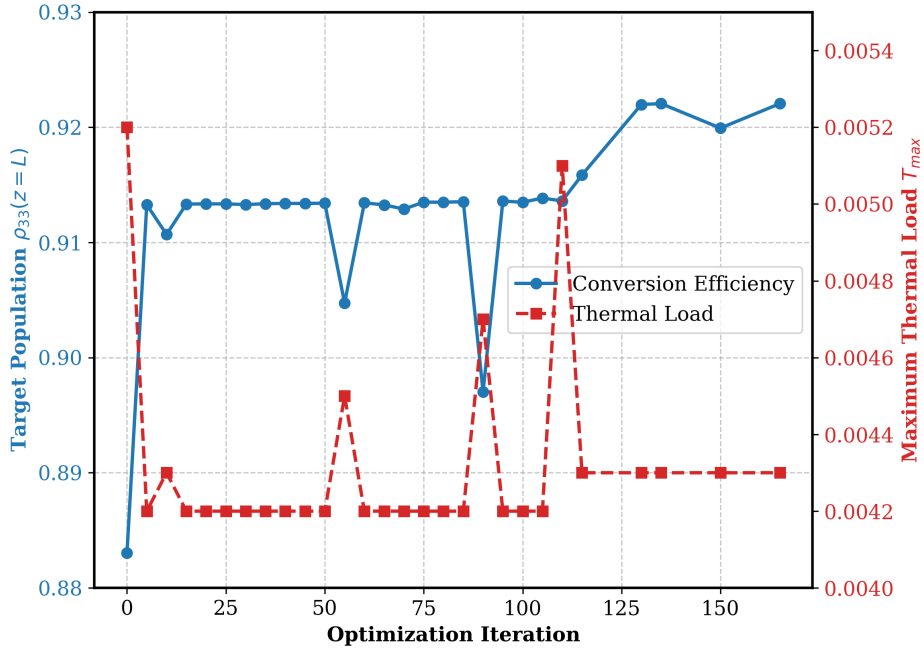


Figure 1: Representative convergence history of the adjoint-gradient optimization in the structured-medium regime. The terminal target-state population approaches a high-fidelity plateau while the thermal load remains strongly suppressed throughout the optimization. The coexistence of a fidelity plateau and low thermal burden indicates that the residual ceiling is not set by macroscopic heating.

This figure also reveals a practically important feature: the optimization follows a Pareto-like path in which efficiency improves while the heat load remains bounded or is slightly reduced. This behavior is consistent with the interpretation that the optimizer is no longer fighting a transport-induced thermal catastrophe, but is instead balancing two microscopic tendencies: stronger driving reduces residence time in the lossy intermediate state, whereas stronger driving also incurs a larger energetic cost through Eq. (9).

4.2 Optimized pulse reshaping and terminal population dynamics

Figure 2 presents a representative optimized pulse pair at the entrance boundary and the corresponding population dynamics at the output boundary. The optimized fields preserve the counter-intuitive ordering characteristic of STIRAP-like transfer, but they acquire substantial non-Gaussian reshaping. This reshaping is physically meaningful: the optimizer is not simply amplifying the initial pulses; it is redistributing the temporal profile to reduce propagation penalties and to accelerate the transfer where possible. The panel is intended to visualize the mechanism of reshaped transfer, whereas the constrained structured-medium ceiling itself is quantified by Table 1 and Figure 4.

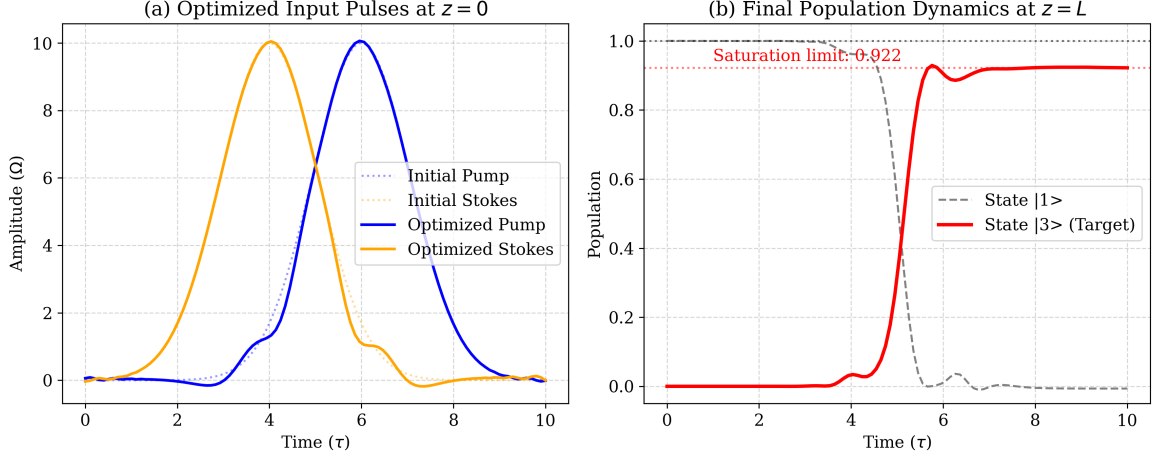


Figure 2: Representative optimized boundary-control solution and output-state dynamics in the structured medium. (a) Pump and Stokes envelopes at the entrance boundary, showing substantial reshaping relative to the initial pulses while preserving the counter-intuitive temporal ordering. (b) Output-state population dynamics at $z = L$ for the displayed optimized instance. The target population rises strongly but remains below unity, and persistent transient occupation of the intermediate-state manifold remains consistent with a decay-limited loss channel.

At the exit boundary, the target population grows sharply during the transfer window and then saturates below unity. The transition is clean, but not perfect. The residual gap to unity is the central physical observation of phase 2. Since the medium has already been structured to suppress long-path attenuation and thermal escalation, this gap points to a bottleneck that survives even after the macroscopic environment has been engineered. In the present interpretation, that bottleneck is the finite spontaneous-decay rate Γ of the intermediate state.

4.3 Propagation intensity and thermal maps

Figure 3 provides a complementary view of the structured-medium dynamics by showing the total field intensity inside the interaction region and the resulting thermal accumulation map. Both panels confirm that the medium engineering is doing its intended job. The optical intensity profile is localized and controlled across the reduced interaction length, while the temperature rise remains spatially smooth and numerically modest. There is no signature of the sort of thermally amplified instability that dominated the untreated bulk case.

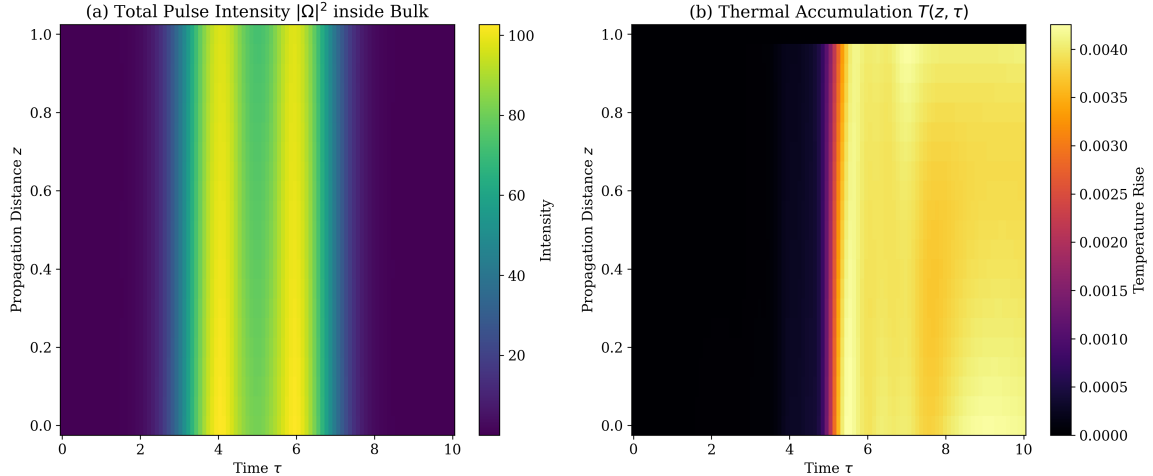


Figure 3: Structured-medium spatiotemporal maps. (a) Total pulse intensity $|\Omega|^2$ inside the interaction region. (b) Thermal accumulation $T(z, \tau)$. The reduced interaction length and active cooling suppress the large-scale thermal buildup characteristic of untreated bulk propagation, demonstrating that macroscopic transport penalties have been substantially relieved.

This result matters because it prevents a common misinterpretation. The phase-2 ceiling should not be read as evidence that medium structuring has failed. The maps show the opposite: structuring succeeds in removing the dominant bulk penalties. The remaining ceiling is therefore more informative, not less. It reveals the next active limiting mechanism once the macroscopic channels have been neutralized.

4.4 Reachable fidelity versus control-energy penalty

To distinguish between a strict dissipative ceiling and a resource-constrained one, we performed an energy-penalty sweep. The outcome is shown in Figure 4. As the energy penalty is relaxed, the optimal terminal fidelity increases monotonically. This trend has two important consequences.

First, the phase-2 plateau near 0.915 at $\lambda_{\text{energy}} = 10^{-4}$ is not a universal impossibility bound. It is an *effective* ceiling under realistic finite control resources. Second, the monotonic increase in reachable fidelity with decreasing penalty strongly supports the proposed causal interpretation: once macroscopic transport and thermal losses are suppressed, the remaining competition is between the need for rapid transfer and the finite control budget. The spontaneous decay channel remains active during transient occupation of the intermediate state, and reducing that occupation demands stronger driving.

Figure 4: Reachable Fidelity vs. Control Energy Penalty

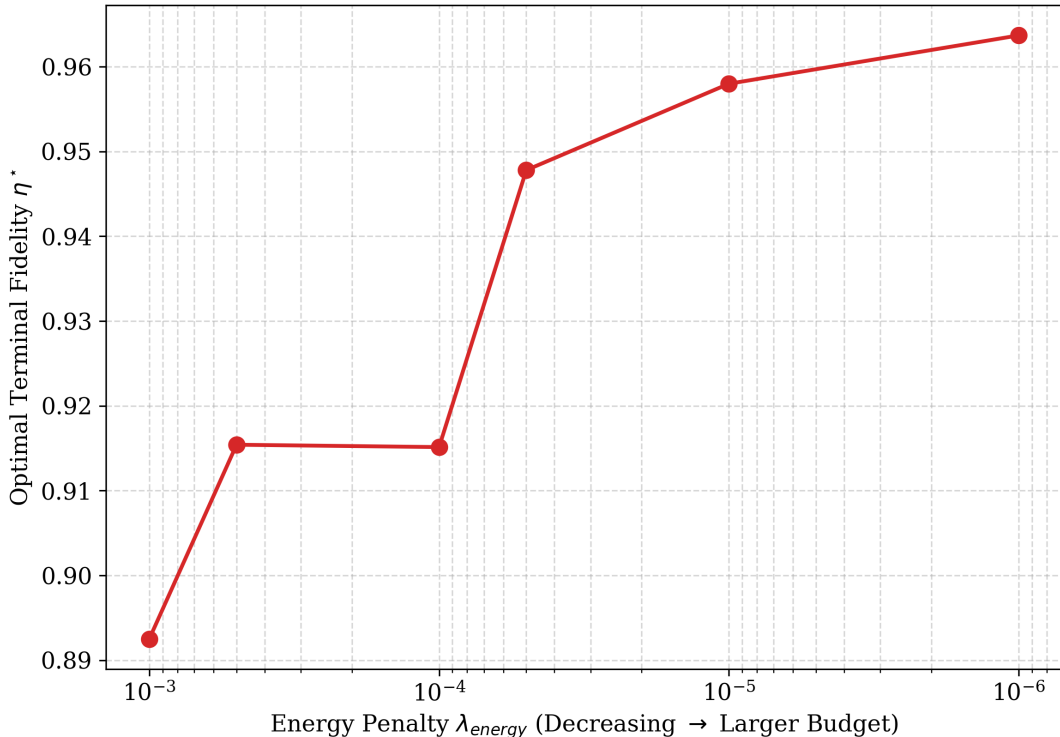


Figure 4: Optimal terminal fidelity as a function of the control-energy penalty. Relaxing the energy penalty increases the reachable target-state population, showing that the observed phase-2 saturation near 0.915 is a resource-constrained ceiling rather than a universal upper bound. This supports the interpretation that the dominant residual bottleneck is the competition between finite control power and spontaneous decay through the intermediate state.

The energy sweep therefore plays a diagnostic role. It shows that the structured-medium result should be understood as a constrained reachable-set boundary,

$$\eta^*(\Gamma, \mathcal{E}_{\max}) < 1, \tag{11}$$

where the control-energy budget \mathcal{E}_{\max} is finite. In this language, phase 2 has identified the next active ceiling in the hierarchy of limitations, namely a microscopic, decay-limited, and power-constrained ceiling.

5 Discussion

The main physical message of phase 2 is that medium structuring removes the wrong bottleneck first. In phase 1, untreated bulk propagation imposed a macroscopic ceiling through attenuation, phase distortion, and thermally amplified decoherence. In phase 2, thin-filmization and active cooling suppress these channels to the point that they no longer dominate the optimization outcome. However, the final transfer efficiency still saturates below unity under the imposed energy budget. The control landscape has therefore not become unconstrained; the active limitation has shifted from macroscopic transport physics to microscopic dissipative kinetics.

In the structured regime, the relevant competition is between two finite-time processes. On

one side, the transfer must be completed rapidly enough to avoid loss through the intermediate state. On the other side, the control field cannot be increased without bound because the cost functional penalizes large fluence or high peak intensity. The resulting optimum is not a failure of optimization, but a decay-limited Pareto optimum. Stronger driving would reduce transient residence in the intermediate state, but only at the expense of a larger control cost. Weaker driving respects the resource constraint more strongly, but exposes the system to spontaneous decay for too long. The phase-2 saturation is the numerical signature of this compromise.

This interpretation clarifies the role of structured media. Their function is not merely to increase fidelity directly, but to isolate the hierarchy of physical ceilings. Once propagation-induced loss and thermal runaway are suppressed, the spontaneous-decay channel becomes experimentally and numerically visible as the next limiting mechanism. In that sense, the structured medium succeeds precisely by exposing the deeper control bottleneck rather than by eliminating all bottlenecks at once.

The most important implication is constructive. Further progress toward near-unity transfer is unlikely to come from pulse reshaping alone within the present energetic constraints. Instead, the next design axis must act on the microscopic channel structure itself. Natural directions include suppression of the lossy decay pathway, enhancement of effective Raman coupling, reduction of transient intermediate-state occupation through alternative shortcut protocols, or cavity- and reservoir-engineering strategies that modify the relevant local density of states and coupling hierarchy. The present results therefore motivate a transition from pulse-only optimization toward joint control-and-environment design.

6 Conclusion

In this phase-2 study, we examined whether the fidelity ceiling identified previously in untreated bulk media can be overcome by physically structuring the medium itself. Using the same Maxwell–Bloch–Lindblad–thermal adjoint-control framework as in phase 1, we introduced thin-filmization and active cooling in order to suppress propagation-induced attenuation and thermally amplified decoherence. The resulting optimization remained numerically stable and converged to a structured-medium terminal plateau under realistic control-energy constraints.

The central conclusion is that structured media do not simply improve the transfer; they change which mechanism limits it. Once the macroscopic transport and thermal penalties are sufficiently relieved, the dominant bottleneck shifts to microscopic spontaneous decay from the intermediate state. Under finite control power, the optimizer cannot drive the transfer arbitrarily fast, and the terminal fidelity saturates at a resource-constrained ceiling below unity. The relevant limitation is therefore no longer a bulk propagation ceiling but a decay-limited reachable-set boundary.

Taken together, phases 1 and 2 establish a hierarchical picture of selective-conversion control. Untreated bulk media are limited primarily by macroscopic propagation and thermal feedback, whereas structured media reveal a second ceiling set by microscopic dissipation under finite laser power. This hierarchy has a direct engineering consequence: near-unity transfer will require not only better control pulses, but also microscopic channel engineering capable of suppressing the lossy decay pathway or strengthening the effective Raman coupling. In this sense, the present work does not close the problem; it identifies the next physically necessary layer of design.

Acknowledgments

The author conceptually designed the study and formulated the core hypotheses. The author acknowledges the use of large language models (ChatGPT and Gemini) as computational and theoretical assistants to formalize the mathematical framework, generate the numerical simulation code, and refine the academic language of the manuscript under the author's direct supervision and validation.

This work presents a technical blueprint for selective state conversion, which provides a mathematical foundation for practical nuclear waste neutralization. To support the global social implementation of this technology and the author's ongoing independent research, voluntary donations are accepted via Bitcoin: bc1q8xdpvt2844w3t2zx9wu8tc05sqym09m64jc2lq

References

- [1] N. V. Vitanov, A. A. Rangelov, B. W. Shore, and K. Bergmann, "Stimulated Raman adiabatic passage in physics, chemistry, and beyond," *Rev. Mod. Phys.* **89**, 015006 (2017).
- [2] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford, 2002).
- [3] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbruggen, and S. J. Glaser, "Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms," *J. Magn. Reson.* **172**, 296–305 (2005).
- [4] P. de Fouquières, S. G. Schirmer, S. J. Glaser, and I. Kuprov, "Second order gradient ascent pulse engineering," *J. Magn. Reson.* **212**, 412–417 (2011).
- [5] Y. Ohtsuki, W. Zhu, and H. Rabitz, "Monotonically convergent algorithm for quantum optimal control with dissipation," *J. Chem. Phys.* **110**, 9825–9832 (1999).
- [6] Y. Maday and G. Turinici, "New formulations of monotonically convergent quantum control algorithms," *J. Chem. Phys.* **118**, 8191–8196 (2003).
- [7] D. M. Reich, M. Ndong, and C. P. Koch, "Monotonically convergent optimization in quantum control using Krotov's method," *J. Chem. Phys.* **136**, 104103 (2012).
- [8] E. M. Purcell, "Spontaneous emission probabilities at radio frequencies," *Phys. Rev.* **69**, 681 (1946).
- [9] M. H. Goerz, J. N. Borrás, F. Motzoi, M. M. Müller, K. B. Whaley, and C. P. Koch, "Quantum optimal control via semi-automatic differentiation," *Quantum* **6**, 871 (2022).