



The D-Wave Advantage2 Prototype

TECHNICAL REPORT

Catherine McGeoch, Pau Farré, and Kelly Boothby

2022-06-09

Overview

This report describes two new key design features in the Advantage2 prototype: the Zephyr connection topology and an increased energy scale. This report also presents four empirical case studies in which the Advantage2 prototype outperformed a current-generation Advantage QPU.

CONTACT

Corporate Headquarters
3033 Beta Ave
Burnaby, BC V5G 4M9
Canada
Tel. 604-630-1428

US Office
2650 E Bayshore Rd
Palo Alto, CA 94303

Email: info@dwavesys.com

www.dwavesys.com

Notice and Disclaimer

D-Wave Systems Inc. (“D-Wave”) reserves its intellectual property rights in and to this document, any documents referenced herein, and its proprietary technology, including copyright, trademark rights, industrial design rights, and patent rights. D-Wave trademarks used herein include D-WAVE®, Leap™ quantum cloud service, Ocean™, Advantage™ quantum system, Advantage2™ quantum system, D-Wave 2000Q™, D-Wave 2X™, D-Wave Launch™, and the D-Wave logo (the “D-Wave Marks”). Other marks used in this document are the property of their respective owners. D-Wave does not grant any license, assignment, or other grant of interest in or to the copyright of this document or any referenced documents, the D-Wave Marks, any other marks used in this document, or any other intellectual property rights used or referred to herein, except as D-Wave may expressly provide in a written agreement.

Summary

Next-generation Advantage2 QPUs will incorporate several design enhancements, including: the Zephyr connection topology, which supports more compact embeddings having shorter chains; and increased energy scale, which lowers error rates due to thermal noise. We examine these new design features as realized on an Advantage2 prototype that is available online for public use starting in June 2022.

The prototype is small and contains approximately one-tenth as many qubits as current-generation Advantage QPUs. Empirical comparisons are challenging because inputs small enough to fit on the prototype do not leave much room for distinguishing performance: both quantum processors have little trouble finding optimal or near-optimal solutions (within a few percentage points of optimal) to the small inputs in our tests.

Nevertheless, we demonstrate that both new design features are effective at improving solution quality and increasing the probability of finding optimal solutions. The Advantage2 prototype outperformed the Advantage QPU in four case studies:

- **More compact embeddings.** For inputs small enough to fit on the prototype, chains in Zephyr embeddings on Advantage2 QPUs are between 6% and 21% shorter than chains in Pegasus embeddings on Advantage QPUs. An empirical comparison using two input classes embedded on the Advantage2 prototype and an Advantage (performance update) QPU showed improvements of up to 25% in mean chain length.
- **Better solutions on embedded inputs.** In tests using two categories of Satisfiability inputs, the Advantage2 prototype solver found better-quality solutions than the Advantage solver in up to 82% of cases.
- **Better solutions on native inputs.** Native inputs are used to isolate the contribution of increased energy scale (as distinct from embedding quality due to Zephyr). For three of four input categories, the Advantage2 prototype solver found better solutions in 77%, 87%, and 89% of cases. (Both solvers effectively tied on the fourth category and always returned optimal or second-best solutions: this case illustrates the challenges of distinguishing performance on small inputs.)
- **Reduced error rates due to increased energy scale.** Three component-level tests of *inverse effective temperature* were applied to qubits, couplers and small chains. These tests showed improvements by 94%, 103%, and 72% on the Advantage2 prototype, compared to the Advantage QPU.

These results demonstrate D-Wave's continuing commitment to discovering and deploying design innovations that improve performance of annealing quantum computers and systems.

Full-sized Advantage2 QPUs due for release in 2023-2024 will incorporate additional design and fabrication improvements not available in the prototype, as discussed in a companion white paper (*Lower-Noise Fabrication Development for the Advantage2 Quantum Computer*).

Contents

1	Introduction	1
2	Advantage2 and Its Prototype	3
3	Performance Comparison	6
3.1	Embedded Chain Lengths	7
3.2	Embedded NAE3SAT Inputs	8
3.3	Native RANd-r Inputs	10
3.4	Inverse Effective Temperature	12
4	Conclusions	13
	References	14

1 Introduction

D-Wave’s next-generation Advantage2 quantum processing units (QPUs) are due for release in 2023 or 2024; this report describes the Advantage2 prototype available in June 2022. The prototype joins several current-generation 2000Q and Advantage QPUs already online and available for public use.¹

Both the Advantage2 QPU and its prototype incorporate a new qubit design that features higher qubit connectivity. This is achieved via a connection topology, named Zephyr, that will replace the Pegasus topology found on Advantage QPUs. In Zephyr, qubits have maximum degree $d = 20$; that is, each qubit has up to 20 couplers connecting it to neighbors. The Pegasus connection topology has $d = 15$. Higher degree means that inputs can be minor-embedded more compactly onto the quantum chip and can have shorter chains.² Shorter chains in embedded inputs can lead to better-quality solutions. For more technical details about Zephyr see [1].

Furthermore, the new qubit design has an increased energy scale. This makes qubits and couplers less sensitive to thermal noise, thereby reducing error rates in the quantum computation. This improvement is especially effective on inputs that require extra precision in their specification.

This technical report presents an overview and small performance evaluation of these Advantage2 design features. Section 2 focuses on the Zephyr topology and explains the relationship between hardware connectivity and solution quality. Section 3 compares performance of the Advantage2 prototype solver³ and the Advantage (performance update) solver, in four case studies.

We remark that the small size of the Advantage2 prototype introduces some challenges to empirical performance evaluation. Significant performance differences can only be detected using large inputs, and tests using inputs small enough to fit on the prototype do not leave much room for differentiation. Despite this limitation, we find clear indications that the new design elements are effective at improving the quality of solutions returned by each solver.

Here are some highlights:

- Section 3.1 presents empirical results showing that chains in Zephyr embeddings are between 6% and 21% shorter than those in Pegasus embeddings, on prototype-sized inputs. A comparison of the Advantage2 prototype and an Advantage (performance update) QPU using two classes of NAE3SAT inputs shows that mean chain lengths can be up to 25% shorter in the prototype.
- Section 3.2 describes tests using two classes of NAE3SAT inputs, which focus on Zephyr and Pegasus embedding quality. Both solvers have little difficulty finding optimal solutions to most of the small inputs in our tests, so comparisons of “best

¹The prototype is small, containing about one-tenth as many qubits as Advantage QPUs. It is provided for demonstration purposes and is not intended to replace current quantum systems for use in application-relevant computations. It is not fully supported in the Ocean SDK.

²Minor embedding is a method for mapping arbitrary inputs onto fixed-connection hardware topologies.

³We use “solver” to refer to a specific QPU chip and its default control settings, together with control parameter ranges made available to users.

solution found” are meaningless. We employ a tiebreaker metric that considers overall distribution of solutions returned by the solvers: on the two classes tested, the prototype solver returned better results in 82% and 81% of inputs, while the Advantage solver prevailed in 0% and 3% of inputs (the remaining cases were ties). The prototype solver also found optimal solutions more frequently, on some inputs.

- Section 3.3 compares performance on four classes of Native inputs, which can be directly mapped to both Zephyr and Pegasus without chains. These inputs isolate the contribution of increased energy scale. In one test (as with NAE3SAT), inputs were too small and easy to distinguish performance: both solvers always found optimal or second-best solutions, and tied in 94% of cases. The other three classes do expose performance differences due to energy scale: the Advantage2 prototype solver found better-quality solutions in 77%, 89%, and 87% of cases, whereas the Advantage solver prevailed in just 13%, 10%, and 11% of cases (the rest were ties).
- Section 3.4 describes the *inverse effective temperature* metric β , which measures the sensitivity of low-level system components to thermal noise. In measurements of three such components — qubits, couplers, and short (2-qubit) chains — the Advantage2 prototype solver showed improvements of 94%, 103%, and 72% percent in mean β , compared to the Advantage QPU.

These results demonstrate D-Wave’s ongoing commitment to building annealing quantum computers that outperform their predecessors in tests using inputs and metrics that are relevant to applications practice.

The new Advantage2 design elements are expected to elicit similar or better performance when deployed on full sized Advantage2 QPUs in 2023 or 2024. Furthermore, Advantage2 QPUs will incorporate additional improvements — such as a new fabrication process described in [2] — that are not available the prototype.

We look forward to more thorough and comprehensive performance studies when full-sized next-generation Advantage2 QPUs become available.

Learn more about D-Wave products and services. The D-Wave Advantage (performance update) and Advantage2 prototype quantum systems described in this report are currently online, as are several other current- and previous-generation quantum systems. Together with the Leap quantum cloud service and Ocean software developer’s toolkit, they are available to the public (in North America, Europe, Japan, Australia, and India) for limited small-scale use at no cost. Larger blocks of QPU and system time are available for purchase from D-Wave or third-party providers.

The D-Wave hybrid solver service (HSS), available on a subscription basis, provides users with solutions to inputs for combinatorial optimization problems that are too large or too complex to fit onto current-generation QPUs. The suite of hybrid solvers incorporates an Advantage performance update system as their back-end quantum query server. Three hybrid solvers are available: BQM reads binary quadratic models; DQM reads discrete quadratic modules; and CQM reads constrained quadratic models defined on binary, integer, and continuous variables. These solvers can accept inputs containing thousands or millions of variables, depending on problem type. A D-Wave technical report [3] describes HSS properties and performance in more detail.

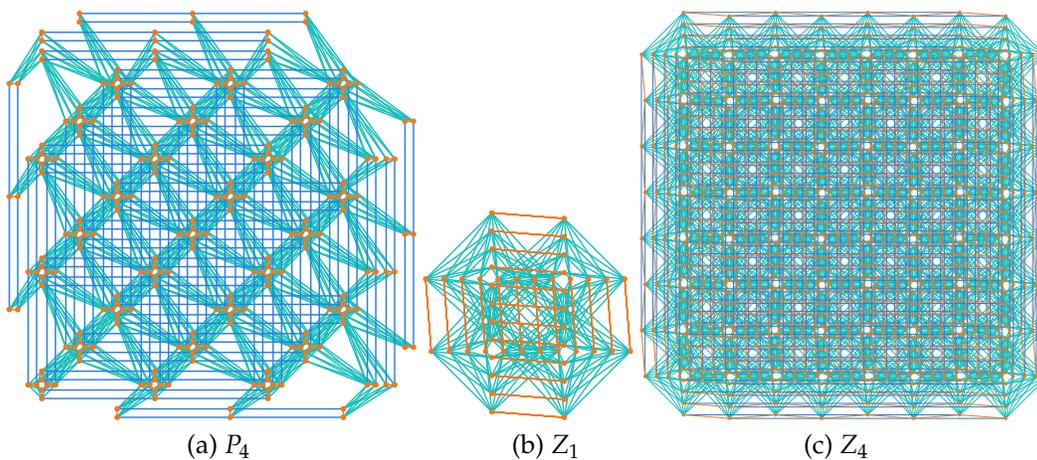


Figure 1: Comparison of Pegasus and Zephyr connection topologies. (a) A Pegasus P_4 graph contains 264 qubits. (b) A Zephyr Z_1 contains 32 qubits and four unit cells that share several couplers. (c) A Zephyr Z_4 graph contains 576 qubits.

Visit dwavesys.com to learn more about D-Wave quantum processors, the Leap and Ocean software stack, and HSS.

2 Advantage2 and Its Prototype

Two key design features that distinguish Advantage2 QPUs from previous-generation QPUs are the Zephyr hardware connection topology and an increased energy scale.

Figure 1 shows images of the Pegasus (Advantage) and Zephyr (Advantage2) topologies, which are constructed on square grids of so-called *unit cells*;⁴ See [1, 4] for technical descriptions of these two topologies. Panel (a) shows a Pegasus P_4 containing 264 qubits. Panel (b) shows a Zephyr Z_1 , which contains four unit cells and 32 qubits. Panel (c) shows a Zephyr Z_4 with 576 qubits — over twice as many as a P_4 — and with 33% more couplers per qubit. The higher qubit count per grid size, and the increased coupler density and complexity of Zephyr graphs are visibly apparent.

Table 1 compares some design features of current and future generations of D-Wave QPUs. As shown in Figure 1 topology size is based on grids containing one or more unit cells; the number of qubits per unit cell can vary. Zephyr has up to 20 couplers per qubit, which is 33 percent more than Pegasus and over three times as many as Chimera. The bottom rows show the minimum number of *active* qubits and couplers available in any product that is made publically available.⁵

The remainder of this section describes the relationship between hardware topology, chain length, and solution quality. A small empirical study compares embedded chain lengths

⁴Pegasus and Zephyr contain unit cells that overlap and share some qubits, and their definitions are not simple to explain. Formal definitions of terms like “number of qubits per unit cell” and “number of unit cells per grid size” may be found in the cited reports.

⁵In any specific QPU a small number of qubits and couplers may be disabled due to not meeting technical specifications.

Processor:	2000Q	Advantage	Advantage2 Prototype	Advantage2
Status	online	online	online	in development
Topology & grid size	Chimera C_{16}	Pegasus P_{16}	Zephyr Z_4	Zephyr Z_{15}
Qubits per unit cell	8	24	8	8
Couplers per qubit	6	15	20	20
Total qubits	> 2000	> 5000	563	> 7000
Total couplers	> 6000	> 35,000	4790	> 60,000

Table 1: Grid sizes, qubit counts, and coupler counts for recent and future D-Wave QPU designs.

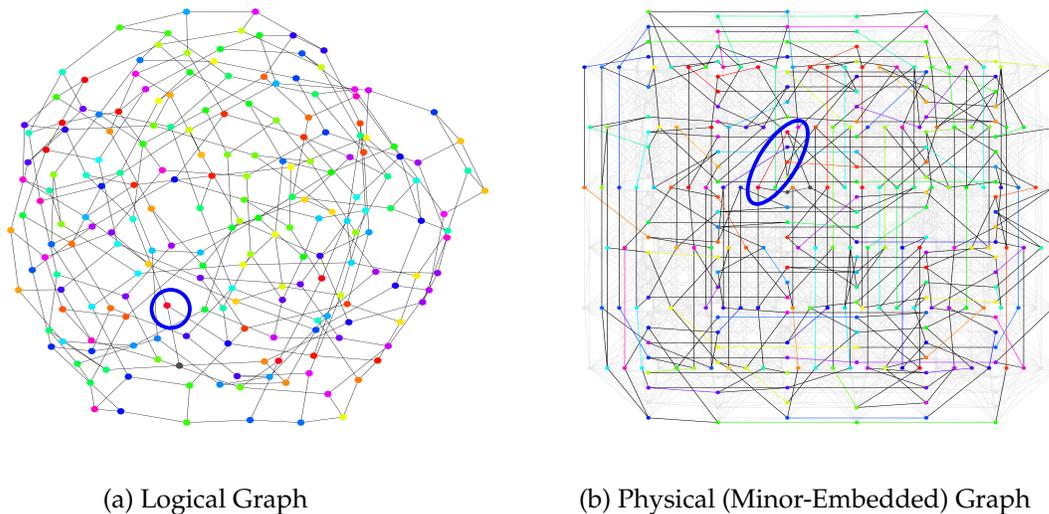


Figure 2: Minor-embedding of a general QUBO graph (a) onto a Z_3 graph (b). Each individual node in (a) maps to a chain in (b). Chain nodes and edges are colored to match their source nodes, and logical edges from the original graph are black. Qubits and edges of the Z_3 that are unused in this embedding are pale gray. The blue circle and oval highlight an example red node that is mapped to a red chain of two qubits.

for Zephyr and Pegasus topologies.

Embeddings and chains. Inputs for D-Wave solvers may be formulated as instances of the quadratic unconstrained binary optimization (QUBO) problem, or for the Ising Model problem (IM). An input of either type is represented by a *logical* graph G , as shown in Figure 2 (a).

In order for an input to be solved directly on a given QPU, the logical graph G must be mapped to a *physical graph* that matches the qubit and coupler connectivity structure, either Chimera, Pegasus, or Zephyr. This mapping is normally performed by the `minorminer` tool available in the Ocean SDK, which has modules for mapping general graphs, cliques (fully connected graphs), and graphs that have inherently spatial layouts. The mapping is called a *minor-embedding* (or informally, an *embedding*).

Figure 2 shows an example graph G and its minor-embedding onto a Zephyr Z_3 . Each

colored node in G corresponds to a *chain* of one or more same-colored qubits in the Z_3 , which are connected by same-colored chain edges.⁶ For example, the blue circle and oval highlight a red node in G that is mapped to a red two-qubit chain on the Z_3 . The logical edges of G are shown in black on the Z_3 ; qubits and couplers not used in the embedding are pale gray.

The higher connectivity of Zephyr graphs — with more qubits per cell and more couplers per qubit — means that graphs such as G can be minor embedded more compactly than is possible on Pegasus (and Chimera) graphs. Compactness is often measured by the maximum (or mean) chain length among all chains in the embedding. For example, Figure 2 (b) shows that most chains are of length $L \leq 2$, i.e., containing one or two qubits, although a green chain of length $L = 3$ may be found on the bottom edge of the Z_3 .

Compact embeddings are desirable for two reasons: they use fewer qubits per logical node, which means that larger and denser graphs can be represented with the same number of qubits; and they tend to have shorter chains, which can improve solution quality, by the following mechanism.

A QUBO input I is defined by real-valued weights $h = h_i$, $J = J_{ij}$ assigned to the nodes i and edges (i, j) of G . The task for the QPU is to find a good-quality solution to I : that is, an assignment of values $b \in \{0, 1\}$ to nodes x_i so as to minimize the solution *energy*, defined by h and J as follows:

$$E(x) = \sum_i h_i x_i + \sum_{(i,j)} J_{ij} x_i x_j. \quad (1)$$

For this to happen, the weights (h, J) in G must be mapped to the embedded graph, and a weight J_{chain} must be assigned to chain edges. We expect J_{chain} to be a large-magnitude negative weight to encourage chained qubits to return identical values (e.g., either 000 or 111 in a 3-qubit chain), which can be mapped back to their corresponding node in the original G .

The choice of chain strength $|J_{chain}|$ plays a role in solution quality, as follows.

- Setting chain strength too low compared to $|J_{ij}|$ increases the probability of finding broken chains in the output, and may even introduce spurious ground states that encourage broken chains. Postprocessing utilities in Ocean SDK can be used to repair broken chains, but repaired solutions tend to be of lower quality than intact solutions from the QPU. The probability of finding broken chains in output can be reduced by increasing the chain strength.
- On the other hand, setting chain strength too high introduces a different issue. Conceptually, when an input is sent to the QPU, all weights are scaled to a fixed physical energy range $[-1, +1]$ according to the largest-magnitude weight, assumed here to be J_{chain} : for example, every logical J_{ij} is scaled down to J_{ij} / J_{chain} in the physical input.⁷ Scaled weights that are too close to one another, or to zero, are not well-distinguished by the analog control system. Errors due to compressed problem scale can be mitigated by reducing the chain strength.

⁶Despite the name, chains are not always connected in a strict sequence; treelike chains with branches may also occur.

⁷This high-level discussion greatly simplifies real operations on current QPUs, which offer user parameters such as `extended_j_range` that offer increased energy scale for negative coupler weights.

Thus, the ideal chain strength lies in a sweet spot between two opposing hazards, creating a dilemma for the user: what is the ideal chain strength for my problem? Fortunately, Ocean utilities are available to suggest appropriate chain strengths based on input properties, and outputs from the QPU tend to be fairly robust with respect to small variations in J_{chain} .

More importantly, the two key features of the Advantage2 prototype QPU — Zephyr topology and increased energy scale — are designed to reduce errors at *both ends* of this spectrum. First, as shown in the next section, Zephyr embeddings tend to have shorter chains than Pegasus embeddings, and shorter chains are more resistant to breakage. Thus, smaller chain strengths can be used, which mitigates reduces compression of logical weights. Second, increased energy scale means that control signals can operate over a wider energy range, which reduces the impact of weight compression by making the computation more robust against thermal noise.

The combined effect is that embedded inputs can have lower chain strengths without an increased probability of chain breakage; and because of the increased energy scale, Advantage2 QPUs are more robust to thermal noise at any chain strength. Empirical results in the next section illustrate and quantify this observation.

3 Performance Comparison

In this section we compare two specific solvers that are currently online: the Advantage2 prototype solver contains a Zephyr Z_4 hardware graph with 563 active qubits, and the Advantage performance update solver contains a Pegasus P_{16} graph with 5627 active qubits.⁸

We report results from four separate case studies:

- Section 3.1 presents an empirical comparison of chain lengths in Zephyr and Pegasus embeddings, and in the Advantage2 prototype and Advantage QPU.
- Section 3.2 looks at solution quality using two categories of embedded NAE3SAT inputs, showing how more compact embeddings (plus increased energy scale) can yield better-quality solutions on embedded problems.
- Section 3.3 considers performance on four categories of so-called *native* inputs, which can be mapped directly onto QPU hardware without chains. Inputs without chains are used to isolate the effects of increased energy scale on full-sized inputs.
- Section 3.4 measures noise reduction in terms of *effective temperature*, a metric that is applied to individual components of the quantum chips, in this case qubits, couplers, and small (2-qubit) chains.

All four case studies demonstrate that the Advantage2 prototype solver finds solutions that are as good as or better than those found by the Advantage solver.

⁸The online names of these solvers are `Advantage2_prototype1.1` and `Advantage_system4.1`. Our tests report results from only these two solvers, which do not necessarily reflect outcomes that might be observed on other Advantage and 2000Q solvers.

3.1 Embedded Chain Lengths

A D-Wave technical report about Zephyr [1] compares chain lengths for a variety of graph categories when they are minor-embedded onto Zephyr, Pegasus and Chimera graphs of approximately equal sizes.

Table 2 below summarizes some results from that report. In this empirical test, a number of graphs from six different categories were minor-embedded 100 times using the heuristic `minorminer.find_embedding` utility, and means over all embeddings were recorded. The target graphs for these embeddings were a fully yielded (no missing qubits) Zephyr Z_4 with 576 qubits — matching the grid size of the Advantage2 prototype — and a Pegasus P_6 with 680 qubits, the nearest grid size.

Let L^P and L^Z denote statistics of chain length distributions in Pegasus and Zephyr embeddings, respectively. The columns show the percent improvements in mean and maximum chain lengths, equal to $P = (L^P - L^Z)/L^P$, for six different graph categories. The biggest improvement shown in this table is about 21 percent, for `random_gnp50` graphs. On average, mean and maximum chain lengths in Pegasus embeddings were reduced by about 15 percent and 18 percent in Zephyr embeddings. See [1] for details.

Input Class	Percent Improvement	
	Mean	Max
<code>small_graphs</code>	9.9	16.1
<code>random_growth</code>	6.5	19.8
<code>random_gnp75</code>	21.0	19.7
<code>random_gnp50</code>	21.4	21.7
<code>random_gnp25</code>	20.2	20.7
<code>random_cubic</code>	10.2	7.6
Average Improvement	14.9	17.6

Table 2: Heuristic embeddings of six input categories onto fully yielded Z_4 and P_6 graphs containing 576 and 680 nodes, respectively. Table entries show percent improvements, in mean and maximum chain lengths, of Z_4 over P_6 .

NAE3SAT embeddings. In addition to the above results, we compare mean chain lengths for two categories of Not-All-Equal 3-Satisfiability (NAE3SAT) inputs. These inputs are minor-embedded onto the Advantage2 prototype and Advantage QPUs used in performance comparisons of later sections.

These are Boolean satisfiability instances defined on n variables and formulated as logical conjunctions of m three-variable clauses, where each clause is satisfied when the three variables do not all have the same value (i.e., neither `ttt` nor `fff`). NAE3SAT is more convenient for quantum annealing benchmark tests than the more familiar CNF3SAT problem class, because instances can be translated directly to QUBO inputs with no increase in the number of variables needed.

We create two input classes parameterized by the clause-to-variable ratio $\rho = m/n$, which controls input hardness. Inputs generated with $\rho = 2.1$ are at the critical phase transition for NAE3SAT, and are interesting for studying performance at finding exact optimal solutions. Inputs generated with $\rho = 3.0$ are unsatisfiable with high probability, and are

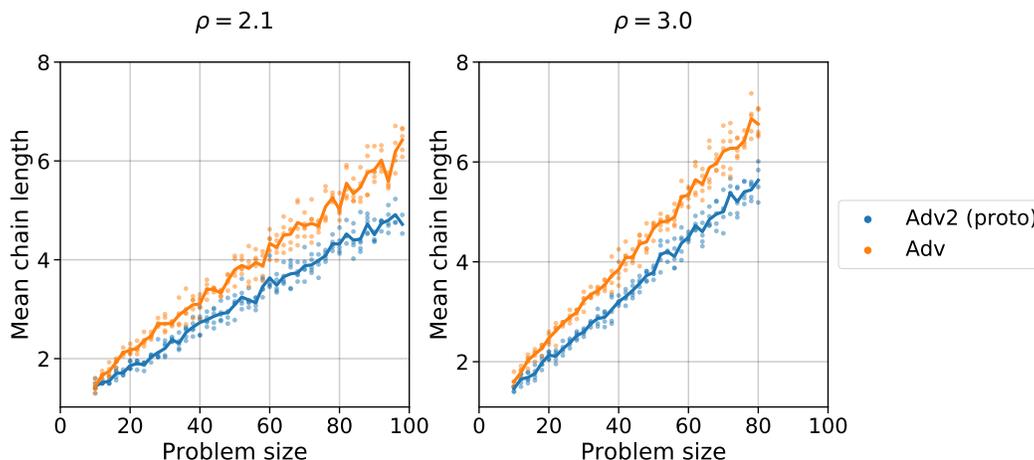


Figure 3: Heuristic embeddings of two classes of NAE3SAT inputs parameterized by $\rho = 2.1$ (left panel) and $\rho = 3.0$ (right panel). Each panel shows mean chain lengths versus input size, color-coded by solver: Advantage2 prototype (blue) and Advantage performance update (orange). Dots show mean chain lengths for five inputs at problem sizes $n = 10, 12, \dots, N_\rho$. Lines connect median points at each n .

interesting for studying approximation performance. For each ρ we generated five random inputs at each size $n = 10, 12, \dots, N_\rho$, where $N_{2.1} = 96$ and $N_{3.0} = 80$. Each input was embedded using the `minorminer.find_embedding` utility available in the Ocean SDK.

Figure 3 compares mean chain lengths in Advantage2 (blue) and Advantage (orange) embeddings, for $\rho = 2.1$ (left panel) and $\rho = 3.0$ (right panel). The dots show outcomes for each input at each size n ; the lines connect median points. At the largest n , mean chain lengths for embeddings on the Z_4 graph are about 25 and 20 percent shorter than embeddings on the full-sized P_{16} graph.

3.2 Embedded NAE3SAT Inputs

This section compares performance of the Advantage solver and the Advantage2 prototype solver at the task of finding good-quality solutions to embedded problems.

We use two sets of NAE3SAT inputs corresponding to $\rho = 2.1$ and $\rho = 3.0$, as described in the previous section. As it turns out, inputs that are small enough to fit on the Z_4 graph are quite easy to solve using both classical and quantum methods. Standard classical techniques were used to find optimal solutions to each instance s . The energy of an optimal solution, as determined by the objective function (1), is denoted S_s^* .

The two quantum solvers were tested using identical parameters: anneal time was $200\mu s$, number of reads was 1000, relative chain strength 3, and postprocessing rule `majority_vote`. For most inputs s , both solvers found optimal solutions at least once; that is, the sample *minimum* was nearly always equal to the optimal energy S_s^* . Solver performance could not be distinguished by comparing sample minimums, since most outcomes were ties. (Figure 6 (a) in the next section illustrates this phenomenon using a different input class.)

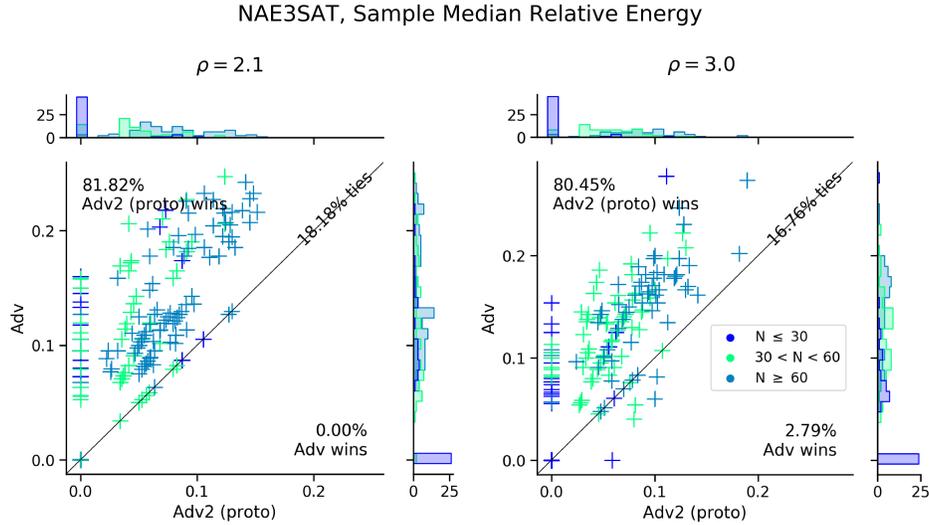


Figure 4: Median relative energies for NAE3SAT inputs with $\rho = 2.1$ (left panel) and $\rho = 3.0$ (right panel). Points are color-coded as small, medium, and large inputs. Points above the diagonal correspond to inputs on which the Advantage2 prototype solver returned strictly better solutions than the Advantage solver. Points on the diagonal line correspond to ties. Histogram plots on the sides of each panel show marginal distributions for each solver.

Instead, we measure *median* energy in each sample, denoted $S_s(A)$ for solver A , which serves as a tiebreaker. The *median relative energy* $R_s(A)$ is the scaled absolute difference between $S_s(A)$ and the optimal energy S^* ,

$$R_s(A) = \frac{|S_s^* - S_x(A)|}{|S_s^*|}. \quad (2)$$

Figure 4 presents input-by-input comparisons of $R_x(A)$ for each solver, for input sets $\rho = 2.1$ (left panel) and $\rho = 3.0$ (right panel). Points are color-coded by size (small, medium, large). A data point for input s has coordinates $x = R_s(\text{ADV2(proto)})$ and $y = R_s(\text{ADV})$. Points above the diagonal line correspond to cases where the Advantage2 prototype solver found strictly better solutions than the Advantage solver; points below the line correspond to cases where the Advantage solver won; and points on the line represent ties.

The dominance of the Advantage2 prototype solver is clear. In both tests, the prototype found better-quality solutions for 82% and 81% of inputs, while the Advantage solver found better solutions in just 0% and 3% of all 371 inputs tested.

Histograms on the sides of each panel show marginal distributions for each solver. The tops of both panels show a clear shift towards lower relative errors for the Advantage2 prototype solver. Also notable are the high concentrations of points at $(0,0)$, and arranged along the y -axis at $x = 0$. Points at $(0,0)$ correspond to inputs for which both solvers found optimal solutions in the median case. The latter points at $(0,y)$ correspond to inputs where the Advantage2 prototype found optimal solutions in sample medians. Such unusually high success probabilities $\pi \geq 0.5$ mean that on average, just one read would suffice to obtain an optimal solution.

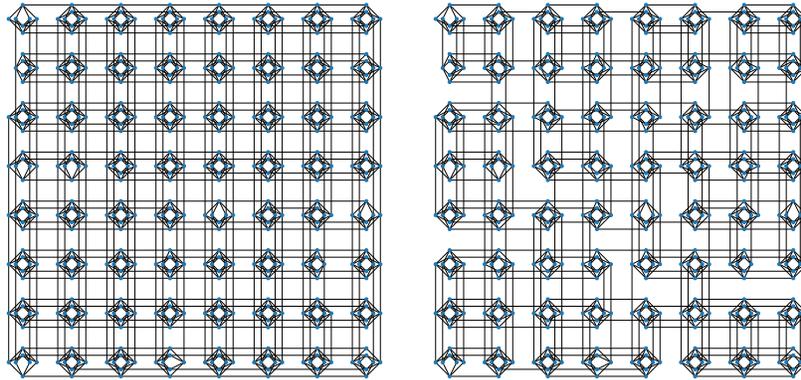


Figure 5: The left panel shows a Chimera C_8 graph with maximum degree $d = 6$; the right panel shows the same graph with edges removed to obtain degree $d = 5$. These graphs can be mapped to both the Pegasus hardware on the Advantage QPU and the Zephyr hardware on the Advantage2 prototype. Missing nodes and edges in the C_8 graphs correspond to inactive qubits and couplers in the intersection of both hardware graphs.

3.3 Native RANd-r Inputs

A native input can be mapped directly onto QPU hardware without the use of chains. Figure 5 (left) shows two examples based on the Chimera graph topology. These example graphs map directly onto the Pegasus and Zephyr hardware graphs inside the Advantage QPU and Advantage2 prototype, without the use of chains. The left panel shows a Chimera C_8 with 503 nodes and maximum degree $d = 6$; the right panel shows the same graph with edges removed to obtain maximum degree $d = 5$. A small number of missing nodes and edges in both graphs map to inactive components in the intersection of the Z_4 and P_{16} hardware graphs of their respective QPUs.

A RANd-r graph based on these Chimera structures has degree $d \in \{5, 6\}$, node weights $h_i = 0$, and edge weights J_{ij} generated uniformly at random from the set $\{\pm 1, \pm 2, \dots, \pm r\}$. Native inputs, having no chains, are interesting because they isolate the contribution of increased energy scale on total performance, as distinct from contributions of embedding quality. Our tests use the single RAN6 graph shown in the figure, and a set of randomly generated RAN5 graphs, of which one example is shown.

In this section we compare performance on four Native classes: RAN5-1, RAN6-10, RAN6-100, and RAN6-1000.⁹ We generated 151 random inputs from each problem class. These inputs are too large for any exact algorithm to find guaranteed optimal solutions in a reasonable amount time; therefore for each input s we use a *reference* solution energy S_x^* (the pooled minimum of these and other solvers not discussed here) as the putative optimal solution.

Both quantum solvers were run with sample size 1001, using three anneal times: $t_a \in$

⁹RAN6-1 graphs are omitted because they have unusual properties that elicit unrepresentative performance profiles from quantum annealing solvers [5]. In a small pilot test on these inputs, both solvers found optimal and near-optimal solutions (fourth-best or better in the solution space) on all inputs, although the sample distribution for the Advantage2 prototype was somewhat more skewed towards the high end.

Sample Minimum Relative Energy

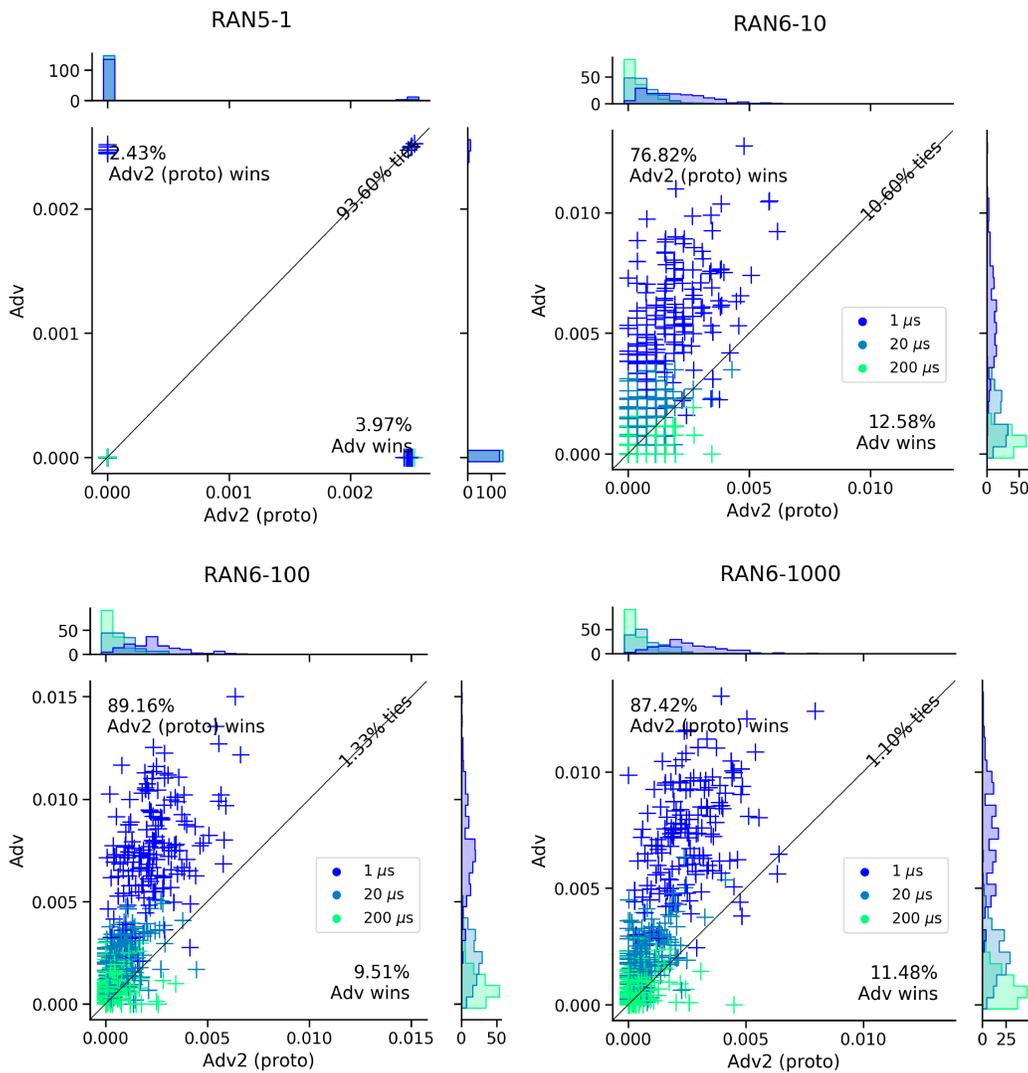


Figure 6: The four panels show results using input from four classes: RAN1-5, RAN6-10, RAN6-100, RAN6-1000. Points are color-coded according to anneal times $t_a \in \{1, 20, 200\} \mu s$. Points above the diagonal are cases where the Advantage2 prototype solver returned strictly better results than the Advantage solver. Points on the diagonal are ties. Histograms on the sides of each panel show marginal distributions for each solver.

$\{1, 20, 200\}$ microseconds. For each solver A and input s , let $S_s(A)$ denote the *minimum* energy found in a sample of 1001 solutions, and let $R_s(A)$ denote the scaled relative error as defined in (2).

Figure 3.3 shows the results for all four input classes. In each panel, for each input s , three different-colored points are shown, located at $x(t_a) = R_s(\text{Adv2}(\text{proto}))$, $y(t_a) = R_s(\text{Adv})$, corresponding to tests using three different anneal times: $t_a = 1\mu\text{s}$ (dark blue), $t_a = 20\mu\text{s}$ (light blue), and $t_a = 200\mu\text{s}$ (green). Points above the diagonal are cases where the Advantage2 prototype solver found strictly better solutions; points on the diagonal correspond to ties. Histograms on the sides show marginal distributions for each solver.

As with NAE3SAT inputs, both solvers cope well with these relatively small inputs. The top left panel shows results for RAN5-1 inputs, on which both solvers tied in 93.6 percent of cases. Over all inputs and anneal times tested here, the sample minimums returned by both solvers matched the reference solution (best) or the second-best solution, as indicated by the concentration of points in the four corners of the panel. The histograms show that both solvers found reference solutions in a large majority of inputs. The remaining differences can be attributed to random sampling error; for these reasons we conclude that these inputs are simply too small and easy to differentiate solver performance.

In contrast to the NAE3SAT results, the larger weight ranges in RAN6-10, RAN6-100, and RAN6-1000 inputs provide fine enough granularity in the distribution of near-optimal solution energies to allow comparisons of sample minimums. Here are some observations:

- The win percentages and locations of point clouds are quite similar across all three panels. This suggests that for both solvers, relative solution energy — that is, solution quality as a proportion of the putative optimal energy — is independent of r , which controls the precision of weights used to define the input.

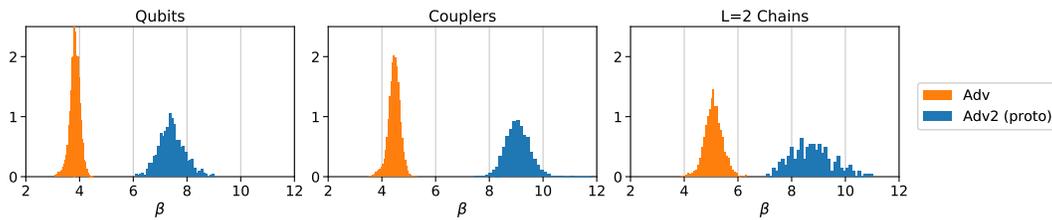
Relative energy for the Advantage2 prototype solver never rise above .008 in all three panels, whereas it can be as high as 0.15 for the Advantage solver. That is, in terms of the approximation ratio $P_s(A) = (1 - R_s(A)) * 100$, the prototype solver always returned solution energies within 99.2 percent of (putative) optimal solutions, whereas the Advantage solver always found solution energies within 98.5 percent.

- Comparisons of the histograms on the top and side of each panel show that the outcome $R_s(A) = 0$ — corresponding to finding the putative optimal solution in the sample minimum — is the most common outcome for both solvers. Both solvers find better results at longer anneal times; the much larger spread of results for the Advantage solver suggests that its performance shows much more variability at the smallest anneal times.

The increased energy range incorporated in future full-sized Advantage2 QPUs is expected to be particularly beneficial on inputs requiring higher precision weights in their specifications, such the ones studied here.

3.4 Inverse Effective Temperature

The *effective temperature* T in an annealing QPU is not a physical temperature, but rather a measure of the impact of thermal noise on its quantum components. Effective temperature



Mean β			
	Advantage	Advantage2	Improvement
Qubit	3.828	7.430	94.1%
Coupler	4.460	9.049	102.9%
Chains L=2	5.094	8.749	71.8%

Figure 7: Single Qubit Temperature. Mean effective inverse temperatures β , for qubits, couplers, and chains in each QPU (higher is better). Beta correlates with the effective suppression of errors due to thermal noise, and with chain strength. The Advantage2 prototype shows improvements in all three categories.

of a single qubit x is calculated from the distribution of outcomes (spin up or spin down) observed over a range of field values h_x . Other low-level system components can similarly be measured using this approach. This quantity is commonly reported in terms of the *inverse effective temperature* $\beta = 1/T$ (lower T and higher β are better). See [6] for technical information about how this quantity is measured.

The histograms in Figure 7 shows measurements of β for three components — qubits (left), couplers (center), and small ($L = 2$) chains (right) — on the Advantage QPU and the Advantage2 prototype. The histograms are normalized by size so that the total area under each curve equals one.

The table in Figure 7 shows mean values of β from each test: this measurement on the Advantage2 prototype is approximately twice that observed on the Advantage QPU, in two cases. This dramatic reduction in sensitivity to thermal noise can be directly attributed to the increased energy scale incorporated into the new qubit design.

4 Conclusions

Next-generation Advantage2 QPUs will incorporate two important design enhancements that are considered in this report: the Zephyr connection topology, which supports more compact embeddings having shorter chains; and increased energy scale, which lowers error rates due to thermal noise.

We examine these new design features as they are realized on an Advantage2 prototype that is available online for public use. The prototype is small, containing about one-tenth as many qubits as current-generation Advantage QPUs.

Empirical comparison of the prototype to current Advantage systems is challenging, because inputs small enough to fit on the prototype do not leave much room for performance

variation: in most cases both QPUs have little trouble finding optimal solutions, or solutions within just a few percentage points of optimal.

Despite this size limitation, we find clear signals that both new design elements are effective at improving the quality of solutions returned, and at increasing the probability of finding optimal solutions in outputs. The Advantage2 prototype shows better performance than the Advantage QPU in four case studies using both embedded and native inputs, as well as in measurements of individual quantum components (qubits, couplers, and short chains).

These results demonstrate D-Wave's continued commitment to discovering and deploying design innovations that improve performance of our annealing quantum computers and systems.

Full-sized Advantage2 QPUs due for release in 2023 or 2024 will incorporate additional design enhancements not available in the prototype, such as a new fabrication process described in [2]. We look forward to more thorough and comprehensive performance studies when Advantage2 QPUs become available.

References

- ¹ K. Boothby, A. D. King, and J. Raymond, *Zephyr Topology of D-Wave Quantum Processors*, 14-1056A-A (D-Wave Technical Report, 2021).
- ² (D-Wave), *Lower Noise Fabrication Development for the Advantage2 Quantum Computer* (D-Wave White Paper (in preparation), 2022).
- ³ (D-Wave), *Hybrid Solver for Constrained Quadratic Models*, 14-1055A-A (D-Wave White Paper, 2019).
- ⁴ K. Boothby, P. Bunyk, J. Raymond, and A. Roy, "Next-generation topology of D-Wave quantum processors," arXiv [quant-ph] [2003.00133](https://arxiv.org/abs/2003.00133) (2020).
- ⁵ A. D. King, E. Hoskinson, T. Lanting, E. Andriyash, and M. H Amin, "Degeneracy, degree, and heavy tails in quantum annealing," *Physical Review A* **93** (2016).
- ⁶ J. Raymond, S. Yarkoni, and E. Andriyas, "Global warming: Temperature estimation in annealers," *Frontiers in ICT* **3** (2016).