



D-Wave Solver Properties and Parameters Reference

USER MANUAL

2018-11-27

Overview

This document describes the solvers available in the D-Wave system, listing their properties and the parameters they accept with a problem submission.

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), its subsidiaries and affiliates, makes commercially reasonable efforts to ensure that the information in this document is accurate and up to date, but errors may occur. NONE OF D-WAVE SYSTEMS INC., its subsidiaries and affiliates, OR ANY OF ITS RESPECTIVE DIRECTORS, EMPLOYEES, AGENTS, OR OTHER REPRESENTATIVES WILL BE LIABLE FOR DAMAGES, CLAIMS, EXPENSES OR OTHER COSTS (INCLUDING WITHOUT LIMITATION LEGAL FEES) ARISING OUT OF OR IN CONNECTION WITH THE USE OF THIS DOCUMENT OR ANY INFORMATION CONTAINED OR REFERRED TO IN IT. THIS IS A COMPREHENSIVE LIMITATION OF LIABILITY THAT APPLIES TO ALL DAMAGES OF ANY KIND, INCLUDING (WITHOUT LIMITATION) COMPENSATORY, DIRECT, INDIRECT, EXEMPLARY, PUNITIVE AND CONSEQUENTIAL DAMAGES, LOSS OF PROGRAMS OR DATA, INCOME OR PROFIT, LOSS OR DAMAGE TO PROPERTY, AND CLAIMS OF THIRD PARTIES.

D-Wave reserves the right to alter this document and other referenced documents without notice from time to time and at its sole discretion. D-Wave reserves its intellectual property rights in and to this document and its proprietary technology, including copyright, trademark rights, industrial design rights, and patent rights. D-Wave trademarks used herein include D-Wave®, D-Wave 2X™, D-Wave 2000Q™, Leap™, and the D-Wave logos (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, 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. This document may refer to other documents, including documents subject to the rights of third parties. Nothing in this document constitutes a grant by D-Wave of any license, assignment, or any other interest in the copyright or other intellectual property rights of such other documents. Any use of such other documents is subject to the rights of D-Wave and/or any applicable third parties in those documents.

All installation, service, support, and maintenance of and for the D-Wave System must be performed by qualified factory-trained D-Wave personnel. Do not move, repair, alter, modify, or change the D-Wave System. If the equipment is used in a manner not specified by D-Wave, the protection provided by the equipment may be impaired. Do not provide access to the customer site to anyone other than authorized and qualified personnel. Failure to follow these guidelines may result in disruption of service, extended downtime, damage to equipment (customer's, D-Wave's, and/or third parties'), injury, loss of life, or loss of property.

Contents

1	About this Document	1
2	QPU (Hardware) Solvers	1
2.1	Properties	1
	anneal_offset_step	1
	anneal_offset_step_phi0	1
	anneal_offset_ranges	1
	annealing_time_range	2
	chip_id	2
	couplers	2
	default_annealing_time	2
	default_programming_thermalization	2
	default_readout_thermalization	2
	extended_j_range	2
	h_range	3
	j_range	3
	max_anneal_schedule_points	3
	num_reads_range	3
	num_qubits	3
	parameters	3
	per_qubit_coupling_range	4
	problem_run_duration_range	4
	programming_thermalization_range	4
	qubits	4
	readout_thermalization_range	4
	supported_problem_types	5
	vfy	5
2.2	Parameters	5
	anneal_offsets	5
	anneal_schedule	5
	annealing_time	6
	answer_mode	6
	auto_scale	7
	beta	7
	chains	7
	flux_biases	8
	flux_drift_compensation	8
	h	8
	initial_state	9
	J	9
	max_answers	10
	num_reads	10
	num_spin_reversal_transforms	11
	postprocess	11
	programming_thermalization	12
	Q	12

qpu_access_overhead_time	13
readout_thermalization	13
reduce_intersample_correlation	14
reinitialize_state	14
reverse_anneal	14
2.3 Answer Format	15
3 QPU-Like Solvers	16
3.1 VFYC Solvers	16
3.2 Optimizing and Sampling Emulators	16
Optimizing Emulators	16
Sampling Emulators	17
4 Ising Heuristic Solvers	18
4.1 Algorithm	18
4.2 Solver Name	19
4.3 Properties	19
4.4 Parameters	19
iteration_limit	19
local_stuck_limit	20
max_bit_flip_prob	20
max_local_complexity	20
min_bit_flip_prob	20
num_perturbed_copies	20
num_variables	21
random_seed	21
time_limit_seconds	21
4.5 Summary of Ising Heuristic Solver Parameters	21
5 Quick Reference Tables	22
5.1 Properties and Parameters Summary by Solver Type	22
5.2 QPU Solver Parameters Matrix	24

1 About this Document

When you submit a problem to the D-Wave system, you do so via a *solver*. Solvers accept certain parameters that affect how the problem is run. This document lists the properties and parameters of:

- QPU (hardware) solvers
- QPU-like solvers, including the VFYC solver, optimizing emulators, and sampling emulators
- Ising heuristic solvers

Note: Not all accounts have access to all solver types.

2 QPU (Hardware) Solvers

These solvers submit problems to and return results from the D-Wave QPU. This section lists the properties of these solvers, lists the parameters they accept with a problem submission, and describes the answer format.

For a quick reference listing all parameters and how they interact with each other, see the [QPU Solver Parameters Matrix](#) section.

2.1 Properties

This section describes the properties of QPU solvers, in alphabetical order.

`anneal_offset_step`

Quantization step size of anneal offset values in normalized units.

`anneal_offset_step_phi0`

Quantization step size in physical units (annealing flux-bias units): Φ_0 .

`anneal_offset_ranges`

Array of ranges of valid anneal offset values, in normalized offset units, for each qubit. The negative values represent the largest number of normalized offset units by which a qubit's anneal path may be delayed. The positive values represent the largest number of normalized offset units by which a qubit's anneal path may be advanced.

Check these ranges before using the [`anneal_offsets`](#) parameter.

annealing_time_range

Range of time, in microseconds, possible for one anneal (read). The lower limit in this range is the fastest quench possible for this solver. When adjusting the anneal schedule, using either the *annealing_time* or *anneal_schedule* parameter, ensure that you do not exceed the limits in this range.

chip_id

Name of the solver.

couplers

Couplers in the working graph. A coupler contains two elements $[q1, q2]$, where both $q1$ and $q2$ appear in the list of working qubits, in the range $[0, num_qubits - 1]$ and in ascending order (i.e., $q1 < q2$). These are the couplers that can be programmed with nonzero J values; for example, $[[0, 4], [1, 4], [2, 4], \dots]$

default_annealing_time

Default time, in microseconds, for one anneal (read). You can change the annealing time for a given problem by using the *annealing_time* or *anneal_schedule* parameters, but do not exceed the upper limit given by the *annealing_time_range* property.

default_programming_thermalization

Default time, in microseconds, that the system waits after programming the QPU for it to return to base temperature. This value contributes to the total *qpu_programming_time*, which is returned by SAPI with the problem solutions.

You can change this value using the *programming_thermalization* parameter, but be aware that values lower than the default accelerate solving at the expense of solution quality.

default_readout_thermalization

Default time, in microseconds, that the system waits after each state is read from the QPU for it to cool back to base temperature. This value contributes to the *qpu_delay_time_per_sample* field, which is returned by SAPI with the problem solutions.

extended_j_range

Extended range of values possible for the coupling strengths (quadratic coefficients), J , for this solver. Strong negative couplings may be necessary for some embeddings;

however, such chains may require additional calibration through the *flux_biases* parameter to compensate for biases introduced by strong negative couplings. See also *per_qubit_coupling_range*.

Note: Flux-bias offsets and extended J ranges cannot be used with autoscaling or spin-reversal transforms.

h_range

Range of values possible for the qubit biases (linear coefficients), h , for this solver. The *auto_scale* parameter enables you to submit problems with values outside these ranges and have the system automatically scale them to fit.

j_range

Range of values possible for the coupling strengths (quadratic coefficients), J , for this solver. The *auto_scale* parameter enables you to submit problems with values outside these ranges and have the system automatically scale them to fit. See also *extended_j_range*.

max_anneal_schedule_points

Maximum number of points permitted in a PWL waveform submitted to change the default anneal schedule. Check this value before defining a new schedule with the *anneal_schedule* parameter.

For reverse annealing, the maximum number of points allowed is one more than the number given in the *max_anneal_schedule_points* property.

num_reads_range

Range of values possible for the number of reads that you can request for a problem; for example, $[1, 1000]$.

num_qubits

Total number of qubits, both working and nonworking, in the QPU; for example, 2048.

parameters

List of the parameters supported for the solver.

per_qubit_coupling_range

Coupling range permitted per qubit for this solver. Check this property when using an extended J range to strongly couple qubits in a chain. Strong negative couplings may be necessary for some embeddings; however, chains may require additional calibration through the *flux_biases* parameter to compensate for biases introduced by strong negative couplings. See also *extended_j_range*.

Note: Flux-bias offsets and extended J ranges cannot be used with autoscaling or spin-reversal transforms.

problem_run_duration_range

Range of time, in microseconds, that a problem can run.

The upper limit of this range is calculated according to the following formula:

$$Duration = ((P_1 + P_2) * P_3) + P_4 \quad (1)$$

where P_1 , P_2 , P_3 , and P_4 are the values specified for the *annealing_time*, *readout_thermalization*, *num_reads* (samples), and *programming_thermalization* parameters, respectively.

programming_thermalization_range

Range of time, in microseconds, possible for the system to wait after programming the QPU for it to cool back to base temperature. This value contributes to the total *qpu_programming_time*, which is returned by SAPI with the problem solutions.

You can change this value using the *programming_thermalization* parameter, but be aware that values lower than the default accelerate solving at the expense of solution quality. The default value for a solver is given in the *default_programming_thermalization* property.

qubits

Indices of the working qubits in the working graph. For example, $[0, 1, 2, 3, \dots]$

Note: In a D-Wave QPU, the set of qubits and couplers that are available for computation is known as the *working graph*. The yield of a working graph is typically less than the total number of qubits and couplers that are fabricated and physically present in the QPU.

readout_thermalization_range

Range of time, in microseconds, possible for the system to wait after each state is read from the QPU for it to cool back to base temperature. This value contributes to the *qpu_delay_time_per_sample* field, which is returned by SAPI with the problem solutions.

Note: The `readout_thermalization` parameter is deprecated and will eventually be removed from the API. Plan code updates accordingly.

supported_problem_types

Indicates what problem types are supported for the solver. QPU and QPU-like solvers support the following energy-minimization problem types:

- `qubo`—Quadratic unconstrained binary optimization (QUBO) problems; use 0/1-valued variables.
- `ising`—Ising model problems; use $-1/1$ -valued variables.

vfyc

Flag indicating whether this solver is a VFYC solver.

2.2 Parameters

This section describes the parameters accepted by QPU solvers, in alphabetical order. See *QPU Solver Parameters Matrix* for a summary.

anneal_offsets

Provides offsets to annealing paths, per qubit.

Provide an array of anneal offset values, in normalized offset units, for all qubits, working or not. Use 0 for no offset. Negative values produce a negative offset (qubits are annealed *after* the standard annealing trajectory); positive values produce a positive offset (qubits are annealed *before* the standard trajectory). Before using this parameter, query the solver properties to see whether the `anneal_offset_ranges` property exists and, if so, to obtain the permitted offset values per qubit.

Format varies by SAPI client:

- For C, this parameter is a pointer to the `sapi_DoubleArray` structure.
- For MATLAB, supply the values as a vector.
- For Python, supply the values as a list or a tuple.

anneal_schedule

Introduces variations to the global anneal schedule. For a reverse anneal, use the `anneal_schedule` parameter with the `initial_state` and `reinitialize_state` parameters.

An anneal schedule variation is defined by a series of pairs of floating-point numbers identifying points in the schedule at which to change slope. The first element in the pair is

time t in microseconds; the second, normalized persistent current s in the range $[0,1]$. The resulting schedule is the piecewise-linear curve that connects the provided points.

The following rules apply to the set of anneal schedule points provided:

- Time t must increase for all points in the schedule.
- For forward annealing, the first point must be $(0,0)$ and the anneal fraction s must increase monotonically.
- For reverse annealing, the anneal fraction s must start and end at $s = 1$.
- In the final point, anneal fraction s must equal 1 and time t must not exceed the maximum value in the *annealing_time_range* property.
- The number of points must be ≥ 2 .
The upper bound is system-dependent; check the *max_anneal_schedule_points* property. For reverse annealing, the maximum number of points allowed is one *more* than the number given by this property.
- Additional rules that govern maximum slope vary by product; check the QPU properties document for your system.

The *max_anneal_schedule_points* property shows the maximum number of points permitted in an anneal schedule, and *default_annealing_time* shows the annealing time range for the solver.

Format varies by SAPI client:

- For C, this parameter is a pointer to the `sapi_AnnealSchedule` structure.
- For MATLAB, provide the schedule as a cell array of two-element matrices; for example: `{[0.0, 0.0], [10.0, 0.3], [20.0, 1.0]}`
- For Python, provide the schedule as a list or tuple of two-element lists or tuples; for example: `[[0.0, 0.0], [10.0, 0.3], [20.0, 1.0]]`

Note: The *annealing_time* and *anneal_schedule* parameters are mutually exclusive.

annealing_time

Sets the duration (in microseconds) of quantum annealing time, per read. This value populates the *qpu_anneal_time_per_sample* field returned in the *timing* structure.

Must be a positive integer falling within the range given by the *annealing_time_range* solver property.

Note: The *annealing_time* and *anneal_schedule* parameters are mutually exclusive.

answer_mode

Indicates how to return answers from the solver; one of:

- `histogram` (default)—Answers returned as a histogram sorted in order of increasing energies. Duplicate answers are merged and include a `num_occurrences` field.
- `raw`—Answers returned individually in the order they were read from the solver. Use this setting if the returned time sequences are an important part of the solution data.

See also the *Answer Format* section below.

auto_scale

Indicates whether h and J values are rescaled:

- `true`— h and J values in the problem are rescaled to use as much of the range of h (`h_range`) and the range of J (`j_range`) as possible. When enabled, h and J values need not lie within the solver's range of h and J , but must still be finite.
- `false`— h and J values in the problem are used as is. If the h and J values are outside the range of the solver, problem submission fails.

Each QPU has an allowed range of values for the biases and strengths of qubits and couplers. Unless you explicitly disable auto-scaling, the D-Wave software adjusts the values defined in a problem to take the entire range available before sending it to the solver.

For problems that use the regular J range of a solver, this parameter is enabled by default. For problems that use the extended J range, it is disabled (and cannot be enabled while the extended range is in use). See also `j_range` and `extended_j_range`.

Note: Flux-bias offsets and extended J ranges cannot be used with autoscaling or spin-reversal transforms.

beta

Provides a value for the Boltzmann distribution parameter. Used when sampling postprocessing is enabled.

Note: As in statistical mechanics, β represents inverse temperature: $1/(k_B T)$, where T is the thermodynamic temperature in kelvin and k_B is Boltzmann's constant. In the D-Wave software, postprocessing refines the returned solutions to target a Boltzmann distribution characterized by β , which is represented by a floating point number without units. When choosing a value for β , be aware that lower values result in samples less constrained to the lowest energy states. For more information on β and how it is used in the sampling postprocessing algorithm, see *Postprocessing Methods on D-Wave Systems*.

chains

Defines which qubits represent the same logical variable. Used only when postprocessing is enabled. Ensures that all qubits in the same chain have the same value within each sample.

Provide the indices of the qubits that are in a chain. Qubits in a chain must be connected, and no qubit may appear more than once. Format varies by SAPI client:

- For C and Python, provide the qubit indices as a list of lists. The inner lists are the indices of the qubits in the same chain; for example, [4, 12, 20, 28, 36, 44]. An example with two parallel horizontal 4-qubit chains is [[4, 12, 20, 28], [5, 13, 21, 29]].
- For MATLAB, provide the qubit indices as a cell array.

flux_biases

List of flux-bias offset values with which to calibrate a chain. Often required when using the extended J range to create a strongly coupled chain for certain embeddings. See also the *extended_j_range* and *per_qubit_coupling_range* properties.

Provide an array of flux-bias offset values, in normalized offset units¹, for all qubits, working or not. Use 0 for no offset. Before using this parameter, query the solver properties to determine whether *extended_j_range* and *per_qubit_coupling_range* properties exist. Format varies by SAPI client:

- For C, this parameter is a pointer to the `sapi_DoubleArray` structure.
- For MATLAB, supply the values as a vector.
- For Python, supply the values as a list or a tuple.

Note: Flux-bias offsets and extended J ranges cannot be used with autoscaling or spin-reversal transforms.

flux_drift_compensation

Boolean flag indicating whether the D-Wave system compensates for flux drift. The procedure it follows to do so is described in detail in Appendix A of *Technical Description of the D-Wave Quantum Processing Unit*.

- `flux_drift_compensation=true` (default)—Compensate for flux drift.
- `flux_drift_compensation=false`—Do not compensate for flux drift. If you disable this parameter, you may want to apply flux-bias offsets manually; see *flux_biases*.

h

For Ising problems, the h values are the linear coefficients (biases) applied to the qubits in a problem.

Because the quantum annealing process² minimizes the energy function of the Hamiltonian, and h is multiplied with the answer, returned states tend toward the opposite sign as that provided as a bias. For example, if you bias a qubit with an h value of -1 , it is more likely to have a final state of $+1$ in the solution.

¹ Flux-biases applied to the qubit body are in units of Φ_0 .

² QPU-like solvers emulate this process.

A problem definition comprises both h and J values. For example, the following is a 7-qubit problem in which q_1 and q_4 are biased with -1 and $q_0, q_2, q_3, q_4,$ and q_6 are biased with $+1$. Qubits q_0 and q_6 have a ferromagnetic coupling between them with a strength of -1 .

```
h = [1, -1, 1, 1, -1, 1, 1]
J = {(0, 6): -1}
```

To see the supported h range for a solver, check its `h_range` property. The `auto_scale` parameter enables you to submit problems with values outside `h_range` and `j_range` and have the system automatically scale them to fit.

Note: For QPU solvers, the programming cycle programs the solver to the specified h and J values for a given Ising problem (or Q values for a given QUBO problem). However, since QPU precision is limited, the h and J values (or Q values) realized on the solver may deviate slightly from the requested values. For more information, see *Technical Description of the D-Wave Quantum Processing Unit*.

Note: For QUBO problems, use **Q** instead of h and J ; see [Q](#) for more information and instructions on converting between the formulations.

initial_state

When using the reverse annealing feature, you must supply the initial state to which the system is set. When you include this parameter with the `anneal_schedule` parameter, this indicates that the requested anneal schedule change is also a reverse anneal.

The `initial_state` parameter specifies the initial classical state of all qubits. Provide (*qubit, state*) pairs, where *qubit* is the qubit index, i , and *state* is:

- -1 or 1 —Ising problems, active qubits
- 0 or 1 —QUBO problems, active qubits
- 3 —Unused or inactive qubits

Format varies by SAPI client; for all, length equals the total number of qubits:

- For C, provide an array using the `sapi_ReverseAnneal` structure.
- For MATLAB, provide the dense set of (*qubit, state*) pairs as a vector.
- For Python, provide the dense set of (*qubit, state*) pairs as a list or tuple.

You have the option to reassert this state for each anneal, although this impacts timing; see also `reinitialize_state`.

J

For Ising problems, the J values are the quadratic coefficients applied to the couplers in a problem. This parameter defines the type and the strength of the coupling between a pair of qubits:

- **J < 0:** Ferromagnetic coupling; coupled qubits prefer to be in the same state, (1, 1) or (-1, -1).
- **J > 0:** Antiferromagnetic coupling; coupled qubits prefer to be in opposite states, (-1, 1) or (1, -1).
- **J = 0:** No coupling; qubit states do not affect each other.

The larger the absolute value of J , the stronger the coupling.

A problem definition comprises both h and J values. The following is a 7-qubit problem with a coupling strength of -1 between q_0 and q_6 :

```
h = [1, -1, 1, 1, -1, 1, 1]
J = {(0, 6): -1}
```

To see the supported J range for a solver, check its `j_range` property. The `auto_scale` parameter enables you to submit problems with values outside `h_range` and `j_range` range and have the system automatically scale them to fit. Some solvers support an extended J range for stronger couplings; however, be aware that you cannot use the extended J range with autoscaling or spin-reversal transforms. See [extended_j_range](#) for more information.

Note: For QPU solvers, the programming cycle programs the solver to the specified h and J values for a given Ising problem (or Q values for a given QUBO problem). However, since QPU precision is limited, the h and J values (or Q values) realized on the solver may deviate slightly from the requested values. For more information, see *Technical Description of the D-Wave Quantum Processing Unit*.

Note: For QUBO problems, use Q instead of h and J ; see [Q](#) for more information and instructions on converting between the formulations.

max_answers

Specifies the maximum number of answers returned from the solver:

- If `answer_mode` is `histogram`—Total number of distinct answers.
- If `answer_mode` is `raw`—Limits the returned values to the first `max_answers` of `num_reads` samples. In this mode, `max_answers` should never be more than `num_reads`.

Must be an integer > 0; default is `num_reads`.

num_reads

Indicates the number of states (output solutions) to read³ from the solver. Must be a positive integer; default is 1.

³ Terms synonymous to `reads` are `anneals` and `samples`.

num_spin_reversal_transforms

Specifies the number of spin-reversal transforms⁴ to perform.

Use this parameter to specify how many spin-reversal transforms to perform on the problem. Valid values range from 0 (do not transform the problem; the default value) to a value equal to but no larger than the *num-reads* specified. If you specify a nonzero value, the system divides the number of reads by the number of spin-reversal transforms to determine how many reads to take for each transform. For example, if the number of reads is 10 and the number of transforms is 2, then 5 reads use the first transform and 5 use the second.

Applying a spin-reversal transform can improve results by reducing the impact of analog errors that may exist on the QPU. This technique works as follows: Given an n -variable Ising problem, we can select a random $g \in \{\pm 1\}^n$ and transform the problem via $h_i \mapsto h_i g_i$ and $J_{ij} \mapsto J_{ij} g_i g_j$. A spin-reversal transform does not alter the mathematical nature of the Ising problem. Solutions s of the original problem and s' of the transformed problem are related by $s'_i = s_i g_i$ and have identical energies. However, the sample statistics can be affected by the spin-reversal transform because the QPU is a physical object with asymmetries.

Spin-reversal transforms work correctly with postprocessing and chains. Majority voting happens on the original problem state, not on the transformed state.

Be aware that each transform reprograms the QPU; therefore, using more than 1 transform will increase the amount of time required to solve the problem. For more information about timing, see *Measuring Computation Time on D-Wave Systems*.

Note: Flux-bias offsets and extended J ranges cannot be used with autoscaling or spin-reversal transforms.

postprocess

Defines what type of postprocessing to run on raw solutions:

- "" (empty string)—No postprocessing (default). Note that if this option is selected for the VFYC solver, sampling postprocessing runs.
- `sampling`—Runs a short Markov-chain Monte Carlo (MCMC) algorithm with single bit-flips starting from each sample. The target probability distribution is a Boltzmann distribution at inverse temperature β .
- `optimization`—Performs a local search on each sample, stopping at a local minimum.

Note: For problems that use the VFYC solver, postprocessing always runs. By default, sampling postprocessing runs using a default β value of 10.

When postprocessing is enabled, qubits in the same chain, defined by the *chains* parameter, are first set to the same value by majority vote. Postprocessing is performed on the logical

⁴ Also known as *gauge transformations*.

problem but qubit-level answers are returned. For more information about postprocessing, see *Postprocessing Methods on D-Wave Systems*.

programming_thermalization

Gives the time (in microseconds) to wait after programming the QPU for it to cool back to base temperature (i.e., post-programming thermalization time). Lower values accelerate solving at the expense of solution quality. Must be an integer. This value contributes to the total *qpu_programming_time*, which is returned by SAPI with the problem solutions.

The default value for a solver is given in the *default_programming_thermalization* property; the range of possible values is given in the *programming_thermalization_range* property.

Q

A quadratic unconstrained binary operation (QUBO) problem is defined using an upper-triangular matrix, \mathbf{Q} , which is an $N \times N$ matrix of real coefficients, and \mathbf{x} , a vector of binary variables. The diagonal entries of \mathbf{Q} are the linear coefficients that bias the qubits (analogous to h , in Ising problems). The nonzero off-diagonal terms are the quadratic coefficients that define the strength of the coupling between variables (analogous to J , in Ising problems).

Input may be full or sparse. Both upper- and lower-triangular values can be used; (i, j) and (j, i) entries are added together. If a Q value is assigned to a coupler not present, an exception is raised. Only entries indexed by working couplers may be nonzero.

Format depends on SAPI client:

- C—See the `sapi_Problem` structure in *Developer Guide for C*.
- MATLAB—Matrix of QUBO coefficients.
- Python—Dictionary of QUBO coefficients.

QUBO problems are converted to Ising format before they run on the solver. Ising format uses h (qubit bias) and J (coupling strength) to represent the problem; see also h and J .

Before sending a problem, check the *h_range* and *j_range* properties for the ranges permitted for the solver. Be aware that problems with values outside the supported range will, by default, be scaled down to fit within the supported range; see the *auto_scale* parameter for more information. Some solvers support an extended J range for stronger couplings; however, be aware that you cannot use the extended J range with autoscaling or spin-reversal transforms. See *extended_j_range* for more information.

Note: For QPU solvers, the programming cycle programs the solver to the specified h and J values for a given Ising problem (or Q values for a given QUBO problem). However, since QPU precision is limited, the h and J values (or Q values) realized on the solver may deviate slightly from the requested values. For more information, see *Technical Description of the D-Wave Quantum Processing Unit*.

Problem Conversion

Conversion between QUBO and Ising problem formulations is as follows:

$$s = q2 - 1. \quad (2)$$

To convert the coefficients from QUBO to Ising:

$$J_{i,j} = \frac{1}{4} Q_{i,j}, \quad (3)$$

$$h_i = \frac{1}{2} Q_{i,i} + \frac{1}{4} \sum_{i < j} Q_{i,j}, \quad (4)$$

or from Ising to QUBO:

$$Q_{i,j} = 4J_{i,j} \quad (5)$$

$$Q_{i,i} = 2h_i - \frac{1}{2} \sum_{i < j} Q_{i,j}. \quad (6)$$

For example, this two-qubit problem in QUBO format,

```
Q = (0,5): -10,
```

is as follows in Ising format:

```
h = [-2.5, 0, 0, 0, 0, -2.5], J = {(0, 5): -2.5}.
```

The energies for the corresponding solutions have the constant difference of

$$\frac{1}{2} \sum_i Q_{i,i} + \frac{1}{4} \sum_{i < j} Q_{i,j}. \quad (7)$$

qpu_access_overhead_time

A small amount of initialization time spent in low-level operations during the processing of a problem; roughly 1 ms. This value is including in the timing information returned by the Solver API for a completed problem.

readout_thermalization

Gives the time (in microseconds) to wait after each state is read from the QPU for it to cool back to base temperature (i.e., post-readout thermalization time). This value contributes to the `qpu_delay_time_per_sample` field, which is returned by SAPI with the problem solutions. Must be an integer.

Note: The `readout_thermalization` parameter is deprecated and will eventually be removed from the API. Plan code updates accordingly.

reduce_intersample_correlation

Reduces sample-to-sample correlations caused by the spin-bath polarization effect⁵ by adding a delay between reads.

Boolean flag indicating whether the system adds a delay.

- `reduce_intersample_correlation=true`—Adds delay.
- `reduce_intersample_correlation=false` (default)—Does not add delay.

Important: Enabling this parameter drastically increases problem run times. To avoid exceeding the maximum problem run time configured for your system, limit the number of reads when using this feature. For more information on timing, see *Measuring Computation Time on D-Wave Systems*.

reinitialize_state

When using the reverse annealing feature, you must supply the initial state to which the system is set; see the *initial_state* parameter. If multiple reads are requested in a single call to the Solver API, you have two options for the starting state of the system. These are controlled by the *reinitialize_state* Boolean parameter:

- `reinitialize_state=true` (default)—Reinitialize the initial state for every anneal-readout cycle. Each anneal begins from the state given in the *initial_state* parameter. The amount of time required to reinitialize varies by system; typical D-Wave 2000Q systems require between 100 and 600 microseconds for this operation.
- `reinitialize_state=false`—Initialize only at the beginning, before the first anneal cycle. Each anneal (after the first) is initialized from the final state of the qubits after the previous cycle. Be aware that even this parameter is disabled, reverse annealing adds a small amount of time ($\approx 10 \mu\text{s}$).

For the C client, this parameter is part of the `sapi_ReverseAnneal` structure.

See also *anneal_schedule*.

reverse_anneal

Relevant only to the C client. Pointer to the `sapi_ReverseAnneal` structure, which defines the initial state of the system at the start of a reverse anneal.

⁵ See the *Technical Description of the D-Wave Quantum Processing Unit* for more information on this effect.

2.3 Answer Format

The format of answers returned by QPU solvers depends on the setting of the *answer_mode* parameter:

- If *raw*—The answer contains two fields, *solutions* and *energies*. The *solutions* field is a list of lists; the inner lists all have length *num_qubits* and entries from $-1, +1$ for Ising problems or $0, 1$ for QUBO problems. Values of 3 denote used or inactive qubits. The *energies* field contains the energy of each corresponding solution.
- If *histogram*—The answer still contains the *solutions* and *energies* fields, but in this case the solutions are unique and sorted in increasing-energy order. An additional field, *num_occurrences*, indicates how many times each solution appeared.

3 QPU-Like Solvers

This section describes the “QPU-like” solvers, including virtual full-yield Chimera (VFYC) solvers, optimizing emulators, and sampling emulators. These solvers have similar properties and accept similar parameters as the QPU solvers. Differences are described below.

Note: Not all accounts have access to these solvers.

3.1 VFYC Solvers

A VFYC solver emulates a fully connected Chimera graph based on an idealized abstraction of the system. Through this solver, variables corresponding to a Chimera-structured graph that are not representable on a specific working graph are determined via hybrid use of the QPU and the integrated postprocessing system, which effectively fills in any missing qubits and couplers that may affect the QPU. For more information on the VFYC solver and how it is integrated with the postprocessing system, see *Postprocessing Methods on D-Wave Systems*.

The properties of a VFYC solver and the parameters it accepts are the same as for the QPU solvers described above, with the following differences:

- For problems that use the VFYC solver, postprocessing always runs. If you do not choose a postprocessing option, sampling postprocessing runs.
- When sampling postprocessing runs, the default β value is 10.

3.2 Optimizing and Sampling Emulators

Software solvers are useful for prototyping optimization or sampling algorithms that make multiple calls to the hardware. Optimizing emulator solvers solve the same type of optimization problems as the QPU, but using a classical software algorithm. Sampling emulator solvers mimic the probabilistic nature of the QPU to draw samples that can be used to train a probabilistic model.

These solvers have names ending with the following strings:

- `-sw_optimize`
- `-sw_sample`

Optimizing Emulators

This class of solvers is entirely deterministic, so the semantics of some parameters are different. The number of solutions returned is always the lesser of `max_answers` and `num_reads` \times `num_programming_cycles`. The solutions returned are the lowest-energy states, sorted in increasing-energy order.

These solvers accept the following parameters, described above:

- *answer_mode*
- *max_answers*
- *num_reads*

The acceptable range and the default value of each field are given in the table below.

Field	Range	Default value
<i>answer_mode</i>	histogram or raw	histogram
<i>max_answers</i>	> 0	<i>num_reads</i>
<i>num_reads</i>	> 0	1

When *answer_mode* is histogram, the *num_occurrences* field contains all ones, except possibly for the lowest-energy solution. That first entry is set so that the sum of all entries is $num_reads \times num_programming_cycles$.

Note: The *num_programming_cycles* parameter is deprecated and will be removed in a future release. Plan code updates accordingly.

Sampling Emulators

This class of solvers mimic the probabilistic nature of the QPU, drawing samples from a Boltzmann distribution; that is, state s is sampled with probability proportional to:

$$\exp(-\beta E(s))$$

where β is some parameter and $E(s)$ is the energy of state s .

These solvers accept the following parameters, described above:

- *answer_mode*
- *beta*
- *max_answers*
- *num_reads*
- *random_seed*

In addition to this list, the sampling emulators accept an additional parameter: *random_seed*. This parameter provides a random number generator seed. If provided, the solver solves the same problem with the same parameters each read thereby producing the same results every time. If no value is provided, a time-based seed is selected, and solutions will vary.

The acceptable range and the default value of each field are given in the following table.

Field	Range	Default value
<i>answer_mode</i>	histogram or raw	histogram
<i>beta</i>	Any finite value	3.0
<i>max_answers</i>	> 0	<i>num_reads</i>
<i>num_reads</i>	> 0	1
<i>random_seed</i>	>= 0	Randomly set

4 Ising Heuristic Solvers

Note: Not all accounts have access to this type of solver.

The Ising heuristic solver is intended to solve complex problem structures. It has some important differences from other solvers:

- There is no fixed problem structure. In particular, this solver does not have the properties *num_qubits*, *qubits*, and *couplers*.
- Only one solution is returned and it is not guaranteed to be optimal.
- Solver properties and parameters are entirely disjoint from those of other solvers.
- It cannot be used with the QSage solver.

Note that this heuristic solver can handle problems of arbitrary structure.

4.1 Algorithm

The core of the heuristic solver is a local search based on optimizing low-treewidth subgraphs. In pseudocode:

```
function local_search(problem, x)
  stuck = 0
  e = evaluate(problem, x)
  while (time limit not exceeded) and (stuck <= local_stuck_limit)
    select random low-treewidth subproblem
    (new_e, x) = solve(subproblem, x)
    if subproblem is the entire problem
      return (new_e, x, exact=true)
    if new_e < e
      stuck = 0
    else
      stuck = stuck + 1
    e = new_e
  return (e, x, exact=false)
```

The *local_stuck_limit* value is a user-supplied parameter. What constitutes a “low-treewidth subproblem” is determined by the parameter *max_local_complexity*. The time limit is provided by the parameter *time_limit_seconds*.

To escape local minima, multiple copies of the solution are made and bits are randomly flipped.

```
function ising_heuristic(problem)
  x = random solution vector
  (e, x, exact) = local_search(problem, x)
  if exact
    return (e, x)
  best_e = current_e = e; best_x = current_x = x
  iter = 0
  while (time limit not exceeded) and (iter < iteration_limit)
```

```

iter = iter + 1
new_e = current_e; new_x = current_x
for i = 1 to num_perturbed_copies
  x = current_x
  flip each bit of x with probability p(i)
  (e, x, exact) = local_search(problem, x)
  if exact
    return (e, x)
  if e < new_e
    new_e = e; new_x = x
current_e = new_e; current_x = new_x
if current_e < best_e
  best_e = current_e; best_x = current_x
return (best_e, best_x)

```

In the `ising_heuristic` function, `iteration_limit` and `num_perturbed_copies` are user-provided parameters. The (i-dependent) bit flip probability $p(i)$ is determined by the parameters `min_bit_flip_prob` and `max_bit_flip_prob`.

4.2 Solver Name

`ising-heuristic`

4.3 Properties

None, except for the common property `supported_problem_types`.

4.4 Parameters

This section describes the parameters accepted by Ising heuristic solvers, in alphabetical order. See *Summary of Ising Heuristic Solver Parameters* for a summary and for the default values.

Note: Many of these parameters described here require a high-level understanding of the heuristic algorithm. Parameter settings can wildly affect solver performance and solution quality. It is difficult in general to know what good values are a priori; defaults are selected to favor quicker run times over aggressive searches. Therefore, experimentation with these values is recommended.

`iteration_limit`

Specifies the maximum number of solver iterations, not including the initial local search. Must be an integer ≥ 0 ; default = 10.

local_stuck_limit

Specifies the number of consecutive local search steps that do not improve solution quality to allow before determining a solution to be a local optimum. Larger values produce more thorough local searches but increase run time. Must be an integer > 0 ; default = 8.

max_bit_flip_prob

Bit flip probability range. The probability of flipping each bit is constant for each perturbed solution copy but varies across copies. The probabilities used are linearly interpolated between *min_bit_flip_prob* and *max_bit_flip_prob*. Larger values allow more exploration of the solution space and easier escapes from local minima but may also discard nearly-optimal solutions.

Must be a number [0.0 1.0] and $\text{min_bit_flip_prob} \leq \text{max_bit_flip_prob}$, default $\text{min_bit_flip_prob} = 1.0 / 32.0$, default $\text{max_bit_flip_prob} = 1.0 / 8.0$.

max_local_complexity

Specifies the maximum complexity of subgraphs used during local search. The run time and memory requirements of each step in the local search are exponential in this parameter. Larger values allow larger subgraphs (which can improve solution quality) but require much more time and space. Subgraph “complexity” here means $\text{treewidth}+1$.

Must be an integer > 0 ; default = 9.

min_bit_flip_prob

Bit flip probability range. The probability of flipping each bit is constant for each perturbed solution copy but varies across copies. The probabilities used are linearly interpolated between *min_bit_flip_prob* and *max_bit_flip_prob*. Larger values allow more exploration of the solution space and easier escapes from local minima but may also discard nearly-optimal solutions.

Must be a number [0.0 1.0] and $\text{min_bit_flip_prob} \leq \text{max_bit_flip_prob}$, default $\text{min_bit_flip_prob} = 1.0 / 32.0$, default $\text{max_bit_flip_prob} = 1.0 / 8.0$.

num_perturbed_copies

Number of perturbed solution copies created at each iteration. Run time is linear in this value.

Must be an integer > 0 ; default = 4.

num_variables

Provides a lower bound on the number of variables. This solver can accept problems of arbitrary structure and the size of the solution returned is determined by the maximum variable index in the problem. The size of the solution can be increased by setting this parameter.

Must be an integer ≥ 0 , default = 0.

random_seed

Provides a random number generator seed. When a value is provided, solving the same problem with the same parameters will produce the same results every read. If no value is provided, a time-based seed is selected.

The use of a wall clock-based timeout may in fact cause different results with the same *random_seed* value. If the same problem is run under different CPU load conditions (or on computers with different performance), the amount of work completed may vary despite the fact that the algorithm is deterministic. If repeatability of results is important, rely on the *iteration_limit* parameter rather than the *time_limit_seconds* parameter to set the stopping criterion.

Must be an integer ≥ 0 ; default is a time-based seed.

time_limit_seconds

Specifies the maximum wall clock time in seconds. Actual run times will exceed this value slightly.

Must be a number ≥ 0.0 ; default = 5.0.

4.5 Summary of Ising Heuristic Solver Parameters

Ising heuristic solver parameters and their default values are summarized in the table below.

Table 1: Ising Heuristic Solver Parameters

Parameter	Range	Default Value
<i>iteration_limit</i>	≥ 0	10
<i>local_stuck_limit</i>	> 0	8
<i>max_bit_flip_prob</i>	[<i>min_bit_flip_prob</i> 1.0]	1/8
<i>max_local_complexity</i>	> 0	9
<i>min_bit_flip_prob</i>	[0.0 1.0]	1/32
<i>num_perturbed_copies</i>	> 0	4
<i>num_variables</i>	≥ 0	0
<i>random_seed</i>	≥ 0	Time-based seed
<i>time_limit_seconds</i>	≥ 0.0	5

5 Quick Reference Tables

This section provides two quick reference tables: *Properties and Parameters Summary by Solver Type* and *QPU Solver Parameters Matrix*.

5.1 Properties and Parameters Summary by Solver Type

The following table summarizes the properties and parameters of each solver described above. For release-specific details of QPU solver parameters, see *QPU Solver Parameters Matrix*.

Note: Not all accounts have access to all solver types.

Solver Type	Properties	Parameters
QPU and VFYC solvers	<ul style="list-style-type: none"> <i>anneal_offset_step</i> <i>anneal_offset_step_phi0</i> <i>anneal_offset_ranges</i> <i>annealing_time_range</i> <i>chip_id</i> <i>couplers</i> <i>default_annealing_time</i> <i>default_programming_thermalization</i> <i>default_readout_thermalization</i> <i>extended_j_range</i> <i>h_range</i> <i>j_range</i> <i>max_anneal_schedule_points</i> <i>num_reads_range</i> <i>num_qubits</i> <i>parameters</i> <i>per_qubit_coupling_range</i> <i>problem_run_duration_range</i> <i>qubits</i> <i>readout_thermalization_range</i> <i>supported_problem_types</i> <i>vfyc</i> 	<ul style="list-style-type: none"> <i>anneal_offsets</i> <i>anneal_schedule</i> <i>annealing_time</i> <i>answer_mode</i> <i>auto_scale</i> <i>beta</i> <i>chains</i> <i>flux_biases</i> <i>flux_drift_compensation</i> <i>initial_state</i> <i>max_answers</i> <i>num_reads</i> <i>num_spin_reversal_transforms</i> <i>postprocess</i> <i>programming_thermalization</i> <i>qpu_access_overhead_time</i> <i>reduce_intersample_correlation</i> <i>readout_thermalization</i> <i>reinitialize_state</i> <i>reverse_anneal</i>
Optimizing emulators	<ul style="list-style-type: none"> <i>couplers</i> <i>num_qubits</i> <i>num_reads_range</i> <i>parameters</i> <i>qubits</i> <i>supported_problem_types</i> 	<ul style="list-style-type: none"> <i>answer_mode</i> <i>max_answers</i> <i>num_reads</i>
Sampling emulators	<ul style="list-style-type: none"> <i>beta_range</i> <i>couplers</i> <i>default_beta</i> <i>num_qubits</i> <i>num_reads_range</i> <i>parameters</i> <i>qubits</i> <i>supported_problem_types</i> 	<ul style="list-style-type: none"> <i>answer_mode</i> <i>beta</i> <i>max_answers</i> <i>num_reads</i> <i>random_seed</i>
Ising heuristic solvers	<ul style="list-style-type: none"> <i>supported_problem_types</i> 	<ul style="list-style-type: none"> <i>iteration_limit</i> <i>local_stuck_limit</i> <i>max_bit_flip_prob</i> <i>max_local_complexity</i> <i>min_bit_flip_prob</i> <i>num_perturbed_copies</i> <i>num_variables</i> <i>random_seed</i> <i>time_limit_seconds</i>

5.2 QPU Solver Parameters Matrix

The following table summarizes the QPU solver parameters, identifies on what platform and in what release they are available, provides guidelines on how parameters interact with each other and on timing,¹ and gives their default settings.

Notice that some parameters are not supported on D-Wave 2X and earlier platforms. Always check solver properties when using a new parameter.

¹ For more information on timing, see *Measuring Computation Time on D-Wave Systems*.

Table 2: QPU Solver Parameters Matrix

Parameter	Pltfrm./ Rel.	Parameter Compatibility	Timing Impact?	Default
<i>anneal_offsets</i>	2000Q, 2.5+		No	No offsets.
<i>anneal_schedule</i>	2000Q, 2.11+	Incompatible with <i>annealing_time</i> . Use with <i>initial_state</i> and <i>reinitialize_state</i> for reverse annealing.	Yes	No variation to the standard anneal schedule.
<i>annealing_time</i>	All	Incompatible with <i>anneal_schedule</i> .	Yes	Check properties.
<i>answer_mode</i>	All		No	histogram
<i>auto_scale</i>	All	Incompatible with <i>flux_biases</i> and extended <i>J</i> range.	Yes	true
<i>beta</i>	All		No	Check properties. For the VFYC solver, this is 10.
<i>chains</i>	All	Use with <i>postprocess</i> .	No	No chains.
<i>flux_biases</i>	2000Q, 3.0+	Incompatible with <i>auto_scale</i> and <i>num_spin_reversal_transforms</i> .	No	No offsets.
<i>flux_drift_compensation</i>	2000Q, 3.0+		No	true
<i>h</i>	All		No	No bias.
<i>initial_state</i>	2000Q, 3.0+	Use with <i>anneal_schedule</i> for reverse annealing. Incompatible with <i>annealing_time</i> .	Yes	No initial state.
<i>J</i>	All		No	No coupling.
<i>max_answers</i>	All		No	<i>num_reads</i>
<i>num_reads</i>	All	The web server supports up to 1000 reads. SAPI clients support up to 10,000.	Yes	1
<i>num_spin_reversal_transforms</i>	All	Incompatible with <i>flux_biases</i> and extended <i>J</i> range.	Yes	0
<i>postprocess</i>	All	Use with <i>chains</i> .	Yes	None. ²
<i>programming_thermalization</i>	All		Yes	Check properties.
<i>Q</i>	All		No	N/A
<i>qpu_access_overhead_time</i>	All		Yes	Introduced by the system. Roughly 1 ms.
<i>reduce_intersample_correlation</i>	2000Q, 3.0+		Yes	false
<i>readout_thermalization</i>	All		Yes	Check properties.
<i>reinitialize_state</i>	2000Q, 3.0+	Use with <i>initial_state</i> and <i>reinitialize_state</i> for reverse annealing.	Yes	true
<i>reverse_anneal</i> ³	2000Q, 3.0+		Yes	NULL; no value set.

² The VFYC solver always runs postprocessing; its default is `sampling`.

³ Relevant only to the C client. Pointer to the *sapi_ReverseAnneal* structure.