From:
doc. Ing. Zdeněk Hurák, Ph.D.
Dept. of Control Engineering, Faculty of Electrical Engineering
Czech Technical University, Technická 2, 166 27 Prague, Czechia
hurak@fel.cvut.cz, http://aa4cc.dce.fel.cvut.cz
Phone: +420 224 357 683

June 21, 2024

To:
doc. Ing. Václav Čuba, Ph.D.
Dean of Faculty of Nuclear Sciences and Physical Engineering

Dear Dean of Faculty of Nuclear Sciences and Physical Engineering,

it is my pleasure to submit the **Review of Jiří Fejlek's dissertation**, which you asked me for earlier this spring.

The ultimate goal set for the submitted dissertation was to develop computational design methods for optimal control of nonlinear dynamical systems such as robots and unmanned vehicles. The approach pursued in the dissertation was to **improve/extend** one particular existing framework described in the literature – the framework of **LQR trees** proposed by Russ Tedrake of MIT and his colleagues some fifteen years ago. The key idea of both the original work and the dissertation is that the state space is split into regions (called funnels here) within which the system is controlled towards the target state (or target region in the state space) by an LQR state feedback regulator tracking a precomputed state trajectory. Finding this partition of the state space and computing the corresponding optimal state trajectory for each funnel along with the acompanying state feedback regulators constitute the major theoretical and computational challenges for the thesis.

I appreciate that the Ph.D. candidate decided to work on improving existing methodology, even though the low hanging fruit had already been picked by the authors of the original LQR-tree framework and what remained was rather challenging technical tasks. The candidate has succeeded in tackling some of those, and in his dissertation he offers both some needed **theoretical results/guarantees** for LQR-tree control design methods as well as **improvements in the computational efficiency** of these methods.

Regarding the **contributions to the theory** of the LQR-tree methods, the dissertation presents a proof that a finite number of iterations are needed to cover the whole state space with funnels and their corresponding trajectories and LQR feedback regulators. This is indeed a result that was missing in the original works on LQR trees.

Regarding the **improved computational efficiency**, one contribution presented in the dissertation is that the simulation needed to verify reachability involves switching among several trajectories, which makes such simulation-based verification of reachability computationally less expensive. The reported reduction in computation time is significant. Another improvement in computation consists in combining the LQR-tree algorithm with the RRT algorithm, which provides better initial guesses for the trajectory optimizer, which in turns makes the trajectory optimization more robust/reliable and even faster. Finally, the dissertation also presents an optimization-based way to compute the funnels that turns out computationally cheaper than one of the original methods based on the sum of squares programming.

The claims about the novelty and correctness of the contribution have also been supported by **two**

papers presented at the two **flagship conferences** of the IEEE Robotics and Automation Society – ICRA and IROS. I can confirm that acceptance of papers at these two conference must certainly be well deserved. One of the two papers already got attention (citation) of a leading researcher in the field. The third paper is reportedly under review, but it has already been "published" at arXiv.org. Although I am not familiar with the rules at Faculty of Nuclear Sciences and Physical Engineering, I regard this publication outcome an adequate achievement for a Ph.D. student.

The text of the dissertation reveals the candidates' high writing culture.

Considering all these aspects, I do not hesitate to recommend the dissertation for the defence. It is indeed a piece of excellent research work on a topic of practical importance. Although the dissertation brings improvements to works published a decade or so ago, it still stands a chance to be appreciated (and cited) by the community. The more so if it is later accompanied by a publicly available software implementation.

After stating this official recommendation, I have a few minor comments. Most of them are just formal/terminological/notational. I am stating them here, although I am aware of the local/national rules that the dissertation author can no longer take benefits from the reviewers' suggestions by incorporating them in the dissertation after the final version is submitted. But perhaps they can be of some use anyway.

1. Concerning the computational superiority claimed and demonstrated in the text, I find it quite unfortunate that the candidate **has not shared his code** with the community. Software implementation of the proposed LQR-tree method(s) is nontrivial and with an abundance of alternative/competing methods, other researchers might not be willing to invest dozens of hours of coding only to check if the methods really work as described in the thesis (and papers). It was certainly my case. These concerns have been voiced in the computational communities by others as well, see the article Stodden, Victoria, Jonatham M. Borwein, and David H. Bailey. "'Setting the Default to Reproducible' in Computational Science Research." SIAM News, June 2013. `https://sinews.siam.org/Details-Page/setting-the-default-to-reproducible-in-computational-science-research`.
2. While I acknowledge the author's right to introduce his or her own terminology, I cannot help but note that encountering **non-traditional terminology** in situations for which the discipline already has an established terminology is tedious. The primary example in this dissertation is the very term **"demonstration"**. I have never encountered it in the literature on optimal control theory (and this dissertation does present contributions to the control theory). The term has not even been used in the papers by Russ Tedrake and his colleagues either. I am afraid that nonstandard terminology can be a barrier for linking the new results with the existing ones. And yet this is completely unnecessary, because the term "trajectory" is quite sufficient, isn't it? Possibly supplemented by the adjectives such as "optimal" or "reference".
3. Continuing to rant about the terminology, I find the usage of the term "direct" rather confusing when writing about optimal control approaches. In the dissertation it seems to be regarded as a synonymum to "online", but the established usage of **"direct" versus "indi-**

rect" has nothing to do with the online/offline distinction. One authoritative resource I can refer to is the section 4.3 named "Direct versus Indirect Methods" in the book Betts, John T. Practical Methods for Optimal Control and Estimation Using Nonlinear Programming. 2nd ed. Philadelphia: SIAM, 2010, which is cited in the dissertation. This is not just a nitpicking, because then in several places in the text it is really unclear, which parts of the computation must be done online and which can be done offline.

4. Similarly, the term "runtime": in the bottom paragraph at page 25 it is stated that some part of the design of the LQR-tree controller is done in runtime. And the Conclusion section contains statement about reduction in the runtime computation. It seems natural to interpret the term "runtime" as "online", that is, during the actualy operation of the controller. But is this really what is meant here?

5. I wouldn't call what is in (1), that is, the $\dot{x} = F(x, u)$, as control system. It is just a (model of a) dynamical system to be controlled. Sometimes it is called a **plant**. The **control system** is that what produces the control u. Alternatively, the control system is the whole thing, including the model, the control law, the reference signal, etc.

6. I find the introduction of the $\text{cost}_{V,t}(x_0, u)$ below (3) on page 12 as highly nonstandard in the optimal control literature (why not just stick to J or V as vast majority of papers and book do?). But even more importantly, it is unclear what is aimed to be achieved with this new notation compared to $V(x, u, t)$ defined in (3), both do the integration over the interval $[0, t]$. On the other hand, the optimal cost-to-go $V^*(x, t)$ presented subsequently in (6) is interprets the input argument t as the beginning of the interval $[t, T]$. This is unnecessarily unintuitive. And once again, here the not perfectly finetuned notation can and does contribute to the confusion.

7. It may be safer to in the sentence before (18) on page 17 to state that the (forward) Euler's method is just one of several possible methods. Direct transcription methods are not tied exclusively to the Euler's method.

8. In the first sentence of page 24, the requirement of controllability is not necessary. Stabilizability is sufficient.

9. I am struggling with the definition of the **funnel** on page 26. It is defined as a set of states (a subset of a state space), albeit time-varying, denoted as $\mathcal{F}(t)$. I get that it comes from the cited Tedrake's papers, but isn't more appropriately defined as a subset of $[t_0, t_f] \times \mathbb{R}^n$ as in the slightly later paper [126]? Doesn't it make more sense? In fact, this was one of major stumbling block for me when trying to get the idea of the funnel.

10. I am not sure I get the assumption 1 on page 29. It reads that it must be possible to compute the trajectory from any considered initial state to the target state (or target set of states). But is that really needed? My understanding of the original Tedrake's papers was that while the LQR-tree is growing, it is only necessary to target the trajectory to the existing (every growing) funnel, and no longer only to the target state. Is this understanding correct?

11. $\Delta x(T)$ in (33) must be a typo. No need for the Δ there.

12. On page 32 we can read that "we assume that the demonstrator somehow takes into account the cost functional (3), although the result does not necessarily have to be optimal." I do not

understand this statement at all.

13. Shouldn't the symbols $\in$ and $\notin$ be replaced by $\subset$ and $\not\subset$ in the middle of page 35?
14. I dare to recommend to consider replacing the dots in numbers such as $2 \cdot 10^5$ by 2×10^5 in order to avoid confusion with the decimal point as in $2.1 \cdot 10^5$.
15. In tables on page 64, tt would do no harm if the units are also included in the first row of the table. It does not appear that there was a shortage of space here.
16. On page 79, the "in this paper" should be replaced by "in this chapter".
17. It is not perfectly obvious how the task solved in Chapter 7 is related to the previous chapters. Is it correct to state that while the previous chapters were improving on the LQR-tree methods where funnels are approximated by sampling and falsification by simulation as in the paper [103], in this chapter the focus is on the more explicit characterization of funnels as in the original papers [123] and [124] and as such the chapter is rather standalone?

Sincerely,

Zdeněk Hurák