

Improved coherence leads to gains in quantum annealing performance

WHITEPAPER

Summary

D-Wave has fabricated a series of quantum processing units (QPUs) possessing the same design but different materials within the QPU. In order to demonstrate the sensitivity of QPU performance to materials-related noise, a common set of identical spin glass problems, similar to those studied in a July 2018 *Science* article [4], has been posed to two such QPUs. The experimental results confirm a positive correlation between reduced noise and improved performance with at least a 25x speed up in solving spin glass problems having been observed.

Introduction

Spin glasses are considered the archetypal representation of a computationally challenging class of problems referred to as non-deterministic polynomial (NP)-hard [1]. Such problems can be extremely challenging to solve using existing classical digital computing technologies. There have been both theoretical [2] and experimental [3] indications that a quantum mechanical algorithm referred to as quantum annealing (QA) may provide a computational advantage in this case.

In [4], researchers at D-Wave studied the physics of phase transitions within quantum spin glasses on a cubic lattice. This study was important because it experimentally verified that quantum mechanics plays a significant role in D-Wave QPUs. Moreover, the presence of those phase transitions is a prerequisite to finding any possible computational advantage afforded by QA.

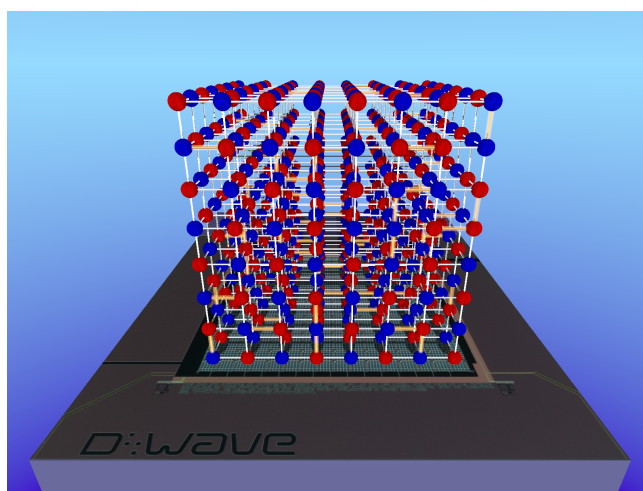


Figure 1: An illustration of one particular $8 \times 8 \times 8$ cubic lattice instance studied in [4]. Red and blue spheres represent the two possible states of the magnetic moments. Silver bars represent antiferromagnetic interactions that favor alternating (red-blue) ordering of the moments. Gold bars represent randomly added ferromagnetic interactions that favor uniform (blue-blue or red-red) ordering. These latter interactions serve to disorder the otherwise antiferromagnetic (alternating) ordering of the moments.

With the existence of the aforementioned phase transitions established, the next critical question is how well does such a system perform as a solver of spin glass problems? Furthermore, do improvements in quantum coherence translate into improvements in that performance? Observation of a positive correlation between coherence and performance would be an indicator that D-Wave's technology roadmap may provide an advantage over classical computation in the long run.

Experiments

D-Wave has maintained an ongoing effort to improve the quality of the materials used in its QPU fabrication process. Modifications and refinements to the process are carefully tested at the single-device level and only those changes that consistently result in reduced noise are then considered qualified for use in a full-scale QPU production process.

In order to determine whether there is a correlation between noise and performance, we have characterized two QPUs possessing the same design but manufactured using two different processes that will be referred to as FAB1 and FAB2 herein. Note that the FAB1 process bears the most resemblance to the process that was used to manufacture the commercially available DW2000Q QPUs to date. Noise was characterized by two standard metrics: macroscopic resonant tunneling (MRT) peak width and low frequency bias noise power spectral density. Performance was characterized using a set of 100 spin glass instances that were defined on a graph given by the intersection of the working physical qubits on each QPU and a perfect $8 \times 8 \times 8$ cubic lattice. Optimal annealing times were found for each instance [5], from which a measure of solution time was calculated.

Results

Within the context of QA as implemented by D-Wave, it has proved convenient to partition noise properties into two broad categories in relation to coherence: Low frequency noise that can be related to dephasing and high frequency noise that can be related to energy exchange (see [4]). The former can be characterized by measuring the time-dependent drift of biases on the individual physical qubits within a QPU [6]. The latter can be characterized by measuring time-dependent relaxation in the limit of small quantum fluctuations. This particular measurement is referred to as macroscopic resonant tunneling (MRT) [6].

Representative noise measurements and relevant best fit parameters for the two QPUs under consideration are summarized in Fig. 2 and Table 1, respectively. The low- and high-frequency noise amplitudes were

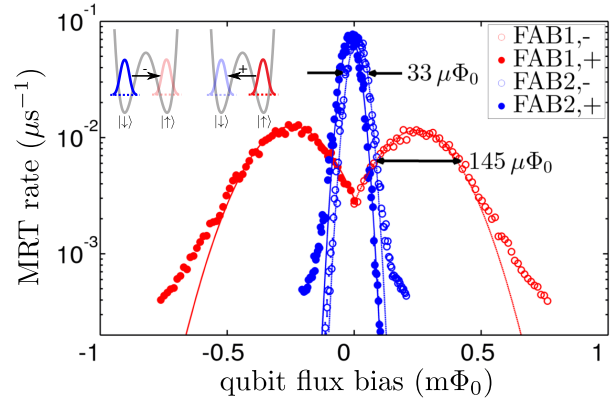


Figure 2: Macroscopic resonant tunneling (MRT) rate versus qubit flux bias for one representative qubit from each of the QPUs used in this study. The illustrations in the top left corner illustrate the physical picture of MRT wherein a qubit prepared in one spin state tunnels at some rate into the opposing state.

QPU	$1/f$ Amp. ($\mu\Phi_0/\sqrt{\text{Hz}}$)	MRT width ($\mu\Phi_0$)
FAB1	6.0 ± 0.5	145 ± 2
FAB2	3.5 ± 0.4	33 ± 1

Table 1: Summary of low- and high-frequency noise metrics for the QPUs used in this study.

roughly a factor of $2\times$ and $5\times$ lower, respectively, in the FAB2 QPU relative to the FAB1 QPU.

Having quantitatively established that the FAB2 process results in lower noise relative to the FAB1 process, the next task was to quantify the relative problem solving performance on cubic spin glass instances. One way to accomplish this task is to measure the probability P_{gs} of observing a ground state solution to a given instance as a function of the anneal time t_a required to complete the QA algorithm. One can then convert these two quantities into a nominal solution time t_s using the intuitive expression

$$t_s = \frac{t_a}{P_{\text{gs}}}. \quad (1)$$

For example, if one only observes a ground state once in 1000 reads ($P_{\text{gs}} = 1/1000$) at $t_a = 1 \mu\text{s}$, then one needs to run the QA algorithm about 1000 times at $t_a = 1 \mu\text{s}$ in order to observe at least one solution, hence the solution time $t_s = 1000 \times 1 \mu\text{s} = 1 \text{ms}$.

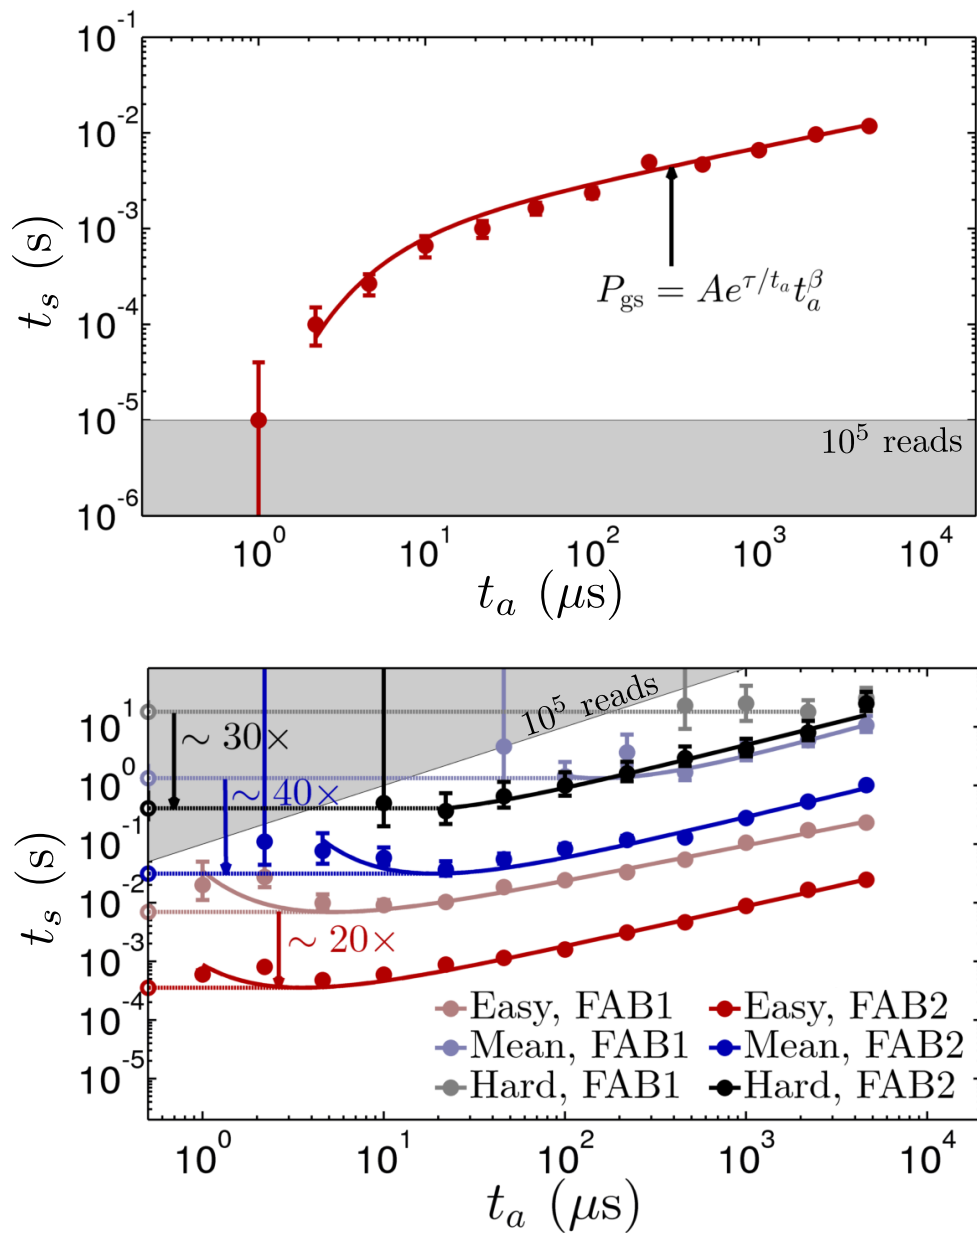


Figure 3: Comparison of QPU performance on cubic spin glass instances. (Top) Ground state (solution) probability P_{gs} versus QA time t_a for an example instance. Solid red curve corresponds to Eq. 2. (Bottom) Calculated solution time t_s versus t_a for representative easy, mean, and hard instances that were run on both the FAB1 QPU and the FAB2 QPU. The minimum value of t_s for each instance is indicated by a horizontal dashed line and point on the vertical axis. Relative reductions in the minimum value of t_s are indicated on the plot for each instance.

An example plot of P_{gs} versus t_a is shown for an instance run on the FAB1 QPU in the top panel of Fig. 3. The measurements indicate that P_{gs} monotonically grows with t_a . The behavior can be captured by the phenomenological form

$$P_{gs} = Ae^{-\tau/t_a}t_a^\beta, \quad (2)$$

which is similar to generalized spin glass relaxation formulae. Substituting Eq. 2 into Eq. 1 gives a functional form that can be used to fit traces of t_s versus t_a .

Example data and fits of t_s versus t_a for three representative instances are shown in the bottom panel of Fig. 3. For all but the representative hard instance run on the FAB1 QPU, the data can be fit to Eq. 2. From these fits and the lowest point for the aforementioned data trace, optimal anneal times and therefore minimum solution times can be identified. For all three instances shown, there is a systematic and well resolved factor of 20 reduction in minimum t_s in favor of the FAB2 process.

A correlation plot of the minimum t_s for all 100 instances run on both QPUs is shown in Fig. 4. For a typical randomly selected instance, the minimum t_s is a factor of 25× lower for the experimental fabrication process. Moreover, this trend holds true over nearly a factor of 10 000 in solution times, with easy instances located at low t_s and hard instances located at high t_s . There is some scatter in the correlation plot that can be bounded by factors of 5× (lower bound) and 100× (upper bound) on the speedup in favor of the experimental fabrication process. These results are very encouraging because they indicate that modest reductions in noise, as summarized in Table 1, lead to significant reductions in solution time.

Conclusions

The results presented herein provide strong validation for continued efforts to reduce noise in D-Wave QPUs. The experimental fabrication stacks examined in this study are but an intermediate step in an ongoing effort to improve D-Wave’s technology. We anticipate that additional and more substantial modifications currently being driven by detailed materials science investigations will imminently be qualified for full-scale manufacturing.

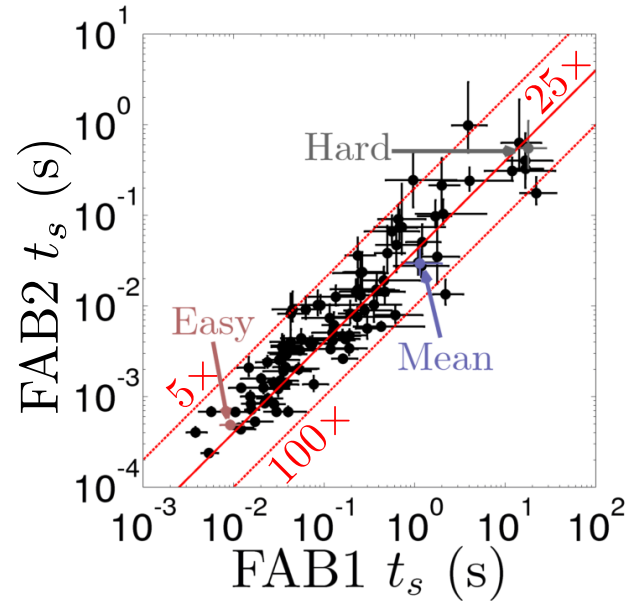


Figure 4: Comparison of QPU performance on 100 cubic spin glass instances. On average, the FAB2 t_s is 25× lower than the FAB1 t_s , and is bounded by lower and upper scale factors of 5× and 100×, respectively.

References

- 1 H. Nishimori, *Statistical physics of spin glasses and information processing: an introduction* (Oxford Science Publications, 2001).
- 2 E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, “A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem,” *Science* **292**, 472–475 (2001), eprint: <http://science.sciencemag.org/content/292/5516/472.full.pdf>.
- 3 J. Brooke, D. Bitko, F. T. Rosenbaum, and G. Aeppli, “Quantum annealing of a disordered magnet,” *Science* **284**, 779–781 (1999), eprint: <http://science.sciencemag.org/content/284/5415/779.full.pdf>.
- 4 R. Harris, Y. Sato, A. J. Berkley, M. Reis, F. Altomare, et al., “Phase transitions in a programmable quantum spin glass simulator,” *Science* **361**, 162–165 (2018), eprint: <https://science.sciencemag.org/content/361/6398/162.full.pdf>.
- 5 T. Albash and D. A. Lidar, “Demonstration of a scaling advantage for a quantum annealer over simulated annealing,” *Phys. Rev. X* **8**, 031016 (2018).
- 6 R. Harris, J. Johansson, A. J. Berkley, M. W. Johnson, T. Lanting, et al., “Experimental demonstration of a robust and scalable flux qubit,” *Phys. Rev. B* **81**, 134510 (2010).