The Quantum Computing Company™

Reverse Quantum Annealing for Local Refinement of Solutions

WHITEPAPER

Summary

The success of classical heuristic search algorithms often depends on the balance between global search for good regions of the solution space (exploration) and local search that refines known good solutions (exploitation). While local refinement of known solutions is not available to the canonical forward quantum annealing algorithm, D-Wave has developed a *reverse annealing* feature that makes this possible by annealing backward from a specified state, then forward to a new state. This enables the use of quantum annealing for the refinement of classical states via local search, making it possible to use quantum annealing as a component in more sophisticated hybrid algorithms. Local quantum search has been analyzed theoretically to explore applications such as protein folding,[1] and has natural application in molecular dynamics, quantum simulation, and quantum chemistry, but has not been available for experiments until now. In a preliminary example, we show that reverse annealing can be used to generate new global optima up to 150 times faster than forward quantum annealing.

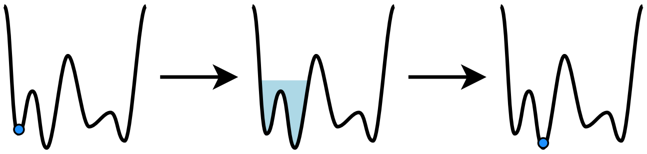

Figure 1: Sketch of the reverse annealing process. The anneal starts at a specified classical state (left), then performs a local quantum annealing search stimulated by an increased transverse field (middle), then settles in a new classical state as the transverse field is removed (right).

Successful classical heuristic search algorithms often combine a global search for good regions of the solution space with local refinement of good solutions that have been found—the complementary tasks of *exploration* and *exploitation*. Global search and local refinement are integral to the canonical forward quantum annealing algorithm, and the balance between them is controlled by the strength of the transverse field. Early in the anneal when the transverse field is strong, tunneling is frequent and the solution space is searched globally. Late in the anneal when the transverse field is weak, the algorithm's search narrows in on local regions of the solution space. A similar transition occurs in simulated classical annealing, wherein a high-temperature random global search early in the annealing algorithm gives way to low-temperature greedy local refinement later.

While forward quantum annealing takes advantage of global search and local refinement, it does so within a black box—nothing between the initial superposition and the final answer can be seen or controlled by the user due to the no-cloning theorem of quantum mechanics.[2] This means that forward quantum annealing cannot be used to perform local refinement of good states found elsewhere, i.e., from a previous anneal or through classical methods.

To address this limitation, the D-Wave 2000Q™ system now provides a *reverse annealing* feature that makes it possible to use quantum annealing for local refinement.

---

[1] A. Perdomo-Ortiz et al., Quantum Inf. Process. 10(1):33–52, 2011.

[2] W. K. Wootters and W. H. Zurek, Nature 299:802–803, 1982.

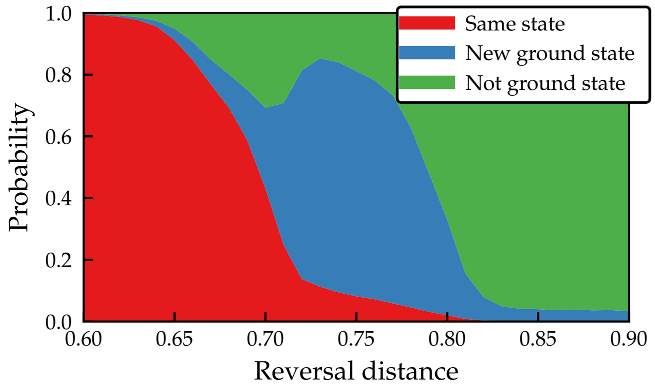

Figure 2: Transition types as a function of reversal distance. If the reversal distance is too low, the algorithm remains in the starting state; if it is too high, the algorithm fails to exploit knowledge of the initial state.

From a specified starting state, a local search is stimulated by increasing the transverse field, effectively annealing backward from a classical state to a mid-anneal quantum superposition; after allowing the search to run for a specified amount of time at this mid-anneal point, the quantum annealing algorithm proceeds forward by removing the transverse field to settle on a new classical output solution (see Figure 1).

Demonstration While the full potential of this novel feature lies in advanced algorithms still in development, we can demonstrate its utility using a relatively simple example. We consider the task of collecting many ground states (global optima) of a problem; we compare the rate at which new ground states are found using forward quantum annealing against the rate they are found using reverse annealing from known ground states. In this way we can determine whether knowledge of one ground state can be exploited in the search for a *new* ground state. The input we examine is a frustrated cluster loop problem[3] mined to have ground states that form distant clusters; within a cluster, ground states are separated by tall, thin energy barriers that make quantum annealing valuable even at a local level.

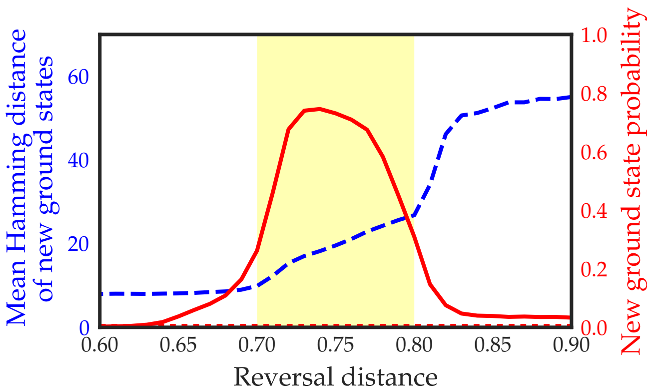

Figure 3: Probability of reaching a new ground state (red) and mean Hamming distance of new ground states from initial states (blue) as functions of reversal distance. Yellow shading is used to indicate the "Goldilocks" region in which we are reversing far enough to escape the initial ground state but not so far that we lose the ability to exploit knowledge of the initial state; the probability of arriving at a new ground state peaks in this region. The mean distance of new ground states from their starting states increases with reversal distance.

In reverse annealing, the breadth of the local search is controlled by the *reversal distance*, which specifies how far we anneal backward, and therefore how strong the transverse field becomes. If the reversal distance is too low, the quantum annealing search is likely to end up exactly where it started. If the reversal distance is too high, the algorithm effectively forgets where it started and is unable to exploit knowledge of the initial state (see Figure 2).

For the input we examined, standard forward annealing gave a ground state probability of 0.5% whereas reverse annealing from a known ground state found a new ground state up to 75% of the time (see Figure 3).

Conclusions Preliminary experiments in reverse quantum annealing found new ground states at 150 times the rate of forward quantum annealing. This shows great promise for reverse annealing as a component of more advanced hybrid algorithms. Two examples that have been proposed are quantum population annealing and quantum parallel tempering.[4] Until now, these algorithms have been hypothetical and only theoretical analysis has been possible. With the advent of the reverse annealing feature of the D-Wave 2000Q system, the algorithms are now open to implementation and experimental analysis. This evolution from canonical forward quantum annealing to advanced hybrid algorithms marks a significant step in the advancement of quantum computing.

---

[3] J. King et al., arXiv:1701.04579, 2017.

[4] N. Chancellor, arXiv:1606.06833, 2016.