RESPONSE TO REFEREE #1

(*) "My main critique of the paper, is that there does not seem to be a clear novel improvement over previous works e.g. Refs. 16 and 17. The author must make clear why their contribution is not an incremental improvement over previous papers."

This work proposes a novel optimization method for quantum thermal machines modeled as Markovian open quantum systems. It should not be seen as an attempt to improve on previous proposals, but as an alternative route. The refs. [16] and [17] (all reference numbers in this response refer to the first manuscript, not the revised version) that the referee points out have been, indeed, the inspiration to work on this project (which is why I also decided to use one of the models that is described in those papers). In those references, the authors present an optimization scheme for QTMs based on reinforcement learning (RL). And, for example in [17], they write:

"The optimization of F in this form is precisely the type of problem that can be readily tackled using optimization techniques, such as the Pontryagin minimum principle or reinforcement learning. In this Letter, we employ the latter."

The authors of Ref. [17] therefore use RL. However, there are very few works that have used the other possibility that they mention in that paragraph: Pontryagin's minimum principle (PMP), the standard workhorse of optimal control theory, to approach this problem (periodically driven quantum open systems modelling quantum thermal machines). The reason is probably that the field of quantum optimal control for *periodic* systems has not received much attention, and is therefore less developed. In essence, the PMP establishes necessary conditions for the extrema of the target function: basically, the gradient of the function should be zero (or the functional derivative, as the argument of the target function is often a function itself). RL, in contrast, is one of the techniques that permit to do the optimization ignoring the gradient.

The only articles (to my knowledge) where the PMP is used explicitly for QTMs modeled with Markovian dynamics are the works of Cavina and collaborators, such as ref. [20]. In that work, they use unconstrained real-time representations of the control functions -- which is the usual way in which PMP operates --, and arrive to the so-called bang-bang (discontinuous) solutions for the problems that they study. The problem of enforcing periodicity does not seem to cause important troubles in that work.

However, a natural route to the optimization of periodic control problems is the use of the spectral method (I have cited ref. [23] as an example). Using a Fourier series as a basis automatically enforces the periodicity, and furthermore, permits to apply bandwidth constraints in a natural and easy way. Strict bang-bang solutions, for example, would not be allowed because they are discontinuous, and only approximations could appear as solutions -- which reflects the real experimental situation. And, using the spectral method, one can derive equations for the gradient of the functional -- coming back to the spirit of the PMP (although this one is normally formulated in real time).

Therefore, my goal has been to explore this path: using gradient-based spectral methods for the optimization of QTMs. I have made this purpose more clear in the introduction of the manuscript. I have also slightly changed the title.

(*) "The introduction is well structured, but does have some strange and redundant phrasing at times. It could be improved upon."

I have rewritten parts of the introduction.

(*) "Fig. 1 is never referenced in the main text that I can see."

Fixed.

(*) "It is not clear what are the main quantum effects being exploited. Does the coherence of the qubit play an important role? What results would one get from an analogous classical master equation?"

That aspect has not been studied.

(*) "Some notation I find a bit awkward such as with f(t) and f_k(t). Maybe the use of vectors would be useful here. Similarly with u and u^{(k)}."

Changed following the suggestion of the referee.

(*) “The author notes that the Lamb-shift “can be ignored in the weak coupling limit”. Recent discussion here [arXiv:2305.08941] suggests that this is a somewhat subtle point. Some more discussion on this is probably warranted."

The method does not depend on whether or not the Lamb shift is ignored. The Hamiltonian that is used for the coherent part of the Hamiltonian may or may not include the Lamb shift, and the working equations for performing the optimizations would be the same.

However, it is true that in order to exemplify the method I have used a model that does not include a Lamb shift. In any case, I have improved the wording of the paragraph in page 4, after Eq. (4), where the sentence that the referee points out was written. Also, I have included the reference pointed out by the referee.

(*) "The notation \rho is used to denote both the general solution to the master equation and the NESS. I would suggest using a subscript for the NESS for clarity."

I have explicitly labeled \rho in Eq. (5), but right afterwards I included a sentence warning that all appearances of rho thereafter correspond to a NESS (otherwise, the notation becomes too heavy).

(*) "There has been recent work e.g. [Quantum 5, 590 (2021) and Sci. Adv.8, eadd0828 (2022)] which show how the usual Lindblad master equation must be modified to ensure thermodynamic consistency when the system is driven. - Since the current work is focussed on quantifying thermodynamics performance, should this modified master equation be employed?"

The method described in Section 3 is generic: for any form of GKSL equation, which is the most general form of the generator of the dynamics of an open system coupled to Markovian baths.

It is true that, in the sample demonstrations that are shown in Section 4, I use a particular form of the GKSL equations, for an also particular model. In this paper I do not enter the discussion about the validity of that model -- I simply rely on the fact that it has been used succesfully in the past. I merely use it as an example to demonstrate the computational feasibility of the method.

(*) "The main focus has been on optimising output power. Does this have any tradeoff relations with the efficiency (even from the numerics)? Would it make sense to use a joint cost function which combines the two?"

It is true that the focus of the sample calculations has been placed on the output power, but in the description of the method I have layed out a general formula (Eq. 34) that shows how one could define, for example, a joint cost function that would combine power output and efficiency, as the referee suggests. Numerically, it would not imply a much larger cost, as the main difficulty is the computation of the gradient of the NESS with respect to the parameters. That is the same ingredient for all the derivatives in the sum appearing in Eq. (34).

(*) "It is claimed that the Lindblad operators can be modulated in time. This is rather unusual. It would be nice if some physical insight into how this is done in practice was included."

The model that I used as an example does not have time-dependent Lindblad operators; it is described for example in [Erdman and Noe, npj Quantum Information (2022) 1]: the dissipator is given in Eq. (9), and then the Lindblad operators in Eq. (11). But note that the dissipator in Eq. (9) allows for time-dependent Lindblad operators; that possibility is used to describe in that same paper a different model -- see Eq. (12).

When deriving Lindblad-type equations for driven open quantum systems, one may arrive to time-dependent Lindblad operators -- see for example, for periodically driven systems, [Kammleitner and Shnirma, PRB 84, 235140 (2011)].

The method described in Section 3 allows, in principle, for that possibility. However, the practical implementation in the code does not allow for it, yet (work in progress).

(*) "“...which in a realistic setup cannot be taken to arbitrarily close-to-zero values”. It would be useful to state how small this can really be made for the setups of interest."

Unfortunately, I lack the experimental knowledge to estimate the real limitations. My point is merely to state the fact, in a realistic setup, transitions have to be continusous and smooth.

(*) "Panel (b) of Figs. 2 and 3 need units for the y-axis."

Fixed. I also changed the scaling in panel (a) to make it more clear.

(*) "For the parameterisation of the control function, would it not be simpler to just use a Fourier series (as in Eq. A2) and just bound the sum of the absolute values of the amplitudes to cap the amplitude?"

Indeed, that would be a simpler expression. And, if one bounds the sum of the absolute values of the Fourier amplitudes, the instantaneous amplitude is also capped at all times. I have tried the suggestion of the referee, and it works, although it seems to lead in practice to amplitudes that are much lower than the bound (the constraint is too "demanding"). As a result, the optimized values that I found are lower.

Note that even if the parametrization reads simpler, it does not imply a real computational advantage. Therefore, I prefer to leave the results obtained with the parametrization that I used before. However, I have added the suggestion of the referee to the text.
