UNIVERSITY OF HELSINKI

Department of Physics
University of Helsinki
Gustaf Hällströmin katu 2
00560 Helsinki, Finland

Report of the PhD thesis

“Spectral gaps, t-designs and ϵ-nets in quantum computing”

by Oskar Szymon Słowik

Supervisor: Prof. dr. hab. Adam Sawicki
Co-supervisor: dr. Oliver Reardon-Smith

1. General Overview

The thesis of Oskar Słowik investigates the computational efficiency of universal quantum gate sets through the lens of spectral gaps, t-designs, and ϵ-nets on compact groups These are three tightly interlocking mathematical structures that sit at the boundary of quantum information theory, harmonic analysis, and Lie group theory. The central organizing question is: Given a universal set S of quantum gates, what is the shortest circuit (of depth ℓ) required to ϵ-approximate any target unitary operation, and how can one bound ℓ via quantities that are actually computable in practice?

The thesis is structured around three published or submitted articles (Papers I–III), each addressing a distinct but interrelated subproblem:

Paper I (J. Phys. A, 2023, with Sawicki) derives lower bounds on the finite-scale spectral gap $\mathrm{gap}_t(S)$ with explicitly calculable constants, using Bourgain’s argument combined with a version of the Solovay–Kitaev theorem with known constants. The main result (Theorem 1) establishes $\mathrm{gap}_t(S) \geq \alpha \cdot g_{t_0}(S) \cdot \log^{-2c}(\beta t)$ where all constants α, β, t_0 are in principle computable numerically, and $g_{t_0}(S)$ (defined in Eq. (137) of the paper) can be estimated via computer simulations.

Paper II (J. Phys. A, 2025, with Reardon-Smith and Sawicki) establishes an improved quantitative correspondence between δ-approximate unitary t-designs and ϵ-nets in $PU(d)$. The key innovation is the use of heat kernels on $SU(d)$, specifically the fundamental solutions of the heat equation, as natural polynomial approximate identities, in place of the periodized Gaussians used in prior work. The main result (Theorem 2) yields $t \simeq d^{5/2}/\epsilon$ and $\delta \simeq (\epsilon^{3/2}/d)^{d^2}$, representing a significant improvement in the δ scaling while preserving the same t scaling.

Paper III (arXiv:2505.00683, 2025, with Dulian and Sawicki) introduces the notion of Quantum Circuit Overhead (QCO) as a dimensionless measure comparing the efficiency of a gate set S to the optimal gate set of the same cardinality. Also the T-Quantum Circuit Overhead (T-QCO) is introduced, where the focus is on the occurrence of specific gates, which are assumed to be considerably more costly than the remaining gates (e.g., this is the case for many fault-tolerant implementations based on the Clifford+T gate set). The QCO can be upper bounded by the quantity $Q(S, \epsilon)$ defined by a formula involving only $\log(|S|)$ and the finite-scale spectral gap, making it numerically accessible. Extensive Monte Carlo simulations reveal, among other things, that the T gate is a markedly non-optimal completion of the single-qubit Clifford group, with $Q_T \approx 52$ versus an optimal value near 3.4.

2. Scientific Quality and Originality

2.1 Significance of the research program

The thesis focuses on deriving efficiency bounds for universal gate sets with explicit, computable constants. This is a problem where rigorous quantitative results have been notably scarce; thus, the thesis addresses a genuine and important gap in the literature. The Solovay–Kitaev theorem and its variants have long been understood at the level of asymptotic scaling, but the constants involved are either unknown or astronomically large, rendering the bounds practically useless for guiding compiler design or hardware choices. The thesis makes a principled and sustained effort to remedy this situation. This is not merely a technical improvement, it shifts the discourse from existence results to quantitative, numerically verifiable claims.

The choice of heat kernels as polynomial approximate identities in Paper II is natural and elegant. Heat kernels on compact Lie groups are classical objects with well-understood analytic properties, and their use simplifies and unifies the proof strategy compared to the ad hoc periodization constructions in earlier work. The improvement in δ scaling is a meaningful technical contribution, even if the improvement in t is limited.

The QCO framework introduced in Paper III provides a conceptually clean way to compare gate sets while normalizing for cardinality, a subtlety that is easy to overlook but that significantly affects practical conclusions. The application to T-gate efficiency, combined with numerical evidence suggesting that certain rotated $P(3\pi/4)$-type gates are substantially superior to (unrotated) $P(\pi/4)$ gates for completing the Clifford group, is a result of potential practical relevance.

2.2 Technical correctness and depth

The mathematical content of the thesis is handled with care. The use of Bourgain's argument in Paper I, the exploitation of the Poisson resummation for the $SU(d)$ heat kernel in Paper II, and the spectral-gap-based volumetric argument for QCO bounds in Paper III are all executed competently. The candidate demonstrates a solid command of harmonic analysis on compact groups, representation theory (Peter–Weyl theorem, Fourier calculus on groups), and the relevant quantum information formalism.

The numerical experiments, particularly those supporting Papers I and III, are substantial. The spectral gap calculations for Paper I consumed approximately two weeks on a 1008-core cluster. The Monte Carlo search over $\sim 10^4$ gate sets per ensemble in Paper III is similarly intensive. This distinguishes the work from purely asymptotic mathematical contributions and demonstrates that the candidate can bridge theory and computation.

The candidate is also to be credited for acknowledging the limitations of his results (e.g., the exponent $2c$ in the spectral gap decay bound of Paper I is non-optimal; the improvement over the prior state of the art in Paper II is primarily in the δ exponent, not in t). Moreover, the errata section for Paper I is also a positive sign: the candidate has carefully re-examined his own published work and identified three errors (two involving incorrect identification of representations and one involving misuse of a volume bound). The corrections appear technically sound.

3. Presentation

The thesis is clearly organized around its three papers, with each paper accompanied by a summary, contribution statement, and the full published text. Chapter 2 (Preliminaries) is comprehensive and serves its purpose well, providing a self-contained reference for the notational and conceptual framework used throughout. The level of detail is appropriate for a specialist reader and sufficient for a qualified non-specialist.

The writing quality is generally good. The English is fluent and technically precise. The introduction situates the work clearly within the broader context of quantum computing, avoiding the common failure of purely mathematical theses to motivate their topic for a physics or computer science audience. The research problems are explicitly formulated in Section 2.4.5, distinguishing addressed from unaddressed subproblems. This is good scientific practice and allows the reader to track exactly what the thesis claims to accomplish. The summary in Chapter 6 revisits these problems faithfully.

The bibliography is thorough and up to date, including Kuperberg's 2023 result breaking the cubic barrier in the Solovay–Kitaev algorithm ($c \approx 1.44$). It is notable that Kuperberg's breakthrough, which appeared during the thesis period, is properly acknowledged.

4. Potential questions and points for Discussion at the Defense

- During your PhD period Kuperberg's 2023 result breaking the cubic barrier in the Solovay–Kitaev algorithm ($c \approx 1.44$) appeared. Could this be used to improve some of the results?
- You use the finite-scale spectral gap $gap_t(S)$ rather than the global gap $gap(S)$. How do the two quantities compare in practice for the gate sets you simulate? Do the numerical computations suggest a rapid convergence to $gap(S)$, or is there significant difference even at the scales accessible to your large-scale numerical computations?
- A central assumption of the spectral gap programme is that every universal gate set has a gap. This is the famous conjecture of Sarnak, which is still unproven in general. All your quantitative bounds depend on the gap being positive, but you do not prove it for any new class of gate sets. How sensitive are the practical conclusions of the thesis to this unproven assumption, and what would change if a universal gate set were found to have zero gap?

5. Conclusion

The thesis of Oskar Słowik represents a coherent and technically solid contribution to the mathematical foundations of quantum gate efficiency. It addresses a well-motivated open problem (the derivation of quantitatively useful, computable bounds on circuit complexity) through three interconnected papers that demonstrate both analytical ability and comfort with large-scale numerical computation. The use of heat kernels as a unifying tool is conceptually satisfying, and the QCO framework provides a practically motivated organizational structure for comparing gate sets. The limitations are real but largely acknowledged, and most fall in the category of opportunities for future work rather than flaws in the existing results.

The candidate has demonstrated the ability to identify significant mathematical problems, develop original proof strategies, implement and analyze demanding numerical experiments, and communicate results clearly in peer-reviewed publications. The contributions are original and represent a genuine advance over the prior state of the art on the specific problems addressed.

Recommendation

The thesis meets the requirements for the degree of Doctor of Philosophy in Physics. The candidate is recommended to proceed to the public defense.

May 25, 2026

Zoltán Zimborás

Zoltán Zimborás
Professor of Quantum Algorithms
Department of Physics, University of Helsinki