Neural and Tensor Networks in the Study of Quantum Annealing Processors

Tomasz Śmierzchalski
Supervisor: dr hab. Bartłomiej Gardas
Co-supervisor: dr hab. inż. Łukasz Pawela

Abstract

Quantum annealing is a hardware realization of adiabatic quantum computation aimed at finding low-energy configurations of Ising/QUBO cost functions and is available in commercial devices such as D-Wave processors. Reliable assessment of these machines requires benchmarking that goes beyond "best solution" and explicitly accounts for strong classical baselines, sampling behavior, and diversity, and the physical cost of computation, including thermodynamic efficiency. This dissertation develops such a framework and applies it to contemporary quantum annealers.

The first contribution is SpinGlassPEPS.jl, a topology-aware tensor-network heuristic designed for optimization and sampling on quasi-two-dimensional graphs aligned with QPU connectivity (e.g., Pegasus/Zephyr). The method reduces problems to local Potts clusters, represents the resulting partition function with PEPS, and performs a branch-and-bound search in probability space, with explicit accuracy cost control via parameters such as inverse temperature and bond dimension. Building on this baseline, the thesis reports head-to-head benchmarks of D-Wave quantum annealers against classical and quantum-inspired solvers on random spin-glass instances defined natively on Pegasus and Zephyr graphs, reaching system sizes of several thousand spins and evaluating both optimization and sampling (including low-energy diversity) using like-for-like metrics. The results establish SpinGlassPEPS.jl as a physically interpretable reference, while also delineating its limitations: for the largest random instances, approximate tensor-network contractions introduce errors that prevent matching highly optimized GPU heuristics at comparable runtimes, positioning tensor-network methods primarily as topology-aware analysis and benchmarking tools for small-to-medium scales.

A second pillar is a thermodynamic characterization of quantum annealers as effective thermal machines, relating computational performance to energy dissipation, entropy production, and effective temperature along programmable schedules. This perspective shows that carefully chosen pauses can simultaneously improve solution quality/success probability and reduce thermodynamic

cost relative to simple reverse annealing, while also identifying regimes in which an applied longitudinal field becomes detrimental once pausing is introduced. The third pillar is a practical reinforcement-learning post-processing approach that treats returned samples as inputs to an RL agent, which learns spin-flip moves that correct typical error patterns. Once trained, it improves success probability and energy accuracy and can be integrated as an orthogonal mitigation layer. Finally, exact simulations of small-system adiabatic dynamics support intuition and modeling choices. Comparison to device data suggests that matching observed output statistics may require stronger effective disorder/noise than expected, indicating additional error sources or stronger environmental coupling in the relevant regimes.

Overall, the dissertation argues for benchmarking quantum annealing technologies using coherent protocols that combine rigorous classical comparators with metrics capturing both algorithmic performance and physical expenditure.