

Simulated Bifurcation Machine (SBM) User Manual

Revision 1.01

© 2019 Toshiba Digital Solutions Corporation

Amazon Web Services, the "Powered by AWS" logo, AWS Marketplace and the AWS Marketplace logo are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries.

All other company names or product names (mentioned herein) may be trademarks of their respective companies.

Contents

Introduction	3
Installation	4
Installation steps	4
Installation notes	4
Using SBM	5
Interface overview	5
Examples of SBM calculation	5
Calculation structure	5
Calculation parameters	6
Calculation result	7
Message strings specification	8
Statistical information	8
Error messages	9
System configuration	9
Specification of each solver	10
ISING	10
Max-CUT	11
Max-SAT	12
Maintenance and Security	13
Log files	13
Authentication	13
Appendix	14
Examples of SBM calculation	14
Getting Help	17
Disclaimer	17

Introduction

Simulated Bifurcation Machine (SBM) is a set of solvers enabling users to quickly obtain nearly optimal solutions for large combinatorial optimization problems. SBM is developed based on the theory described in the following paper:

Goto, H., Tatsumura, K., & Dixon, A. R. (2019) Combinatorial optimization by simulating adiabatic bifurcations in nonlinear Hamiltonian systems, *Science Advances*, 5(4), DOI:10.1126/sciadv.aav2372

SBM is provided as an AMI (See Figure 1). The main characteristics are:

- REST-API
- Solves ISING, Max-CUT, and Max-SAT optimization problems.
- Quickly obtains nearly optimal solutions for large optimization problems; however, it does not guarantee optimal solutions.
- Easy to use. Complicated parameter setting is not required.
- Each solver has its maximum size limits. For details, refer to the section "Using SBM."




Figure 1 Functional diagram of the SBM AMI

Installation

Installation steps

1. Select the SBM AMI and its associated instance in the AWS Marketplace using the AWS console. Only a p3.2xlarge instance is supported.
2. Set the security group. The inbound rule is shown in Table 1. The port number (8000) of SBM calculation service can be customized by modifying the system configuration file. For details, refer to the section "System configuration."
3. Launch the SBM instance. Once the SBM instance starts, SBM calculation service will start automatically.

Table 1: Inbound rule of the security group

Type	Protocol	Port Range
Custom TCP rule	TCP	8000
SSH	TCP	22

Installation notes

When connecting to the SBM instance via REST-API, it is recommended to connect from the instance in the VPC where the SBM instance resides. Connecting directly from the Internet is not recommended because API requests are not authenticated.

Using SBM

Interface overview

- Method
 - POST
- URL
 - `http://{ip}:{port}/solver/{solver}?{parameter}`
 - ✧ ip: IP address of the SBM instance.
 - ✧ port: Port number of the API server.
 - ✧ solver: Solver name (one of `ising`, `maxcut`, or `maxsat`)
 - ✧ parameter: Specified as a query string, i.e. consisting of key=value pairs separated by "&".
 - Example

`http://123.45.67.89:8000/solver/ising?steps=1000&loops=10`
- Request Header
 - Set the Content-Type field in the Request Header to "application/octet-stream."
- Request Body
 - Contains a problem file for each solver.
 - Every solver has its own format.
- Response Body
 - Outputs the calculation results.
 - Encoded in the JSON format (`application/json`).

Examples of SBM calculation

Refer to the section "Examples of SBM calculation" in the Appendix.

Calculation structure

A calculation request consists of three calculation units: step, run, and loop in ascending order (Figure 2). The smallest unit is a "step." A group of steps constitutes a "run"; a group of runs constitutes a "loop"; a group of loops constitutes a "request." The number of steps and loops can be specified using the calculation parameters described in the next section, but the number of runs is solver-specific. See the section "Specification of each solver."




Figure 2 Calculation structure

Here SBM executes a group of runs in parallel while each run produces a nearly optimal solution. This means a request produces $(\text{number of loops in a request}) \times (\text{number of runs in a loop})$ optimal solutions in total. SBM calculation returns the best of these solutions as the result of a request together with statistical information on these solutions.

The number of steps and loops affects the accuracy and speed of SBM calculation. The auto step functionality will help you find a better setting. For detailed parameter settings, refer to the next section.

Calculation parameters

The parameters specified in the query string of the URL are as follows:

Table 2: Calculation parameters

Parameter	Default	Description
steps	0	Specifies the number of steps in SBM calculation. The value 0 means automatic setup where the number of steps is dynamically determined.
loops	1	Specifies the number of loops in SBM calculation. SBM searches for a better solution through repeated calculations. If 0 is specified, calculation will be repeated to the maximum calculation time.

Parameter	Default	Description
timeout	<i>configurable</i>	Specifies the maximum calculation time in seconds. If it is likely that this <i>timeout</i> is reached before the next calculation loop finishes, the calculation stops and the best solution obtained so far will be the calculation result.
maxwait	<i>configurable</i>	Specifies the maximum waiting time in seconds for each request. An accepted request waits in a queue for the preceding requests being processed. If this time passes during this wait, the request ends with a timeout error.
target	<i>unspecified</i>	Specifies the end condition of calculation. When the value of the objective function reaches this value, the calculation will stop.
prefer	auto	Select "speed" or "auto." "speed" refers to a setup where calculation time is prioritized over accuracy. "auto" refers to a standard setup. In either setup, an optimal result and faster calculation time are not guaranteed however.
stats	<i>configurable</i>	Specifies how to output statistical information by choosing "none," "summary," or "full." "none" outputs nothing. "summary" outputs the average, standard deviation, and the best frequency. "full" outputs the average, standard deviation, and a histogram.

Calculation result

Calculation results included in the Response Body are shown in Table 3:

Table 3: Calculation results

Name	Description
time	Net calculation time in seconds, not including overheads, such as data transfer time.
result	The best solution found.
value	The value of the objective function calculated using the solution returned in the "result."
message	Completion conditions, etc.
wait	Time to wait for the preceding requests to finish in seconds.

Name	Description
runs	The number of the processed executable units.
steps	The number of steps by which the "result" has been obtained.
id	Internally assigned request ID.
stats	Statistical information. If stats="none", no output.

Note that the input and output data are not stored in the SBM instance after returning the results.

Message strings specification

Table 4: Message strings specification

String	Specification
finished	Completed the calculation of the specified number of steps and loops.
timeout	Stopped the calculation because the calculation time has reached its timeout.
reached	Stopped the calculation because the "value" has reached its target.

Statistical information

The average, standard deviation, the best frequency, and/or a histogram are outputted according to the stats parameter settings. A histogram is a two-dimensional array which consists of values and their frequencies. The array is sorted in the ascending order of values. If the histogram becomes too large, its classes will be merged. In this case, the class width of the histogram is outputted as *histogram_width*.

Examples of histograms

```
"stats":{"avg":150.8766,"stddev":4.501432,"histogram":[[144,39],[145,23],
[146,84],[147,172],[148,226],[149,148],[150,73],[151,52],[152,37],[153,44
],[154,44],[155,62],[156,74],[157,61],[158,63],[159,27],[160,28],[161,7],
[162,10],[163,3],[164,2],[167,1]]]}
```

```
"stats":{"avg":150.6531,"stddev":4.407379,"histogram_width":10,"histogram
":[[140,707],[150,539],[160,34]]]}
```

Error messages

If the calculation ends with an error, an HTTP response status code other than 200 will be returned. The HTTP Response Body shows the cause of the error.

Table 5: Error messages

Status code	Message	Description
400	(The message varies.)	The format of Request Body is invalid. The message gives the detailed information.
404	N/A	The requested URL is invalid.
413	N/A	The payload size of the HTTP request has exceeded <i>payload_limit</i> . See the section "System configuration."
503	busy	The number of requests has exceeded the request queue length.
504	timeout	<i>maxwait</i> time has passed before the request is processed.

System configuration

`/home/ec2-user/config` sets the following system parameters. The values in the parentheses denote the default settings. Usually there is no need to change the settings. If you do change the setting, restart the SBM instance for the changes to take effect.

Table 6: System configuration parameters

Parameter	Description
<code>timeout_upper_bound</code> (600)	Upper limit of the calculation time in seconds. Regardless of the timeout setting in the request, calculation ends at this bound.
<code>payload_limit</code> (1,000,000)	Maximum size of Request Body in bytes.
<code>max_requests</code> (20)	Number of requests acceptable at one time (queue length). SBM accepts multiple requests concurrently but calculates them one by one.
<code>server_port</code> (8000)	Port number of the HTTP server.
<code>number_of_workers</code> (4)	Number of worker threads which accept HTTP requests.

Parameter	Description
autoshutdown (disabled)	Time to power off the SBM instance automatically in seconds. The auto shutdown function executes /sbin/poweroff if there is no processing request during the specified time. This function is disabled if the parameter is commented out in the configuration file.
max_histogram (1,000)	Upper limit of the number of histogram classes in the statistical information. Classes will be merged so as not to exceed the limit.
default_timeout (60)	Default value of the calculation parameter "timeout"
default_maxwait (60)	Default value of the calculation parameter "maxwait"
default_stats (none)	Default value of the calculation parameter "stats"

Specification of each solver

ISING

- Request Body
 - Place the content of an ISING problem file in the Request Body.
 - The energy of the problem is defined using the QUBO model.
 - ✧ The energy E is represented by the following formula:

$$E = \sum_{i \leq j} Q_{ij} \cdot x_i \cdot x_j$$

where Q_{ij} denotes the element of the matrix, and $x_i \in \{0, 1\}$ is a spin variable.

- ✧ This formula can be transformed into the following formula:

$$E = \sum_{i < j} Q_{ij} \cdot x_i \cdot x_j + \sum Q_{ii} \cdot x_i$$

- ✧ Coefficients Q_{ij} represent single-precision floating point numbers.
- The file format must be:
 - ✧ MatrixMarket format (symmetric form only)
(<https://math.nist.gov/MatrixMarket/formats.html>)

- Response Body
 - value: The energy of QUBO.
 - result: An array of the binary value of each node
 - ✧ Example (Number of nodes = 5): [0,1,1,0,0]
- Maximum size of the problem
 - Maximum number of spins: 10,000
 - Maximum number of non-zero coefficients in the matrix: 1,000,000
- Additional note
 - The number of runs per loop is 320.

Max-CUT

- Request Body
 - Place the content of a Max-CUT problem file in the Request Body.
 - ✧ The weight of each edge must be ± 1 .
 - The file format must be one of the followings.
 - ✧ MatrixMarket format (symmetric type only)
 - (<https://math.nist.gov/MatrixMarket/formats.html>)
 - ✧ Gset format (<https://web.stanford.edu/~yyye/yyye/Gset/>)
 - The content of the Max-CUT problem file represents a symmetric matrix with zero diagonal components.
- Response Body
 - value: The cut value
 - result: An array of the binary value of each node
 - ✧ Example (Number of nodes = 5): [0,1,1,0,0]
- Maximum size of the problem
 - Maximum number of nodes: 10,000
 - Maximum number of edges: 1,000,000
- Additional note
 - The number of runs per loop is 320.

Max-SAT

- Request Body
 - Place the content of a Max-SAT problem file in the Request Body.
 - The weights of a weighted problem represent a single-precision floating point numbers.
 - The file format must be as follows:
 - ✧ DIMACS format¹ (<http://maxsat.ia.udl.cat/requirements/>)
- Response Body
 - value: The number of violated clauses; in the case of weighted problems, the sum of weights of violated clauses.
 - result: An array of the binary value of each variable (false: 0, true: 1)
 - ✧ Example (Number of nodes = 5): [0, 1, 1, 0, 0]
- Maximum size of the problem
 - Maximum number of variables: 1,000
 - Maximum number of clauses: 3,000
 - Maximum number of variables allowed in the Max-SAT formula: 10,000
 - Maximum number of variables in a clause: 32
- Additional note
 - The number of runs per loop is 1,280.

¹ Partial Max-SAT and Weighted Partial Max-SAT input formats are not supported.

Maintenance and Security

Log files

Log files are created under `/home/ec2-user/logs`. Obtain them by ssh login to the SBM instance as `ec2-user`.

Authentication

The SBM calculation service does not provide authentication. It is NOT recommended to connect SBM instances to the Internet directly.

Getting Help

- User Manual (this document)
- General inquiries: tdsl-sbm@toshiba-sol.co.jp

Disclaimer

THE SOFTWARE AND SERVICES ARE PROVIDED "AS IS." YOU AND YOUR USERS ACCEPT AND ASSUME THE ENTIRE RISK AS TO THE QUALITY, PERFORMANCE AND RESULTS OF ACCESS OR USE OF THE SOFTWARE AND SERVICES. TOSHIBA DIGITAL SOLUTIONS CORPORATION (TDSL) AND TDSL'S AFFILIATES AND LICENSORS MAKE NO REPRESENTATIONS OR WARRANTIES OF ANY KIND, WHETHER EXPRESS, IMPLIED, STATUTORY OR OTHERWISE REGARDING THE SOFTWARE AND SERVICES, INCLUDING ANY WARRANTY THAT THE SOFTWARE AND SERVICES WILL BE UNINTERRUPTED, ERROR FREE OR FREE OF HARMFUL COMPONENTS, OR THAT ANY CONTENT IN YOUR AWS ACCOUNT, INCLUDING YOUR CONTENT AND DATA, WILL BE SECURE OR NOT OTHERWISE LOST OR DAMAGED. EXCEPT TO THE EXTENT PROHIBITED BY LAW, TDSL AND TDSL'S AFFILIATES AND LICENSORS DISCLAIM ALL WARRANTIES, INCLUDING ANY IMPLIED WARRANTIES OF PERFORMANCE, MERCHANTABILITY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR QUIET ENJOYMENT, AND ANY WARRANTIES ARISING OUT OF ANY COURSE OF DEALING OR USAGE OF TRADE.