

LINE: Queueing Analysis Algorithms

User manual for Python

Disclaimer: LINE for Python is in alpha version, some features are still under development and therefore marked in the manual as TODO.

Last revision: July 12, 2024

Contents

1	Introduction	5
1.1	What is LINE?	5
1.2	Obtaining the latest release	6
1.3	References	6
1.4	Contact and credits	8
1.5	Copyright and license	8
1.6	Acknowledgement	8
2	Getting started	9
2.1	Installation and support	9
2.1.1	Software requirements	10
2.1.2	Documentation	11
2.1.3	Getting help	11
2.2	Getting started examples	11
2.2.1	Controlling verbosity	12
2.2.2	Model gallery	12
2.2.3	Example 1: A M/M/1 queue	12
2.2.4	Example 2: A multiclass M/G/1 queue	14
2.2.5	Example 3: Machine interference problem	16
2.2.6	Example 4: Round-robin load-balancing	18
2.2.7	Example 5: Modelling a re-entrant line	19
2.2.8	Example 6: A queueing network with caching	22
2.2.9	Example 7: Response time distribution and percentiles	24
2.2.10	Example 8: Optimizing a performance metric	25
2.2.11	Example 9: Studying a departure process	26
2.2.12	Example 10: Evaluating a CTMC symbolically	27
3	Network models	29
3.1	Network object definition	30

3.1.1	Creating a network and its nodes	30
3.1.2	Advanced node parameters	34
3.1.3	Job classes	34
3.1.4	Routing strategies	37
3.1.5	Class switching	40
3.1.6	Service and inter-arrival time processes	42
3.2	Internals	45
3.2.1	Representation of the model structure	45
3.3	Debugging and visualization	47
3.4	Model import and export	48
3.4.1	Creating a LINE model using JMT	49
3.4.2	Supported JMT features	50
4	Analysis methods	52
4.1	Performance metrics	52
4.2	Steady-state analysis	52
4.2.1	Station average performance	52
4.2.2	Station response time distribution	54
4.2.3	System average performance	54
4.3	Specifying states	54
4.3.1	Station states	55
4.3.2	Network states	57
4.3.3	Initialization of transient classes	58
4.3.4	State space generation	58
4.4	Transient analysis	58
4.4.1	Computing transient averages	58
4.4.2	First passage times into stations	59
4.5	Sample path analysis	59
4.6	Sensitivity analysis and numerical optimization	60
4.6.1	Fast parameter update	60
4.6.2	Refreshing a network topology with non-probabilistic routing	61
4.6.3	Saving a network object before a change	61
5	Network solvers	62
5.1	Overview	62
5.2	Solution methods	63
5.2.1	LINE	66
5.2.2	CTMC	66
5.2.3	FLUID	67
5.2.4	JMT	67

5.2.5	MAM	68
5.2.6	MVA	68
5.2.7	NC	69
5.2.8	SSA	70
5.3	Supported language features and options	70
5.3.1	Solver features	70
5.3.2	Class functions	70
5.3.3	Node types	73
5.3.4	Scheduling strategies	73
5.3.5	Statistical distributions	74
5.3.6	Solver options	75
5.4	Solver maintenance	77
6	Layered network models	78
6.1	Basics about layered networks	78
6.2	LayeredNetwork object definition	79
6.2.1	Creating a layered network topology	79
6.2.2	Describing host demands of entries	80
6.2.3	Debugging and visualization	81
6.3	Internals	82
6.3.1	Representation of the model structure	82
6.3.2	Decomposition into layers	83
6.4	Solvers	83
6.4.1	LQNS	85
6.4.2	QNS	85
6.4.3	LN	86
6.5	Model import and export	86
7	Random environments	87
7.1	Environment object definition	87
7.1.1	Specifying the environment	87
7.1.2	Specifying a reset policy	88
7.1.3	Specifying system models for each stage	89
7.2	Solvers	89
7.2.1	ENV	89
A	Examples	94

Chapter 1

Introduction

1.1 What is LINE?

LINE is an open-source software package to analyze queueing models via analytical methods and simulation. The tool aims at simplifying the computation of performance and reliability metrics in models of systems such as software applications, business processes, or computer networks. LINE decomposes a high-level system model into one or more stochastic models, typically extended queueing networks, that are subsequently analyzed using either numerical algorithms or simulation. The development of the stand-alone Python version covered in this manual is presently in alpha version (<https://github.com/imperial-qore/line-solver/tree/main/python>) and is mainly a wrapper for the Java version (<https://github.com/imperial-qore/line-solver/tree/main/java>). Other tool versions can be found at <https://github.com/imperial-qore/line-solver>.

A key feature of LINE is that the solver decouples the model description from the solvers used for its solution. That is, LINE implements model-to-model transformations that automatically translate the model specification into the input format (or data structure) accepted by the target solver. External solvers supported by LINE include Java Modelling Tools (JMT; <http://jmt.sf.net>) and LQNS (<http://www.sce.carleton.ca/rads/lqns/>). Native model solvers are instead based on formalisms and techniques such as:

- Continuous-time Markov chains (CTMC)
- Fluid ordinary differential equations (FLUID)
- Matrix analytic methods (MAM)
- Normalizing constant analysis (NC)
- Mean-value analysis (MVA)
- Stochastic simulation (SSA)

Each solver encodes a general solution paradigm and can implement both exact and approximate analysis methods. For example, the *MVA* solver implements both exact mean value analysis (MVA) and approximate mean value analysis (AMVA). The offered methods typically differ for accuracy, computational cost, and the subset of model features they support. A special solver (*AUTO*) is supplied that provides an automated recommendation on which solver to use for a given model.

The above techniques can be applied to models specified in the following formats:

- *LINE modeling language*. This is a domain-specific object-oriented language designed to resemble the abstractions available in JMT’s queueing network simulator (JSIM).
- *Layered queueing network models (LQNS XML format)*. LINE is able to solve a sub-class of layered queueing network models, either specified using the LINE modeling language or according to the XML metamodel of the LQNS solver.
- *JMT simulation models (JSIMg, JSIMw formats)*. LINE is able to import and solve queueing network models specified using JSIMgraph and JSIMwiz. LINE models can be exported to, and visualized with, JSIMgraph and JSIMwiz.
- *Performance Model Interchange Format (PMIF XML format)*. LINE is able to import and solve closed queueing network models specified using PMIF v1.0.

1.2 Obtaining the latest release

This document contains the user manual for LINE version 2.0.x, which can be obtained from:

<http://line-solver.sourceforge.net/>

LINE 2.0.x has been tested using Python version 3.10, and IntelliJ DataSpell 2024.1.1 for the Jupyter notebooks.

1.3 References

To cite the LINE solver, we recommend to reference:

- G. Casale. “Integrated Performance Evaluation of Extended Queueing Network Models with LINE”, in *Proc. of WSC 2020*, ACM Press, Dec 2020. *This paper presents the technical approach used to develop Line 2.0.x.*

The following papers discuss LINE or use it in applications:

- R.-A. Dobre, Z. Niu, G. Casale. Approximating Fork-Join Systems via Mixed Model Transformations. *Proc. of WOSP-C, 2024. This paper presents the fork-join approximation technique implemented in SolverMVA.*

- G. Casale, Y. Gao, Z. Niu, L. Zhu. LN: a Meta-Solver for Layered Queueing Network Analysis, Proceedings of QEST, 22 pages, Sep 2022. *This paper gives a short introduction to the Layered Queueing Network solver available in Line 2.0.x.* Later extended into ACM Transactions on Modeling and Computer Simulation, 2024.
- Y. Gao, G. Casale. JCSP: Joint Caching and Service Placement for Edge Computing Systems, in Proc. of IEEE/ACM IWQoS, 10 pages, June 2022. *This work introduces caching layers in Layered Queueing Networks within Line's LN solver.*
- J. Ruuskanen, T. Berner, K.-E. Årzén et al., Improving the mean-field fluid model of processor sharing queueing networks for dynamic performance models in cloud computing, Performance Evaluation (2021). *This work applies Line-JMT model-to-model transformations to validate mean-field approximations for product-form queueing networks..*
- Z. Niu, G. Casale. A Mixture Density Network Approach to Predicting Response Times in Layered Systems, in Proceedings of IEEE MASCOTS, 8 pages, Nov 2021. *This work applies mixture density networks for end-to-end latency percentile estimations in Layered Queueing Networks using Line.*
- Y. Chen, G. Casale. Deep Learning Models for Automated Identification of Scheduling Policies, in Proceedings of IEEE MASCOTS, 8 pages, Nov 2021. *This work uses LINE's SSA solver to generate scheduling traces for training neural networks to infer scheduling policies.*
- G. Russo Russo, V. Cardellini, G. Casale, F. Lo Presti. MEAD: Model-Based Vertical Auto-Scaling for Data Stream Processing, in Proc. of IEEE/ACM CCGRID, 10 pages, May 2021. *This work uses Line's MAM solver for burstiness-aware auto-scaling in data streaming systems.*
- G. Casale. "Automated Multi-paradigm Analysis of Extended and Layered Queueing Models with LINE", in Proc. of ACM/SPEC 2019, ACM Press, Apr 2019. *This paper gives a short introduction to Line 2.0.0.*
- J. F. Pérez and G. Casale. "LINE: Evaluating Software Applications in Unreliable Environments", in IEEE Transactions on Reliability, Volume 66, Issue 3, pages 837-853, Feb 2017. *This paper introduces the core algorithms behind Line 1.0.0.*
- C. Li and G. Casale. "Performance-Aware Refactoring of Cloud-based Big Data Applications", in Proceedings of 10th IEEE/ACM International Conference on Utility and Cloud Computing, 2017. *This paper uses Line to model stream processing systems.*
- D. J. Dubois, G. Casale. "OptiSpot: minimizing application deployment cost using spot cloud resources", in Cluster Computing, Volume 19, Issue 2, pages 893-909, 2016. *This paper uses Line to determine bidding costs in spot VMs.*

- R. Osman, J. F. Pérez, and G. Casale. “Quantifying the Impact of Replication on the Quality-of-Service in Cloud Databases”. Proceedings of the IEEE International Conference on Software Quality, Reliability and Security (QRS), 286-297, 2016. *This paper uses Line to model the Amazon RDS database.*
- C. Müller, P. Rygielski, S. Spinner, and S. Kounev. Enabling Fluid Analysis for Queueing Petri Nets via Model Transformation, *Electr. Notes Theor. Comput. Sci*, 327, 71–91, 2016. *This paper uses Line to analyze Descartes models used in software engineering.*
- J. F. Pérez and G. Casale. “Assessing SLA compliance from Palladio component models,” in Proceedings of the 2nd Workshop on Management of resources and services in Cloud and Sky computing (MICAS), IEEE Press, 2013. *This paper uses Line to analyze Palladio component models used in model-driven software engineering.*

1.4 Contact and credits

Project coordinator: Giuliano Casale, Department of Computing, Imperial College London, 180 Queen’s Gate, SW7 2AZ, London, United Kingdom. Web: <http://wp.doc.ic.ac.uk/gcasale/>

Please refer to the following file for detailed credits:

1.5 Copyright and license

Copyright Imperial College London (2012-Present). LINE is freeware and open-source, released under the 3-clause BSD license. Additional licensing information is available in the file:

<https://github.com/imperial-qore/line-solver/blob/main/LICENSE>

1.6 Acknowledgement

LINE has been partially funded by the European Commission grants FP7-318484 (MODAClouds), H2020-644869 (DICE), H2020-825040 (RADON), and by the EPSRC grant EP/M009211/1 (OptiMAM).

Chapter 2

Getting started

2.1 Installation and support

This is the fastest way to get started with LINE:

1. Obtain the latest release:

- Stable release (zip file): <https://sourceforge.net/projects/line-solver/files/latest/download>
- Development release (git): <https://github.com/imperial-qore/line-solver/>

Ensure that the files are decompressed (or checked out) in the installation folder.

2. From now on, you will need to run all the commands from the `python` folder. Install the necessary Python libraries running

```
pip install -r requirements.txt
```

3. LINE is now ready to use. For example, you can run a basic M/M/1 model using

```
python3 mm1.py
```

4. Jupyter notebooks are also available under the `examples` and `gettingstarted` folders.

To run line within your Python program, import the `line_solver` module at the beginning of the file, e.g.,

```
from line_solver import *
```

2.1.1 Software requirements

Certain features of LINE depend on external tools and libraries. The recommended dependencies are:

- Python version 3.10 or later, with the following packages:
 - `enum_tools`
 - `jpype1`
 - `numpy`
 - `pandas`
 - `matplotlib`
 - `scipy`
- Jupyter notebooks have been developed and tested using IntelliJ DataSpell 2024.1.1.

Partial Java ports of these libraries have been implemented or are automatically downloaded or shipped with LINE:

- Java Modelling Tools (<http://jmt.sf.net>): version 1.2.4 or later. The latest version is automatically downloaded at the first call of the JMT solver.
- KPC-Toolbox (<https://github.com/kpctoolboxteam/kpc-toolbox>): version 0.3.4 or later. This release is already included under the `lib` subfolder.
- M3A (<https://github.com/imperial-qore/M3A>): version 1.0.0. This release is already included under the `lib` subfolder.
- BuTools (<https://github.com/ghorvath78/butools>): version 2.0 or later. This release is already included under the `lib` subfolder.
- Q-MAM (<https://win.uantwerpen.be/~vanhoudt/tools/QBDfiles.zip>): This release is already included under the `lib` subfolder.

Optional dependencies recommended to utilize all features available in LINE are as follows:

- LQNS (<https://github.com/layeredqueuing/V6>): version 6.2.28 or later. System paths need to be configured such that the `lqns` and `lqnsim` solvers need are available on the command line (i.e., can be invoked from MATLAB via the `dos` or `unix` commands without need to specify the paths of the executables).

2.1.2 Documentation

This manual introduces the main concepts to define models in LINE and run its solvers. The document includes in particular several tables that summarize the features currently supported in the modeling language and by individual solvers. Additional resources are as follows:

- PDF versions of all manuals (Java, MATLAB, Python): <https://github.com/imperial-qore/line-solver/tree/main/doc>
- An online wiki version of the manual (MATLAB version): <https://github.com/line-solver/LINE/wiki>.
- Java API class hierarchy: <https://htmlpreview.github.io/?https://raw.githubusercontent.com/imperial-qore/line-solver/main/java/docs/javadoc/index.html>.

2.1.3 Getting help

For discussions, bug reports, new feature requests, please create a thread on one of the following Sourceforge boards:

- General discussion: <https://sourceforge.net/p/line-solver/discussion/help/>
- Bugs and issues: <https://sourceforge.net/p/line-solver/tickets/>
- Feature requests: <https://sourceforge.net/p/line-solver/feature-requests/>

2.2 Getting started examples

In this section, we present some examples that illustrate how to use LINE. The relevant scripts are included under the `gettingstarted` folder.

Systems can be described in LINE using one of the available classes of stochastic models:

- `Network` models are extended queueing networks. Typical instances are open, closed and mixed queueing networks, possibly including advanced features such as class-switching, finite capacity, priorities, non-exponential distributions, and others. Technical background on these models can be found in books such as [5, 34] or in tutorials such as [1, 33].
- `LayeredNetwork` models are layered queueing networks, i.e., models consisting of layers, each corresponding to a `Network` object, which interact through synchronous and asynchronous calls. Technical background on layered queueing networks can be found in [50].

The goal of the remainder of this chapter is to provide simple examples that explain the basics on how these models can be analyzed in LINE. More advanced forms of evaluation, such as probabilistic or transient analyses, are discussed in later chapters. Additional examples are supplied under the `examples` and `gallery` folders.

2.2.1 Controlling verbosity

Solver verbosity may be configured at program start using, e.g.:

```
GlobalConstants.setVerbose(VerboseLevel.DEBUG)
```

Alternative verbosity levels are `VerboseLevel.STD` and `VerboseLevel.SILENT`.

2.2.2 Model gallery

LINE includes a collection of classic, commonly occurring, queueing models under the `gallery` folder. They include single queueing systems (e.g., $M/M/1$, $M/H_2/1$, $D/M/1$, ...), tandem queueing systems, and basic queueing networks. For example, to instantiate and estimate the mean response time for a tandem network of $M/M/1$ queues we may run

```
SolverMVA(gallery_mm1_tandem(), 'method', 'exact').getAvgTable()
```

Obtaining the following pandas DataFrame

	Station	JobClass	QLen	Util	RespT	ResidT	ArvR	Tput
0	mySource	myClass	0.0000	0.0	0.0000	0.0000	0.0	1.0
1	Queue1	myClass	8.9992	0.9	8.9992	8.9992	1.0	1.0
2	Queue2	myClass	8.9992	0.9	8.9992	8.9992	1.0	1.0

The examples in the gallery may also be used as templates to accelerate the definition of basic models. Example 9 shows later an example of gallery instantiation of a $M/E_2/1$ queue.

2.2.3 Example 1: A M/M/1 queue

The $M/M/1$ queue is a classic model of a queueing system where jobs arrive into an infinite-capacity buffer, wait to be processed in first-come first-served (FCFS) order, and then leave after service completion. Arrival and service times are assumed to be independent and exponentially distributed random variables.

In this example, we wish to compute average performance measures for the $M/M/1$ queue. We assume that arrivals come in at rate $\lambda = 1$ job/s, while service has rate $\mu = 2$ job/s. It is known from theory that the exact value of the server utilization in this case is $\rho = \lambda/\mu = 0.5$, i.e., 50%, while the mean response time for a visit is $R = 1/(\mu - \lambda) = 1$ s. We wish to verify these values using JMT-based simulation, instantiated through LINE.

The general structure of a LINE script consists of four blocks:

1. Definition of nodes
2. Definition of job classes and associated statistical distributions
3. Instantiation of model topology

4. Solution

For example, the following script solves the M/M/1 model

```

model = Network("M/M/1 model")
# Block 1: nodes
source = Source(model, "Source")
queue = Queue(model, "Queue", SchedStrategy.FCFS)
sink = Sink(model, "Sink")
# Block 2: classes
jobclass = OpenClass(model, "Class1")
source.setArrival(jobclass, Exp(1.0))
queue.setService(jobclass, Exp(2.0))
# Block 3: topology
model.addLink(source, queue)
model.addLink(queue, sink)
# Block 4: solution
avgTable = SolverJMT(model, 'seed', 23000).getAvgTable() # pandas.DataFrame

```

In the example, `source` and `sink` are arrival and departure points of jobs; `queue` is a queueing station with FCFS scheduling; `jobclass` defines an open class of jobs that arrive, get served, and leave the system; `Exp(2.0)` defines an exponential distribution with rate parameter $\lambda = 2.0$; finally, the `getAvgTable` command solves for average performance measures with JMT's simulator, using for reproducibility a specific seed for the random number generator.

The result is a table with mean performance measures including: the number of jobs in the station either queueing or receiving service (`QLen`); the utilization of the servers (`Util`); the mean response time for a visit to the station (`RespT`); the mean residence time, i.e. the mean response time cumulatively spent at the station over all visits (`ResidT`); the mean throughput of departing jobs (`Tput`)

	Station	JobClass	QLen	Util	RespT	ResidT	Tput
0	Source	Class1	0.000000	0.00000	0.000000	0.000000	0.990016
1	Queue	Class1	0.950088	0.48791	0.967911	0.967911	0.997006

One can verify that this matches JMT results by first typing

```
model.jsimView()
```

which will open the model inside JSIMgraph, as shown in Figure 2.1. From this screen, the simulation can be started using the green “play” button in the JSIMgraph toolbar. A pre-defined gallery of classic models is also available, for example

```
model = gallery_mm1()
```

returns a M/M/1 queue with 50% utilization.

If we want to select a particular row of the `AvgTable` data structure, we can use the `tget` (table get) command, for example

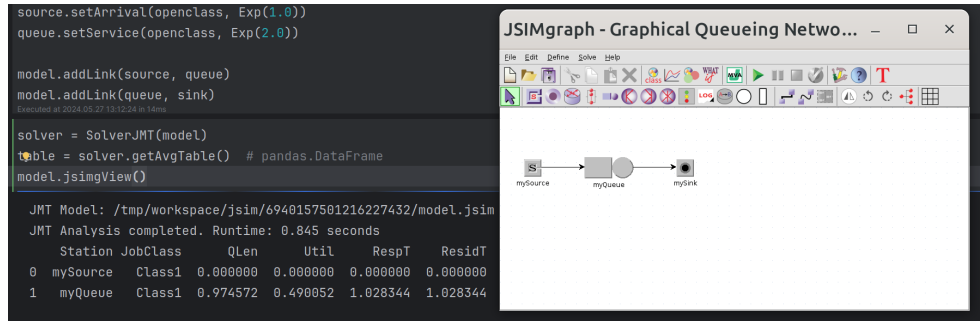


Figure 2.1: M/M/1 example in JSIMgraph launched from the DataSpell IDE

```
ARow = tget(AvgTable, 'Queue', 'Class1')
```

gives output

	Station	JobClass	QLen	Util	RespT	ResidT	Tput
1	Queue	Class1	0.955501	0.48736	0.954293	0.954293	0.999868

If we specify only 'Queue' or 'Class1', `tget` will return all entries corresponding to that station or class. Moreover, the following syntax is also valid

```
ARow = tget(AvgTable, queue, jobclass)
```

if we specify only queue or only jobclass, `tget` will return all entries corresponding to that station or class.

2.2.4 Example 2: A multiclass M/G/1 queue

We now consider a more challenging variant of the first example. We assume that there are two classes of incoming jobs with non-exponential service times. For the first class, service times are Erlang distributed with unit rate and variance $1/3$; they are instead read from a trace for the second class. Both classes have exponentially distributed inter-arrival times with mean $2s$.

To run this example, let us first change the working directory to the `examples` folder. Then we specify the node block

```
model = Network('M/G/1')
source = Source(model, 'Source')
queue = Queue(model, 'Queue', SchedStrategy.FCFS)
sink = Sink(model, 'Sink')
```

The next step consists in defining the classes. We fit automatically from mean and squared coefficient of variation (i.e., $SCV = \text{variance}/\text{mean}^2$) an Erlang distribution and use the `Replayer` distribution to request that the specified trace is read cyclically to obtain the service times of class 2

```
jobclass1 = OpenClass(model, 'Class1')
jobclass2 = OpenClass(model, 'Class2')

source.setArrival(jobclass1, Exp(0.5))
source.setArrival(jobclass2, Exp(0.5))

queue.setService(jobclass1, Erlang.fitMeanAndSCV(1, 1/3))
queue.setService(jobclass2, Replayer('example_trace.txt'))
```

Note that the `example_trace.txt` file consists of a single column of doubles, each representing a service time value, e.g.,

```
1.2377474e-02
4.4486055e-02
1.0027642e-02
2.0983173e-02
...
```

We now specify a linear route through source, queue, and sink for both classes

```
P = model.initRoutingMatrix()
P.set(jobclass1, Network.serialRouting(source, queue, sink))
P.set(jobclass2, Network.serialRouting(source, queue, sink))
model.link(P)
```

and solve the model with JMT

```
jmtAvgTable = SolverJMT(model, 'seed', 23000).getAvgTable()
```

which gives

```
JMT Model: /tmp/workspace/jsim/7955719514502154869/model.jsim
JMT Analysis completed. Runtime: 1.06 seconds
  Station JobClass   QLen   Util   RespT   ResidT   Tput
0 Source   Class1  0.000000  0.000000  0.000000  0.000000  0.501685
1 Source   Class2  0.000000  0.000000  0.000000  0.000000  0.506538
2 Queue    Class1  0.880661  0.494749  1.737321  1.737321  0.506679
3 Queue    Class2  0.438433  0.053432  0.817414  0.817414  0.507503
```

We wish now to validate this value against an analytical solver. Since `jobclass2` has trace-based service times, we first need to revise its service time distribution to make it analytically tractable, e.g., we may ask `LINE` to fit an acyclic phase-type distribution [4] based on the trace

```
queue.setService(jobclass2, Replayer('example_trace.txt').fitAPH())
```

We can now use a Continuous Time Markov Chain (CTMC) to solve the system, but since the state space is infinite in open models, we need to truncate it to be able to use this solver. For example, we may restrict to states with at most 2 jobs in each class, checking with the `verbose` option the size of the resulting state space

```
TODO: under development
```

However, we see from the comparison with JMT that the errors of `SolverCTMC` are rather large. Since the truncated state space consists of just 46 states, we can further increase the cutoff to 4, trading a slower solution time for higher precision

```
TODO: under development
```

To gain more accuracy, we could either keep increasing the cutoff value or, if we wish to compute an exact solution, we may call the matrix-analytic method (MAM) solver instead. `SolverMAM` uses the repetitive structure of the CTMC to exactly analyze open systems with an infinite state space, calling

```
SolverMAM(model).getAvgTable()
```

we get

	Station	JobClass	QLen	Util	RespT	ResidT	ArvR	Tput
0	Source	Class1	0.000000	0.000000	0.000000	0.000000	0.5	0.5
1	Source	Class2	0.000000	0.000000	0.000000	0.000000	0.5	0.5
2	Queue	Class1	0.876460	0.500000	1.752920	1.752920	0.5	0.5
3	Queue	Class2	0.426996	0.050536	0.853991	0.853991	0.5	0.5

The current MAM implementation is primarily constructed on top of Java ports of the BuTools solver [31] and the SMC solver [3].

2.2.5 Example 3: Machine interference problem

Closed models involve jobs that perpetually cycle within a network of queues. The machine interference problem is a classic example, in which a group of repairmen is tasked with fixing machines as they break and the goal is to choose the optimal size of the group. We here illustrate how to evaluate the performance of a given group size. We consider a scenario with $S = 2$ repairmen, with machines that break down at a rate of 0.5 failed machines/week, after which a machine is fixed in an exponential distributed time with rate 4.0 repaired machines/week. There are a total of $N = 3$ machines.

Suppose that we wish to obtain an exact numerical solution using Continuous Time Markov Chains (CTMCs). The above model can be analyzed as follows:

```
model = Network('MRP')
delay = Delay(model, 'WorkingState')
queue = Queue(model, 'RepairQueue', SchedStrategy.FCFS)
queue.setNumberOfServers(2)
```



```
cclass = ClosedClass(model, 'Machines', 3, delay)
delay.setService(cclass, Exp(0.5))
queue.setService(cclass, Exp(4.0))
model.link(Network.serialRouting(delay, queue))
```

TODO: under development

Here, `delay` appears in the constructor of the closed class to specify that a job will be considered completed once it returns to the delay (i.e., the machine returns in working state). We say that the delay is thus the *reference station* of `cclass`. The above code prints the following result

TODO: under development

As before, we can inspect and analyze the model in JSIMgraph using the command

```
model.jsimView()
```

Figure 2.2 illustrates the result, demonstrating the automated definition of the closed class.

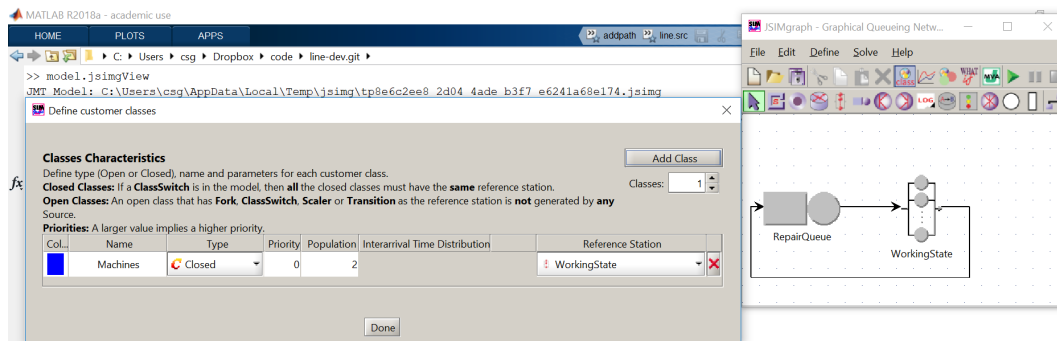


Figure 2.2: Machine interference model in JSIMgraph

We can now also inspect the CTMC more in the details as follows

TODO: under development

which produces in output the state space of the model and the infinitesimal generator of the CTMC

TODO: under development

For example, the first state $(0 \ 1 \ 2)$ consists of two components: the initial 0 denotes the number of jobs in service in the `delay`, while the remaining part is the state of the FCFS queue. In the latter, the 1 means that a job of class 1 (the only class in this model) is in the waiting buffer, while the 2 means that there are two jobs in service at the queue.

As another example, the second state (1 0 2) is similar, but one job has completed at the queue and has then moved to the `delay`, concurrently triggering an admission in service for the job that was in the queue buffer. As a result of this, the buffer is now empty. The corresponding transition rate in the infinitesimal generator matrix is row 1 and column 2 of `InfGen`, which has value 8.0, that is the sum of the completion rates at the queue for each server in the first state, and where indexes 1 and 2 are the rows in `StateSpace` associated to the source and destination states.

On this and larger infinite generators, we may also list individual non-zero transitions as follows

```
TODO: under development
```

The above printout helps in matching the state transitions to their rates.

To avoid having to inspect the `StateSpace` variable to determine to which station a particular column refers to, we can alternatively use the more general invocation

```
TODO: under development
```

which automatically splits the state space into its constituent parts for each stateful node.

A further observation is that `model.getStateSpace()` forces the regeneration of the state space at each invocation, whereas the equivalent function in the CTMC solver, `solver.getStateSpace()`, returns the state space cached during the solution of the CTMC.

2.2.6 Example 4: Round-robin load-balancing

In this example we consider a system of two parallel processor-sharing queues and we wish to study the effect of load-balancing on the average performance of an open class of jobs. We begin as usual with the node block, where we now include a special node, called the `Router`, to control the routing of jobs from the source into the queues:

```
model = Network('RRLB')
source = Source(model, 'Source')
lb = Router(model, 'LB')
queue1 = Queue(model, 'Queue1', SchedStrategy.PS)
queue2 = Queue(model, 'Queue2', SchedStrategy.PS)
sink = Sink(model, 'Sink')
```

Let us then define the class block by setting exponentially-distributed inter-arrival times and service times, e.g.,

```
jobclass = OpenClass(model, 'Class1')
source.setArrival(jobclass, Exp(1.0))
queue1.setService(jobclass, Exp(2.0))
queue2.setService(jobclass, Exp(2.0))
```

We now wish to express the fact that the router applies a round-robin strategy to dispatch jobs to the queues. Since this is now a non-probabilistic routing strategy, we need to adopt a slightly different style to declare

the routing topology as we cannot specify anymore routing probabilities. First, we indicate the connections between the nodes, using the `addLinks` function:

```
model.addLinks([[source, lb],
               [lb, queue1],
               [lb, queue2],
               [queue1, sink],
               [queue2, sink]])
lb.setRouting(jobclass, RoutingStrategy.RAND)
```

At this point, all nodes are automatically configured to route jobs with equal probabilities on the outgoing links (`RoutingStrategy.RAND` policy). If we solve the model at this point, we see that the response time at the queues is around 0.66s. Running JMT

```
jmtAvgTable = SolverJMT(model, 'seed', 23000).getAvgTable()
```

we get

Station	JobClass	QLen	Util	RespT	ResidT	ArvR	Tput
0	Source	Class1	0.000000	0.000000	0.000000	0.000000	1.013486
1	Queue1	Class1	0.316119	0.246825	0.654111	0.327056	0.500997
2	Queue2	Class1	0.334030	0.250757	0.684064	0.342032	0.504135

After resetting the internal data structures, which is required before modifying a model we can require `LINE` to solve again the model using this time a round-robin policy at the router.

```
model.reset()
lb.setRouting(oclass, RoutingStrategy.RROBIN)
```

A representation of the model at this point is shown in Figure 2.3.

Lastly, we run again JMT and find that round-robin produces a visible decrease in response times, which are now around 0.56s.

```
SolverJMT(model, 'seed', 23000).getAvgTable()
```

gives

Station	JobClass	QLen	Util	RespT	ResidT	ArvR	Tput
0	Source	Class1	0.000000	0.000000	0.000000	0.000000	1.008868
1	Queue1	Class1	0.304291	0.261181	0.584815	0.292408	0.505255
2	Queue2	Class1	0.292822	0.243971	0.572931	0.286466	0.505264

2.2.7 Example 5: Modelling a re-entrant line

Let us now consider a simple example inspired to the classic problem of modeling *re-entrant lines*. This arises in manufacturing systems where parts (i.e., jobs) re-enter multiple times a machine (i.e., a queueing

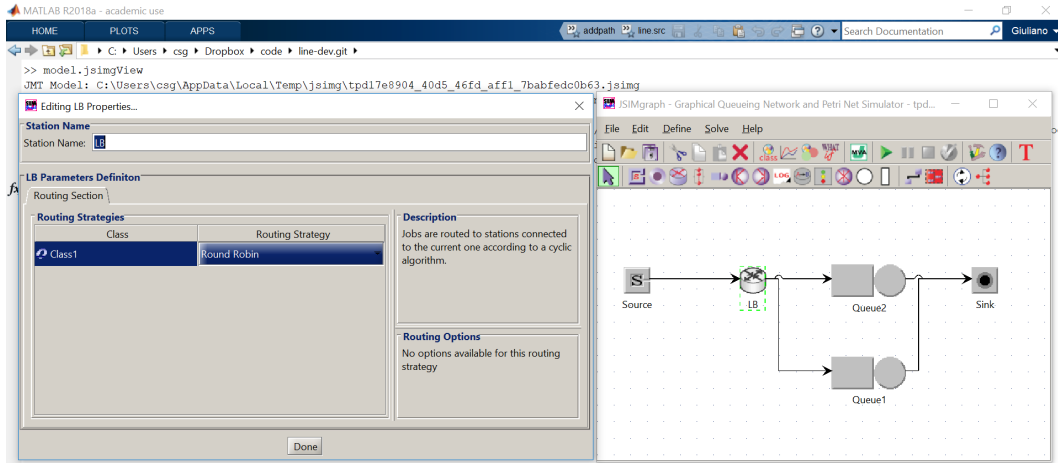


Figure 2.3: Load-balancing model

station), asking at each visit a different class of service. This implies, for example, that the service time at every visit could feature a different mean or a different distribution compared to the previous visits, thus modeling a different stage of processing.

To illustrate this, consider for example a degenerate model composed by a single FCFS queue and K classes. In this model, a job that completes processing in class k is routed back at the tail of the queue in class $k + 1$, unless $k = K$ in which case the job re-enters in class 1.

We take the following assumptions: $K = 3$ and class k has an Erlang-2 service time distribution at the queue with mean equal to k ; the system starts with $N_1 = 1$ jobs in class 1 and zero jobs in all other classes.

```

model = Network('RL')
queue = Queue(model, 'Queue', SchedStrategy.FCFS)
K = 3
N = (1, 0, 0)
jobclass = []
for k in range(K):
    jobclass.append(ClosedClass(model, 'Class' + str(k), N[k], queue))
    queue.setService(jobclass[k], Erlang.fitMeanAndOrder(1+k, 2))

P = model.initRoutingMatrix()
P.set(jobclass[0], jobclass[1], queue, queue, 1.0)
P.set(jobclass[1], jobclass[2], queue, queue, 1.0)
P.set(jobclass[2], jobclass[0], queue, queue, 1.0)
model.link(P)

```

The corresponding JMT model is shown in Figure 2.4, where it can be seen that the class-switching rule is automatically enforced by introduction of a `ClassSwitch` node in the network.

We can now simulate the performance indexes for the different classes, for example using LINE's normal-




Figure 2.4: Re-entrant lines as an example of class-switching

izing constant solver (SolverNC)

```
ncAvgTable = SolverNC(model).getAvgTable()
```

gives

	Station	JobClass	QLen	Util	RespT	ResidT	ArvR	Tput
0	Queue	Class1	0.166667	0.166667	1.0	0.333333	0.166667	0.166667
1	Queue	Class2	0.333333	0.333333	2.0	0.666667	0.166667	0.166667
2	Queue	Class3	0.500000	0.500000	3.0	1.000000	0.166667	0.166667

Suppose now that the job is considered completed, for the sake of computation of system performance metrics, only when it departs the queue in class K (here Class3). By default, LINE will return *system-wide* performance metrics using the `getAvgSysTable` method, i.e.,

```
new SolverNC(model).getAvgSysTable().print();
```

gives

	Chain	JobClasses	SysRespT	SysTput
0	Chain1	(Class1 Class2 Class3)	2.0	0.5

This method identifies the model *chains*, i.e., groups of classes that can exchange jobs with each other, but not with classes in other chains. Since the job can switch into any of the three classes, in this model there is a single chain comprising the three classes.

We see that the throughput of the chain is 0.5, which means that LINE is counting every departure from the queue in any class as a completion for the whole chain. This is incorrect for our model since we want to count completions only when jobs depart in Class3. To require this behavior, we can tell to the solver that passages for classes 1 and 2 through the reference station should not be counted as completions

```
jobclass[0].completes = False
jobclass[1].completes = False
```

This modification then gives the correct chain throughput, matching the one of Class3 alone

```
ncAvgSysTable = SolverNC(model).getAvgSysTable()
```

gives

Chain	JobClasses	SysRespT	SysTput
0 Chain1	(Class1 Class2 Class3)	6.0	0.166667

2.2.8 Example 6: A queueing network with caching

In this more advanced example, we show how to include in a queueing network a cache adopting a least-recently used (LRU) replacement policy. Under LRU, upon a cache miss the least-recently accessed item will be discarded to make room for the newly requested item.

We consider a cache with a capacity of 50 items, out of a set of 1000 cacheable items. Items are accessed by jobs visiting the cache according to a Zipf-like law with exponent $\alpha = 1.4$ and defined over the finite set of items. A client cyclically issues requests for the items, waiting for a reply before issuing the next request. We assume that a cache hit takes on average $0.2ms$ to process, while a cache miss takes $1ms$. We ask for the average request throughput of the system, differentiated across hits and misses.

Node block As usual, we begin by defining the nodes. Here a delay node will be used to describe the time spent by the requests in the system, while the cache node will determine hits and misses:

```
model = Network('model')
clientDelay = Delay(model, 'Client')
cacheNode = Cache(model, 'Cache', 1000, 50, ReplacementStrategy.LRU)
cacheDelay = Delay(model, 'CacheDelay')
```

Class block We define a set of classes to represent the incoming requests (`clientClass`), cache hits (`hitClass`) and cache misses (`missClass`). These classes need to be closed to ensure that there is a single outstanding request from the client at all times:

```
clientClass = ClosedClass(model, 'ClientClass', 1, clientDelay, 0)
hitClass = ClosedClass(model, 'HitClass', 0, clientDelay, 0)
missClass = ClosedClass(model, 'MissClass', 0, clientDelay, 0)
```

We then assign the processing times, using the `Immediate` distribution to ensure that the client issues immediately the request to the cache:

```
clientDelay.setService(clientClass, Immediate())
cacheDelay.setService(hitClass, Exp.fitMean(0.2))
cacheDelay.setService(missClass, Exp.fitMean(1.0))
```

The next step involves specifying that the request uses a Zipf-like distribution (with parameter $\alpha = 1.4$) to select the item to read from the cache, out of a pool of 1000 items

```
cacheNode.setRead(clientClass, Zipf(1.4, 1000))
```

Finally, we ask that the job should become of class `hitClass` after a cache hit, and should become of class `missClass` after a cache miss:

```
cacheNode.setHitClass(clientClass, hitClass)
cacheNode.setMissClass(clientClass, missClass)
```

Topology block Next, in the topology block we setup the routing so that the request, which starts in `clientClass` at the `clientDelay`, then moves from there to the cache, remaining in `clientClass`

```
P = model.initRoutingMatrix()
P.set(clientClass, clientClass, clientDelay, cacheNode, 1.0)
```

Internally to the cache, the job will switch its class into either `hitClass` or `missClass`. Upon departure in one of these classes, we ask it to join in the same class `cacheDelay` for further processing

```
P.set(hitClass, hitClass, cacheNode, cacheDelay, 1.0)
P.set(missClass, missClass, cacheNode, cacheDelay, 1.0)
```

Lastly, the job returns to `clientDelay` for completion and start of a new request, which is done by switching its class back to `clientClass`

```
P.set(hitClass, clientClass, cacheDelay, clientDelay, 1.0)
P.set(missClass, clientClass, cacheDelay, clientDelay, 1.0)
```

The above routing strategy is finally applied to the model

```
model.link(P)
```

Solution block To solve the model, since JMT does not support cache modeling, we use the native MATLAB-based simulation engine provided within LINE, the SSA solver:

```
TODO: under development
```

The above script produces the following result

```
TODO: under development
```

The departing flows from the `CacheDelay` are the miss and hit rates. Thus, the hit rate is 2.4554 jobs per unit time, while the miss rate is 0.50892 jobs per unit time.

Let us now suppose that we wish to verify the result with a longer simulation, for example with 10 times more samples. To this aim, we can use the automatic parallelization of SSA

```
TODO: under development
```

This gives us a rather similar result, when run on a dual-core machine

```
TODO: under development
```

The execution time is longer than usual at the first invocation of the parallel solver due to the time needed by MATLAB to bootstrap the parallel pool, in this example around 22 seconds. Successive invocations of parallel SSA normally take much less, with this example around 7 seconds each.

2.2.9 Example 7: Response time distribution and percentiles

In this example we illustrate the computation of response time percentiles in a queueing network model. We begin by instantiating a simple closed model consisting of a delay followed by a processor-sharing queueing station.

```
model = Network("Model")

node = np.empty(2, dtype=object)
node[0] = Delay(model, 'Delay')
node[1] = Queue(model, 'Queue1', SchedStrategy.PS)
```

There is a single class consisting of 5 jobs that circulate between the two stations, taking exponential service times at both.

```
jobclass = np.empty(2, dtype=object)
jobclass[0] = ClosedClass(model, 'Class1', 5, node[0], 0)
node[0].setService(jobclass[0], Exp(1.0))
node[1].setService(jobclass[0], Exp(0.5))

model.link(Network.serialRouting(node[0], node[1]))
```

We now wish to compare the response time distribution at the PS queue computed analytically with a fluid approximation against the simulated values returned by JMT. To do so, we call the `getCdfRespT` method

```
TODO: under development
```

```
# RDfluid = SolverFluid(model).getCdfRespT()
RDsim = SolverJMT(model, 'seed', 23000, 'samples', 10000).getCdfRespT()
```


TODO . The first column represents the cumulative distribution function (CDF) value $F(t) = Pr(T \leq t)$, where T is the random variable denoting the response time, while t is the percentile appearing in the corresponding entry of the second column.

For example, to plot the complementary CDF $1 - F(t)$ we can use the following code

```
TODO: under development
```

which produces the graph shown in Figure ??

```
TODO: under development
```

The graph shows that, although the simulation refers to a transient, while the fluid approximation refers to steady-state, there is a tight matching between the two response time distributions.

We can also readily compute the percentiles from the `RDfluid` and `RDsim` data structures, e.g., for the 95th and 99th percentiles of the simulated distribution

```
TODO: under development
```

That is, 95% of the response times at the PS queue (node 2, class 1) are less than or equal to 27.0222 time units, while 99% are less than or equal to 41.8743 time units.

2.2.10 Example 8: Optimizing a performance metric

In this example, we show how to optimize with the help of `LINE` a performance metric. We wish to find the optimal routing probabilities that minimize average response times for two parallel processor sharing queues. We assume that jobs are fed by a delay station, arranged with the two queues in a closed network topology.

We first define a Python function with header

```
def objFun(p):
```

Within the function definition, we instantiate the two queues and the delay station

```
model = Network('LoadBalCQN')
delay = Delay(model, 'Think')
queue1 = Queue(model, 'Queue1', SchedStrategy.PS)
queue2 = Queue(model, 'Queue2', SchedStrategy.PS)
```

We assume that 16 jobs circulate among the nodes, and that the service rates are $\sigma = 1$ jobs per unit time at the delay, and $\mu_1 = 0.75$ and $\mu_2 = 0.50$ at the two queues:

```
cclass = ClosedClass(model, 'Job1', 16, delay)
delay.setService(cclass, Exp(1))
queue1.setService(cclass, Exp(0.75))
queue2.setService(cclass, Exp(0.50))
```

We initially setup a topology with arbitrary values for the routing probabilities between delay and queues, ensuring that jobs completing at the queues return to the delay:

```
P = model.initRoutingMatrix()
P.set(cclass, cclass, queue1, delay, 1.0)
P.set(cclass, cclass, queue2, delay, 1.0)
model.link(P)
```

We now return the system response time for the jobs as a function of the routing probability p to choose queue 1 instead of queue 2:

```
P.set(cclass, cclass, delay, queue1, p)
P.set(cclass, cclass, delay, queue2, 1.0 - p)
model.reset()
model.link(P)
# Block 4: solution
R = SolverMVA(model, 'method', 'exact', 'verbose', False).getAvgSysRespT()
return R[0]
```

Lastly, we optimize the function we defined

```
p_opt = optimize.fminbound(objFun, 0, 1)
print(p_opt[0])
```

We are now ready to run the example. The execution returns the optimal value 0.6104878504366782.

2.2.11 Example 9: Studying a departure process

This examples illustrates LINE's support for extracting simulation data about particular events in an extended queueing network, such as departures from a particular queue.

Our goal is to obtain the squared coefficient of variation of the inter-departure times from a $M/E_2/1$ queue, which has Poisson arrivals and 2-phase Erlang distributed service times.

Because this is a classic model, we can find it in LINE's model gallery. The additional return parameters (e.g., `source`, `queue`, ...) provide handles to the entities within the model.

```
TODO: under development
```

We now extract 50,000 samples from simulation based on the underpinning continuous-time Markov chain

```
TODO: under development
```

The returned data structure supplies information about the stateful nodes (here `source` and `queue`) at each of the 50,000 instants of sampling, together with the events that have been collected at these instants.

```
TODO: under development
```

As an example, the first two events occur both at timestamp 0 and indicate a departure event from node 1 (the type `EventType.DEP` maps to `event : DEP`) followed by an arrival event at node 2 (the type `EventType.ARV` maps to `event : ARV`) which accepts it always (`prob : 1`).

```
TODO: under development
```

We may also plot the first 300 events as follows

```
TODO: under development
```

We are now ready to filter the timestamps of events related to departures from the `queue` node

```
TODO: under development
```

Followed by a calculation of the time series of inter-departure times

```
TODO: under development
```

We may now for example compute the squared coefficient of variation of this process

```
TODO: under development
```

which evaluates to 0.8750. Using Marshall's exact formula for the $GI/G/1$ queue [35], we get a theoretical value of 0.8750.

2.2.12 Example 10: Evaluating a CTMC symbolically

In this example, we will use MATLAB's symbolic toolbox to examine the continuous-time Markov chain underlying a simple closed queueing network. The network consists of a delay station and a processor sharing station arranged in a cyclic topology. We may generate it using LINE's demo gallery as follows

```
TODO: under development
```

Here, the first argument adds a single station to the next, while the second argument requires presence of a delay station. The network has a single class with 4 circulating jobs.

The `getSymbolicGenerator` method of the CTMC solver can now be called to obtain the symbolic generator

```
TODO: under development
```

The first returned argument is the symbolic infinitesimal generator

```
TODO: under development
```

There are therefore 5 states, corresponding to all possible way of distributing the 4 jobs across the two stations

```
TODO: under development
```

An event is represented in LINE as a synchronization between an active agent and a passive agent. Typically, the station that completes a job is an active agent, whereas the one that receives it is a passive agent. In this sense, x_1 and x_2 may be seen as the rates at which the two agents synchronize to perform the two actions.

To learn the meaning of the symbolic variables x_1 and x_2 we can now use the `syncInfo` data structure

```
TODO: under development
```

In the above, we see that x_1 is a class-1 departure from station 1 (they delay) into station 2 (the processor sharing queue), and viceversa x_2 is a departure from station 2 that joins station 1.

```
TODO: under development
```

Chapter 3

Network models

Throughout this chapter, we discuss the specification of `Network` models, which are extended queueing networks. `LINE` currently supports open, closed, and mixed networks with non-exponential service and arrivals, and state-dependent routing. All solvers support the computation of basic performance metrics, while some more advanced features are available only in specific solvers. Each `Network` model requires in input a description of the nodes, the network topology, and the characteristics of the jobs that circulate within the network. In output, `LINE` returns performance and reliability metrics.

The default metrics supported by all solvers are as follows:

- Mean queue-length (`QLen`). This is the mean number of jobs residing at a node when this is observed at a random instant of time.
- Mean utilization (`Util`). For nodes that serve jobs, this is the mean fraction of time the node is busy processing jobs. In both single-server and multi-server nodes, this is a number normalized between 0 and 1, corresponding to 0% and 100%. In infinite-server nodes, the utilization is set by convention equal to the mean queue-length, therefore taking the interpretation of the mean number of jobs in execution at the station.
- Mean response time (`RespT`). This is the mean time a job spends traversing a node within a network. If the node is visited multiple times, the response time is the time spent for a single visit to the node.
- Mean residence time (`ResidT`). This is the total time a job accumulates, on average, to traverse a node within a network. If the node is visited multiple times, the residence time is the time accumulated overall visits to the node prior to returning to the reference station or arriving to a sink.
- Mean throughput (`Tput`). This is the mean departure rate of jobs completed at a resource per time unit. Typically, this matches the mean arrival rate, unless the node switches the class of the jobs in which case the arrival rate of a class may not match its departure rate.

The above metrics refer to the performance characteristics of individual nodes. Response times and throughputs can also be system-wide, meaning that they can describe end-to-end performance during the visit to the network. In this case, these metrics are called *system* metrics.

3.1 Network object definition

3.1.1 Creating a network and its nodes

A queueing network can be described in LINE using the `Network` class constructor with a unique string identifying the model name:

```
model = Network('myModel')
```

The returned object of the `Network` class offers functions to instantiate and manage resource *nodes* (stations, delays, caches, ...) visited by jobs of several types (*classes*).

A node is a resource in the network that can be visited by a job. A node must have a unique name and can either be *stateful* or *stateless*, the latter meaning that the node does not require state variables to determine its state or actions. If jobs visiting a stateful node can be required to spend time in it, the node is also said to be a *station*. A list of nodes available in `Network` models is given in Table 3.1.1.

Table 3.1: Nodes available in `Network` models.

Node	Description
Cache	A node to switch job classes based on hits/misses in its cache
ClassSwitch	A node to switch job classes based on a static probability matrix
Delay	A station where jobs spend time without queueing
Fork	A node that forks jobs into tasks
Join	A node that joins sibling tasks into the original job
Logger	A node that logs passage of jobs
Queue	A node where jobs queue and receive service
Router	A node that routes jobs to other nodes
Sink	Exit point for jobs in open classes
Source	Entry point for jobs in open classes

We now provide more details on each of the nodes available in `Network` models.

Queue node. A `Queue` specifies a queueing station from its name and scheduling strategy, e.g.

```
queue = Queue(model, 'Queue1', SchedStrategy.FCFS)
```

specifies a first-come first-served queue. It is alternatively possible to instantiate a queue using the `QueueingStation` constructor, which is merely an alias for `Queue`.

Queueing stations have by default a single server. The `setNumberOfServers` method can be used to instantiate multi-server stations.

Valid scheduling strategies are specified within the `SchedStrategy` static class and include:

- First-come first-serve (`SchedStrategy.FCFS`)
- Infinite-server (`SchedStrategy.INF`)
- Processor-sharing (`SchedStrategy.PS`)
- Service in random order (`SchedStrategy.SIRO`)
- Discriminatory processor-sharing (`SchedStrategy.DPS`)
- Generalized processor-sharing (`SchedStrategy.GPS`)
- Shortest expected processing time (`SchedStrategy.SEPT`)
- Shortest job first (`SchedStrategy.SJF`)
- Head-of-line priority (non-preemptive) (`SchedStrategy.HOL`)
- Polling (`SchedStrategy.POLLING`)

If a strategy requires class weights, these can be specified directly as an argument to the `setService` function or using the `setStrategyParam` function, see later the description of DPS scheduling for an example.

Delay node. Delay stations, also called infinite server stations, may be instantiated either as objects of `Queue` class, with the `SchedStrategy.INF` scheduling strategy, or using the following specialized constructor

```
delay = Delay(model, 'ThinkTime')
```

As for queues, for readability it is possible to instantiate delay nodes using the `DelayStation` class which is an alias for the `Delay` class.

Source and Sink nodes. As seen in the M/M/1 getting started example, these nodes are mandatory elements for the specification of open classes. Their constructor only requires a specification of the unique name associated to the nodes:

```
source = Source(model, 'Source')  
sink = Sink(model, 'Sink')
```

Fork and Join nodes. The fork and join nodes are currently available only for the JMT solver. The `Fork` splits an incoming job into a set of sibling tasks, sending out one task for each outgoing link. These tasks inherit the class of the original job and are served as normal jobs until they are reassembled at a `Join` station.

Their specification of Fork and Join nodes only requires the name of the node

```
fork = Fork(model, 'Fork')
join = Join(model, 'Join', fork)
```

The number of tasks sent by a `Fork` on each output link can be set using the `setTasksPerLink` method of the `fork` object. To enable effective analytical approximations, presently LINE requires that every join node is bound to a specific fork node, although specific solvers will ignore this information (e.g., JMT).

Also note that the routing probabilities out of the `Fork` node need to be set to 1.0 towards every other node connected to the `Fork`. For example, a `Fork` sending jobs in class 1 to nodes *A*, *B* and *C*, cannot send jobs in class 2 only to *A* and *B*: it must send them to all three connected nodes *A*, *B* and *C*. A new fork node visited only by class-2 jobs needs to be created in order to send that class of jobs only to *A* and *B*.

After splitting a job into tasks, LINE takes the convention that visit counts refer to the average number of passages at the target resources for the original job, scaled by the number of tasks. For example, if a job is split into two tasks at a fork node, each visiting respectively nodes *A* and *B*, the average visit count at *A* and *B* will be 0.5.

ClassSwitch node. This is a stateless node to change the class of a transiting job based on a static probabilistic policy. For example, it is possible to specify that all jobs belonging to class 1 should become of class 2 with probability 1.0, or that a transiting job of class 2 should become of class 1 with probability 0.3. This example is instantiated as follows

```
cs = ClassSwitch(model, 'ClassSwitchPoint', [[0.0, 1.0], [0.3, 0.7]])
```

Cache node. This is a stateful node to store one or more items in a cache of finite size, for which it is possible to specify a replacement policy. The `cache` constructor requires the total cache capacity and the number of items that can be referenced by the jobs in transit, e.g.,

```
cacheNode = Cache(model, 'Cache1', nitems, capacity, ReplacementStrategy.LRU)
```

If the capacity is a scalar integer (e.g., 15), then it represents the total number of items that can be cached and the value cannot be greater than the number of items. Conversely, if it is a vector of integers (e.g., [10,5]) then the node is a list-based cache, where the vector entries specify the capacity of each list. We point to [26] for more details on list-based caches and their replacement policies.

Available replacement policies are specified within the `ReplacementStrategy` static class and include:

- First-in first-out (`ReplacementStrategy.FIFO`)
- Random replacement (`ReplacementStrategy.RR`)
- Least-recently used (`ReplacementStrategy.LRU`)
- Strict first-in first-out (`ReplacementStrategy.SFIFO`)

Upon cache hit or cache miss, a job in transit is switched to a user-specified class. More details are given later in Section 3.1.5.

Router node. This node is able to route jobs according to a specified `RoutingStrategy`, which can either be probabilistic or not (e.g., round-robin). Upon entering a `Router`, a job neither waits nor receives service; it is instead directly forwarded to the next node according to the specified routing strategy. A `Router` can be instantiated as follows:

```
router = Router(model, 'RouterNode')
```

An example of use of this node is given in Section 2.2.6. Routing strategies need to be specified for each class using the `setRouting` method and valid choices are as follows

- Random routing (`RoutingStrategy.RAND`)
- Round robin (`RoutingStrategy.RROBIN`)
- Probabilistic routing (`RoutingStrategy.PROB`)
- Join-the-shortest-queue (`RoutingStrategy.JSQ`)

For example, assume that `oclass` is a class of jobs. In order to route jobs in this class with equal probabilities to every outgoing link we set

```
router.setRouting(oclass, RoutingStrategy.RAND)
```

It should be noted that `setRouting` is also available for all other nodes such as queueing stations, delays, etc. Therefore, the added value of the `Router` node is the ability to represent certain system elements that centralize the routing logic, such as load balancers.

Logger node. A logger node is a node that closely resembles the logger node available in the `JSIMgraph` simulator within `JMT`. At present, models that include this element can only be solved using the `JMT` solver.

A `Logger` node records information about passing jobs in a `csv` file, such as the timestamp of passage and general information about the jobs. The node can be instantiated as follows

```
TODO: under development
```

The routing behavior of jobs can be set up as explained for regular nodes such as queues or delay stations.

3.1.2 Advanced node parameters

Scheduling parameters

Upon setting service distributions at a station, one may also specify scheduling parameters such as weights as additional arguments to the `setService` function. For example, if the node implements discriminatory processor sharing (`SchedStrategy.DPS`), the command

```
queue.setService(class2, Cox2.fitMeanAndSCV(0.2,10), 5.0)
```

assigns a weight 5.0 to jobs in class 2. The default weight of a class is 1.0.

Finite buffers

The functions `setCapacity` and `setChainCapacity` of the `Station` class are used to place constraints on the number of jobs, total or for each chain, that can reside within a station. Note that LINE does not allow one to specify buffer constraints at the level of individual classes unless chains contain a single class, in which case `setChainCapacity` is sufficient for the purpose.

For example,

```
TODO: under development
```

creates an example model with two chains and three classes (specified in `example_closedModel_3.m`) and requires the second station to accept a maximum of one job in each chain. Note that if we were to ask for a higher capacity, such as `setChainCapacity([1, 7])`, which exceeds the total job population in chain 2, LINE would have automatically reduced the value 7 to the chain 2 job population (2). This automatic correction ensures that functions that analyze the state space of the model do not generate unreachable states.

The `refreshCapacity` function updates the buffer parameterizations, performing appropriate sanity checks. Since `example_closedModel_3` has already invoked a solver prior to our changes, the requested modifications are materially applied by LINE to the network only after calling an appropriate `refreshStruct` function, see the sensitivity analysis section. If the buffer capacity changes were made before the first solver invocation on the model, then there would not be the need for a `refreshCapacity` call, since the internal representation of the `Network` object used by the solvers is still to be created.

```
TODO: under development
```

3.1.3 Job classes

Jobs travel within the network placing service demands at the stations. The demand placed by a job at a station depends on the class of the job. Jobs in *open classes* arrive from the external world and, upon completing the visit, leave the network. Jobs in *closed classes* start within the network and are forbidden to ever leave it, perpetually cycling among the nodes.

Open classes

The constructor for an open class only requires the class name and the creation of special nodes called `Source` and `Sink`

```
source = Source(model, 'Source')
sink = Sink(model, 'Sink')
```

Sources are special stations holding an infinite pool of jobs and representing the external world. Sinks are nodes that route a departing job back into this infinite pool, i.e., into the source. Note that a network can include at most a single `Source` and a single `Sink`.

Once source and sink are instantiated in the model, it is possible to instantiate open classes using

```
class1 = OpenClass(model, 'Class1')
```

LINE does not require explicitly associating source and sink with the open classes in their constructors, as this is done automatically. However, the LINE language requires explicitly creating these nodes since the routing topology needs to indicate the arrival and departure points of jobs in open classes. However, if the network does not include open classes, the user will not need to instantiate a `Source` and a `Sink`.

Closed classes

To create a closed class, we need instead to indicate the number of jobs that start in that class (e.g., 5 jobs) and the *reference station* for that class (e.g., `queue`), i.e.:

```
class2 = ClosedClass(model, 'Class2', 5, queue)
```

The reference station indicates a point in the network used to calculate certain performance indexes, called *system performance indexes*. The end-to-end response time for a job in an open class to traverse the system is an example of a system performance index (system response time). The reference station of an open class is always automatically set by LINE to be the `Source`. Conversely, the reference station needs to be indicated explicitly in the constructor for closed classes since the point at which a class job completes execution depends on the semantics of the model.

LINE also supports a special class of jobs, called *self-looping jobs*, which perpetually loop at the reference station, remaining in their class. The following example shows the syntax to specify a self-looping job, which is identical to closed classes but there is no need later to specify routing information.

```
model = Network('model')
delay = Delay(model, 'Delay')
queue = Queue(model, 'Queue1', SchedStrategy.FCFS)
cclass = ClosedClass(model, 'Class1', 10, delay, 0)
slclass = SelfLoopingClass(model, 'SLC', 1, queue, 0)
delay.setService(cclass, Exp(1.0))
queue.setService(cclass, Exp(1.5))
```

```

queue.setService(slclass, Exp(1.5))
P = model.initRoutingMatrix()
P[0] = [[0.7, 0.3], [1.0, 0.0]]
model.link(P)

```

Note that any routing information specified for the self-looping class will be ignored.

Mixed models

LINE also accepts models where a user has instantiated both open and closed classes. The only requirement is that, if two classes communicate by means of a class-switching mechanism, then the two classes must either be all closed or all open. In other words, classes in the same chain must either be both closed or open. Furthermore, for all closed classes in the same chain, it is required for the reference station to be the same.

Class priorities

If a class has a priority, with 0 representing the highest priority, this can be specified as an additional argument to both `OpenClass` and `ClosedClass`, e.g.,

```

class2 = ClosedClass(model, 'Class2', 5, queue, 0)

```

In `Network` models, priorities are intended as hard priorities and the only supported priority scheduling strategy (`SchedStrategy.HOL`) is non-preemptive. Weight-based policies such as DPS and GPS may be used, as an alternative, to prevent starvation of jobs in low-priority classes.

Class switching

In LINE, jobs can switch their class while they travel between nodes (including self-loops on the same node). For example, this feature can be used to model queueing properties such as re-entrant lines in which a job visiting a station a second time may require a different average service demand than at its first visit.

A chain defines the set of reachable classes for a job that starts in the given class r and over time changes class. Since class switching in LINE does not allow a closed class to become open, and vice-versa, chains can themselves be classified into *open chains* and *closed chains*, depending on the classes that compose them.

Jobs in open classes can only switch to another open class. Similarly, jobs in closed classes can only switch to a closed class. Thus, class switching from open to closed classes (or vice-versa) is forbidden. More details about class-switching are given in Section [3.1.5](#).

Reference station

Before we have shown that the specification of classes requires choosing a reference station. In LINE, reference stations are properties of chains, thus if two closed classes belong to the same chain they must

have the same reference station. This avoids ambiguities in the definition of the completion point for jobs within a chain.

For example, the system throughput for a chain is defined as the sum of the arrival rates at the reference station for all classes in that chain. That is, the solver counts a return to the reference station as a completion of the visit to the system. In the case of open chains, the reference station is always the `Source` and the system throughput corresponds to the rate at which jobs arrive at the sink `Sink`, which may be seen as the arrival rate seen by the infinite pool of jobs in the external world. If there is no class switching, each chain contains a single class, thus per-chain and per-class performance indexes will be identical.

Reference class

Occasionally, it is possible to encounter situations where a job needs to change class while remaining inside the same station. In this case, `LINE` modifies the network automatically to introduce a class-switching node for the job to route out of the station and immediately return to it in the new class.

One complication of the approach is that, by departing the node and returning to it, the job visits the station one additional time, affecting the visit count to the station and therefore performance metrics such as the residence time. To cope with this issue, `LINE` offers a method for the class objects, called `setReferenceClass`, that allows users to specify whether the visit of that class to the reference station should be considered upon computing the residence times across the network for the chain to which the class belongs. By default, all classes traversing the reference station are used in the visit count calculation.

3.1.4 Routing strategies

Probabilistic routing

Jobs travel between nodes according to the network topology and a routing strategy. Typically a queueing network will use a probabilistic routing strategy (`RoutingStrategy.PROB`), which requires specifying routing probabilities among the nodes. The simplest way to specify a large routing topology is to define the routing probability matrix for each class, followed by a call to the `link` function. This function will automatically add certain nodes to the network to ensure the correct class switching for jobs moving between stations (`ClassSwitch` elements).

In the running case, we may instantiate a routing topology as follows:

```
P = model.initRoutingMatrix()
P.set(class1, class1, source, queue, 1.0)
P.set(class1, class1, queue, queue, 0.3) # self-loop with probability 0.3
P.set(class1, class1, queue, delay, 0.7)
P.set(class1, class1, delay, sink, 1.0)
P.set(class2, class2, delay, queue, 1.0) # note: closed class jobs start at delay
P.set(class2, class2, queue, delay, 1.0)
model.link(P)
```

When used as arguments to a cell array or matrix, class, and node objects will be replaced by a corresponding numerical index. Normally, the indexing of classes and nodes matches the order in which they are instantiated in the model and one can therefore specify the routing matrices using this property. In this case, we would have

```
P = model.initRoutingMatrix()
pmatrx = np.empty(K, dtype=object)
pmatrx[0] = [[0,1,0,0], [0,0.3,0.7,0], [0,0,0,1], [0,0,0,0]]
pmatrx[1] = [[0,0,0,0], [0,0,1,0], [0,1,0,0], [0,0,0,0]]
P.setRoutingMatrix(jobclass, node, pmatrx)
```

Where needed, the `getClassIndex` and `getNodeIndex` functions return the numerical index associated with a node name, for example `model.getNodeIndex('Delay')`. Class and node names in a network *must be unique*. The list of names already assigned to nodes in the network can be obtained with the `getClassNames`, `getStationNames`, and `getNodeNames` functions of the `Network` class.

It is also important to note that the routing matrix in the last example is specified between *nodes*, instead of between just stations or stateful nodes, which means that for example elements such as the `Sink` need to be explicitly considered in the routing matrix. The only exception is that `ClassSwitch` elements do not need to be explicitly instantiated in the routing matrix, provided that one uses the `link` function to instantiate the topology. Note that the routing matrix assigned to a model can be printed on the screen in a human-readable format using the `printRoutingMatrix` function, e.g.,

```
model.printRoutingMatrix()
```

prints

```
Delay [Class1] => Queue1 [Class1] : Pr=1.0
Delay [Class2] => Queue1 [Class2] : Pr=0.001
Queue1 [Class1] => Queue1 [Class1] : Pr=0.3
Queue1 [Class1] => Source [Class1] : Pr=0.7
Queue1 [Class2] => Source [Class2] : Pr=1.0
Source [Class1] => Sink [Class1] : Pr=1.0
Source [Class2] => Queue1 [Class2] : Pr=1.0
Sink [Class2] => Source [Class2] : Pr=1.0
```

Other routing strategies

The above routing specification style is only for models with probabilistic routing strategies between every pair of nodes. A different style should be used for scheduling policies that do not require to explicit routing probabilities, as in the case of state-dependent routing. Currently supported strategies include:

- Round robin (`RoutingStrategy.RROBIN`). This is a deterministic strategy that sends jobs to outgoing links in a cyclic order.

- Random routing (`RoutingStrategy.RAND`). This is equivalent to a standard probabilistic strategy that for each class assigns identical values to the routing probabilities of all outgoing links. When a target is invalid its probability is kept to zero, e.g., random routing will not send a job in a closed class to a sink.
- Join-the-Shortest-Queue (`RoutingStrategy.JSQ`). This is a non-probabilistic strategy that sends jobs to the destination with the smallest total number of jobs in it (either queueing or receiving service). If multiple stations have the same total number of jobs, then the destination is chosen at random with equal probability.

For the above policies, the function `addLink` should be first used to specify pairs of connected nodes

```
model.addLink(queue, queue) #self-loop
model.addLink(queue, delay)
```

Then an appropriate routing strategy should be selected at every node, e.g.,

```
queue.setRouting(class1, RoutingStrategy.RROBIN)
```

assigns round robin among all outgoing links from the `queue` node.

A model could also include both classes with probabilistic routing strategies and classes that use round-robin or other non-probabilistic strategies. To instantiate routing probabilities in such situations one should then use, e.g.,

```
queue.setRouting(class1, RoutingStrategy.PROB)
queue.setProbRouting(class1, queue, 0.7)
queue.setProbRouting(class1, delay, 0.3)
```

where `setProbRouting` assigns the routing probabilities to the two links.

Routing probabilities for Source and Sink nodes

In the presence of open classes, and in mixed models with both open and closed classes, one needs only to specify the routing probabilities *out* of the source. The probabilities out of the sink can all be set to zero for all classes and destinations (including self-loops). The solver will take care of adjusting these inputs to create a valid routing table.

Simplified definition of tandem and cyclic topologies

Tandem networks are open queueing networks with a serial topology. LINE provides functions that ease the definition of tandem networks of stations with exponential service times. For example, the getting started Example 1 on the M/M/1 queue illustrates a simplified way to specify a serial routing topology, i.e.,

```
model.link(Network.serialRouting(source, queue, sink))
```

In a similar fashion, we can also rapidly instantiate a tandem network consisting of stations with PS and INF scheduling as follows

```

TODO: under development
lambda = [10,20] # lambda(r) - arrival rate of class r
D = [[11,12], [21,22]] # D(i,r) - class-r demand at station i (PS)
Z = [[91,92], [93,94]] # Z(i,r) - class-r demand at station i (INF)
modelPsInf = Network.tandemPsInf(lambda,D,Z)

```

The above snippet instantiates an open network with two queueing stations (PS), two delay stations (INF), and exponential distributions with the given inter-arrival rates and mean service times. The `Network.tandemPs`, `Network.tandemFcfs`, and `Network.tandemFcfsInf` functions provide static constructors for networks with other combinations of scheduling policies, namely only PS, only FCFS, or FCFS and INF.

A tandem network with closed classes is instead called a cyclic network. Similar to tandem networks, LINE offers a set of static constructors:

- `Network.cyclicPs`: cyclic network of PS queues
- `Network.cyclicPsInf`: cyclic network of PS queues and delay stations
- `Network.cyclicFcfs`: cyclic network of FCFS queues
- `Network.cyclicFcfsInf`: cyclic network of FCFS queues and delay stations

These functions only require replacing the arrival rate vector `A` by a vector `N` specifying the job populations for each of the closed classes, e.g.,

```

TODO: under development

```

3.1.5 Class switching

Depending on the specified probabilities, a job will be able to switch its class only among a subset of the available classes. Each subset is called a *chain*. Chains are computed in LINE as the weakly connected components of the routing probability matrix of the network when this is seen as an undirected graph. The function `model.getChains()` produces the list of chains for the model, inclusive of a list of their composing classes.

The definition of class switching in a model is integrated in the specification of the routing between stations as described in the next subsection.

Probabilistic class switching

In models with class switching and probabilistic routing at all nodes, a routing matrix is required for each possible pair of source and target classes. For instance, suppose that in the previous example the job in the

closed class `class2` switches into a new closed class (`class3`) while visiting the queue node. We can specify this routing strategy as follows:

TODO: under development

Importantly, LINE assumes that a job switches class an instant *after* leaving a station, thus the performance metrics of a class at the node refer to the class that jobs had upon arrival to that node.

Class switching with non-probabilistic routing strategies

In the presence of non-probabilistic routing strategies, such as round-robin or join-the-shortest-queue, one may need to manually specify the details of the class switching mechanism. This can be done through addition to the network topology of `ClassSwitch` nodes.

The constructor of the `ClassSwitch` node requires to specify a probability matrix C such that the element in row r and column s is the probability that a job of class r arriving into the node switches to class s during the visit. For example, in a 2-class model, the following node will switch all visiting jobs into class 2

```
# Block 1: nodes
...
csnode = ClassSwitch(model, 'ClassSwitch 1')
# Block 2: classes
jobclass = np.empty(2, dtype=object)
jobclass[0] = OpenClass(model, 'Class1', 0)
jobclass[1] = OpenClass(model, 'Class2', 0)
...
# Block 3: topology
C = csnode.initClassSwitchMatrix()
C[0][1] = 1.0
C[1][1] = 1.0
csnode.setClassSwitchingMatrix(C)
```

Note that for a network with M stations, up to M^2 `ClassSwitch` elements may be required to implement class-switching across all possible links, including self-loops.

Cache-based class-switching

An advanced feature of LINE available for example within the `Cache` node, is that the class-switching decision can dynamically depend on the state of the node (e.g., cache hit/cache miss). However, in order to statically determine chains, LINE requires that every class-switching node declares the pair of classes that can potentially communicate with each other via a switch. This is called the *class-switching mask* and it is automatically computed. The boolean matrix returned by the `model.getClassSwitchingMask` function provides this mask, which has an entry in row r and column s set to true only if jobs in class r can switch into class s at some node in the network.

Upon cache hit or cache miss, a job in transit is switched to a user-specified class, as specified by the `setHitClass` and `setMissClass`, so that it can be routed to a different destination based on whether it found the item in the cache or not. The `setRead` function allows the user to specify a discrete distribution (e.g., `Zipf`, `DiscreteSampler`) for the frequency at which an item is requested. For example,

```
refModel = Zipf(0.5,nitems)
cacheNode.setRead(initClass, refModel)
cacheNode.setHitClass(initClass, hitClass)
cacheNode.setMissClass(initClass, missClass)
```

Here `initClass`, `hitClass`, and `missClass` can be either open or closed instantiated as usual with the `OpenClass` or `ClosedClass` constructors.

3.1.6 Service and inter-arrival time processes

A number of statistical distributions are available to specify job service times at the stations and inter-arrival times from the `Source` station. The class `PhaseType` offers distributions that are analytically tractable, which are defined using absorbing Markov chains consisting of one or more states (*phases*) and called phase-type distributions. They include as special cases the following distributions:

- Exponential distribution: `Exp(λ)`, where λ is the rate of the exponential
- n -phase Erlang distribution: `Erlang(α, n)`, where α is the rate of each of the n exponential phases
- 2-phase hyper-exponential distribution: `HyperExp(p, λ_1, λ_2)`, that returns an exponential with rate λ_1 with probability p , and an exponential with rate λ_2 otherwise.
- n -phase hyper-exponential distribution: `HyperExp(p, λ)`, that builds a n -phase hyper-exponential from a rate vector $\lambda = [\lambda_1, \dots, \lambda_n]$ and phase selection probabilities $p = [p_1, \dots, p_n]$.
- 2-phase Coxian distribution: `Coxian(μ_1, μ_2, ϕ_1)`, which assigns phases μ_1 and μ_2 to the two rates, and completion probability from phase 1 equal to ϕ_1 (the probability from phase 2 is $\phi_2 = 1.0$).
- n -phase Coxian distribution: `Coxian(μ, ϕ)`, which builds an arbitrary Coxian distribution from a vector $\mu = [\mu_1, \dots, \mu_n]$ of n rates and a completion probability vector $\phi = [\phi_1, \dots, \phi_n]$ with $\phi_n = 1.0$.
- n -phase acyclic phase-type distribution: `APH(α, T)`, which defines an acyclic phase-type distribution with initial probability vector $\alpha = [\alpha_1, \dots, \alpha_n]$ and transient generator T .

For example, given mean $\mu = 0.2$ and squared coefficient of variation $SCV=10$, where $SCV=\text{variance}/\mu^2$, we can assign to a node a 2-phase Coxian service time distribution with these moments as

```
queue.setService(class2, Cox2.fitMeanAndSCV(0.2,10.0))
```

where `Cox2` is a static class to fit 2-phase Coxian distributions. Inter-arrival time distributions can be instantiated in a similar way, using `setArrival` instead of `setService` on the `Source` node. For example, if the `Source` is node 3 we may assign the inter-arrival times of class 2 to be exponential with mean 0.1 as follows

```
source.setArrival(class2, Exp.fitMean(0.1))
```

It is also possible to plot the structure of a phase-type distribution using `PhaseType.plot` static method.

Non-Markovian distributions are also available, but typically they can restrict the available solvers to the JMT simulator. They include the following distributions:

- Deterministic distribution: `Det(μ)` assigns probability 1.0 to the value μ .
- Uniform distribution: `Uniform(a, b)` assigns uniform probability $1/(b - a)$ to the interval $[a, b]$.
- Gamma distribution: `Gamma(α, k)` assigns a gamma density with shape α and scale k .
- Pareto distribution: `Pareto(α, k)` assigns a Pareto density with shape α and scale k .

Lastly, we discuss two special distributions. The `Disabled` distribution can be used to explicitly forbid a class to receive service at a station. This may be useful to declare in models with sparse routing matrices to debug the model specification. Performance metrics for disabled classes will be set to NaN.

Conversely, the `Immediate` class can be used to specify instantaneous service (zero service time). Normally, LINE solvers will replace zero service times with small positive values ($\epsilon = \text{GlobalConstants.FineTol}$).

Fitting a distribution

The `fitMeanAndSCV` function is available for all distributions that inherit from the `PhaseType` class. This function provides exact or approximate matching of the first two moments, depending on the theoretical constraints imposed by the distribution. For example, an Erlang distribution with $\text{SCV}=0.75$ does not exist, because in a n -phase Erlang it must be $\text{SCV}=1/n$. In a case like this, `Erlang.fitMeanAndSCV(1, 0.75)` will return the closest approximation, e.g., a 2-phase Erlang ($\text{SCV}=0.5$) with unit mean. The Erlang distribution also offers a function `fitMeanAndOrder(μ, n)`, which instantiates a n -phase Erlang with given mean μ .

In distributions that are uniquely determined by more than two moments, `fitMeanAndSCV` chooses a particular assignment of the residual degrees of freedom other than mean and SCV. For example, `HyperExp` depends on three parameters, therefore it is insufficient to specify mean and SCV to identify the distribution. Thus, `HyperExp.fitMeanAndSCV` automatically chooses to return a probability of selecting phase 1 equal to 0.99. Compared to other choices, this particular assignment corresponds to a higher probability mass in the tail of the distribution. `HyperExp.fitMeanAndSCVBalanced` instead assigns p in a two-phase hyper-exponential distribution so that $p/\mu_1 = (1 - p)/\mu_2$.

Inspecting and sampling a distribution

To verify that the fitted distribution has the expected mean and SCV it is possible to use the `getMean` and `getSCV` functions, e.g.,

```
TODO: under development
```

Moreover, the `sample` function can be used to generate values from the obtained distribution, e.g. we can generate 3 samples as

```
TODO: under development
```

The `evalCDF` and `evalCDFInterval` functions return the cumulative distribution function at the specified point or within a range, e.g.,

```
TODO: under development
```

For more advanced uses, the distributions of the `PhaseType` class also offer the possibility to obtain the standard (D_0, D_1) representation used in the theory of Markovian arrival processes by means of the `getRepresentation` function [5].

Load-dependent service

A queueing station i is called *load-dependent* whenever its service rate is a function of the number n_i of resident jobs at the station, summed across the ones in service and the ones in the waiting buffer. For example, a multi-server station with c identical servers, each with processing rate μ , may be shown to behave similarly to a single-server load-dependent station where the service rate is $\mu(n_i) = \mu\alpha(n_i) = \mu \min(n_i, c)$.

LINE presently supports *limited load-dependence* [12], meaning that it is possible to specify the form of the load-dependent service up to a finite range of n_i . As such, the support is currently limited to closed models, which are guaranteed to have a finite population at all times.

To specify a load-dependence service for a queueing station over the range $n_i \in [1, N]$ it is sufficient to call the `setLoadDependence` method, passing a vector of size N in its input with the scaling factor values for each n_i . For example, to instantiate a c -server node we write

```
queue.setLoadDependence(np.minimum(np.arange(0,N,1),c)) % multi-server with c ...  
servers
```

where the i -th element of the vector argument of `setLoadDependence` is the scaling factor $\alpha(n_i)$. It is assumed by default that $\alpha(0) = 1$.

Class-dependent service

A generalization of load-dependent service is *class-dependent* service, where the service rate is now a function of the vector $n_i = [n_{i,1}, \dots, n_{i,R}]$, where $n_{i,r}$ is the current number of class- r jobs at station i .

LINE supports class-dependence in the MVA solver, provided that this is specified as a function handle. The solver implicitly assumes that the function is smooth and defined also for fractional values of $n_{i,r}$. For example, in a two-class model we may write

```
TODO: under development
```

applies a multiserver-type only to class-2 jobs, but not to the others.

```
TODO: under development
```

Temporal dependent processes

It is sometimes useful to specify the statistical properties of a *time series* of service or inter-arrival times, as in the case of systems with short- and long-range dependent workloads. When the model is stochastic, we refer to these as situations where one specifies a *process*, as opposed to only specifying the *distribution* of the service or inter-arrival times. In LINE processes inherit from the `PointProcess` class, and include the 2-state Markov-modulated Poisson process (MMPP2) and empirical traces read from files (`Replayer`).

For the latter, LINE assumes that empirical traces are supplied as text files (ASCII), formatted as a column of numbers. Once specified, the `Replayer` object can be used as any other distribution. This means that it is possible to run a simulation of the model with the specified trace. However, analytical solvers will require tractable distributions from the `PhaseType` class.

3.2 Internals

In this section, we discuss the internal representation of the `Network` objects used within the LINE solvers. By applying changes directly to this internal representation it is possible to considerably speed up the sequential evaluation of models.

3.2.1 Representation of the model structure

For efficiency reasons, once a user requests to solve a `Network`, LINE calls internally generate a static representation of the network structure using the `refreshStruct` function. This function returns a representation object that is then passed on to the chosen solver to parameterize the analysis.

The representation used within LINE is the `NetworkStruct` class, which describes an extended multiclass queueing network with class-switching and acyclic phase-type (APH) service times. APH generalizes known distributions such as Coxian, Erlang, Hyper-Exponential, and Exponential. The representation can be obtained as follows

```
sn = model.getStruct()
```

The next tables present the properties of the `NetworkStruct` class. Table ?? gives the invariant properties of the class, which specify the initial parameterization of the nodes, which we refer to as static parameters. Table ?? further details parameters that require algorithmic calculations to derive from the static parameters.

TODO: under development

TODO: under development

For advanced nodes, such as `Cache` and `Transition`, additional parameters are specified under the `nodeparam` cell array for the corresponding node. Table 3.2.1 illustrates the properties specified within the `nodeparam{i}` cell array for `Transition` node i . Table 3.2.1 similarly illustrates the properties in `nodeparam{i}` array for a `Cache` node i .

Table 3.2: `nodeparam{i}` fields for a `Transition` node i

Field	Type	Description
<code>enabling{m}(k,r)</code>	integer	Enabling condition for mode m at transition node i with respect to class r jobs at linked node k
<code>firing{m}(k,r)</code>	integer	Firing outputs of class r jobs for mode m at transition node i towards linked node k
<code>firingprio(m)</code>	integer	Firing priority for mode m at transition node i
<code>firingid(m)</code>	integer	Firing type at node i for mode m (e.g., <code>TimingStrategy.IMMEDIATE</code>)
<code>firingphases(m)</code>	integer	Number of phases for firing process of mode m at node i
<code>firingproc{m}</code>	cell	Matrix representation of the mode m firing process at transition node i
<code>firingprocid(m)</code>	integer	Firing process type id for mode m (e.g., <code>ProcessType.ID.HYPEREXP</code>)
<code>fireweight(m)</code>	integer	Firing weight for mode m at transition node i
<code>inhibiting{m}(k,r)</code>	integer	Inhibiting condition for mode m at transition node i with respect to class r jobs at linked node k
<code>modenames{m}</code>	string	Name of mode m at transition node i
<code>nmodes</code>	integer	Number of modes for transition node i
<code>nmodeservers{m}</code>	integer	Number of servers for mode m

Table 3.3: `nodeparam{i}` fields for a `Cache` node i

Field	Type	Description
<code>accost{r,i}(l,h)</code>	double	Access cost for class r to move item i from list l to list h .
<code>hitclass(r)</code>	integer	Class switching for class r under a cache hit
<code>itemcap(l)</code>	integer	Item capacity in list l
<code>missclass(r)</code>	integer	Class switching for class r under a cache miss
<code>nitems</code>	integer	Number of items
<code>pread{r}(i)</code>	double	Probability for class r of reading item i
<code>rpolicy</code>	integer	Replacement policy id (e.g., <code>ReplacementPolicy.ID.LRU</code>)

As shown in the tables, internally to `LINE` there is an explicit differentiation between properties of nodes, stations, and stateful nodes. This distinction has impact in particular over routing and class-switching mechanisms, and also allows solvers to better differentiate between different kinds of nodes.

In some cases, one may want to access some properties of nodes that are contained in `NetworkStruct` fields that are however referenced by station or stateful node index. To help this and similar situations, the `NetworkStruct` class also provides static methods to quickly convert the indexing of nodes, stations, and stateful nodes, which is used in referencing its data structures:

- `nodeToStateful`
- `nodeToStation`
- `stationToNode`
- `stationToStateful`
- `statefulToNode`

As an example, we can determine the portion of the `nodevisits` field that refers to stateful nodes in chain $c = 1$ as follows

```
TODO: under development
```

3.3 Debugging and visualization

JSIMgraph is the graphical simulation environment of the JMT suite. LINE can export models to this environment for visualization purposes using the command

```
TODO: under development
```

An example is shown in Figure [Figure 3.1](#) below. Using a related function, `jsimwView`, it is also possible to export the model to the JSIMwiz environment, which offers a wizard-based interface.

Another way to debug a LINE model is to transform it into a MATLAB graph object, e.g.

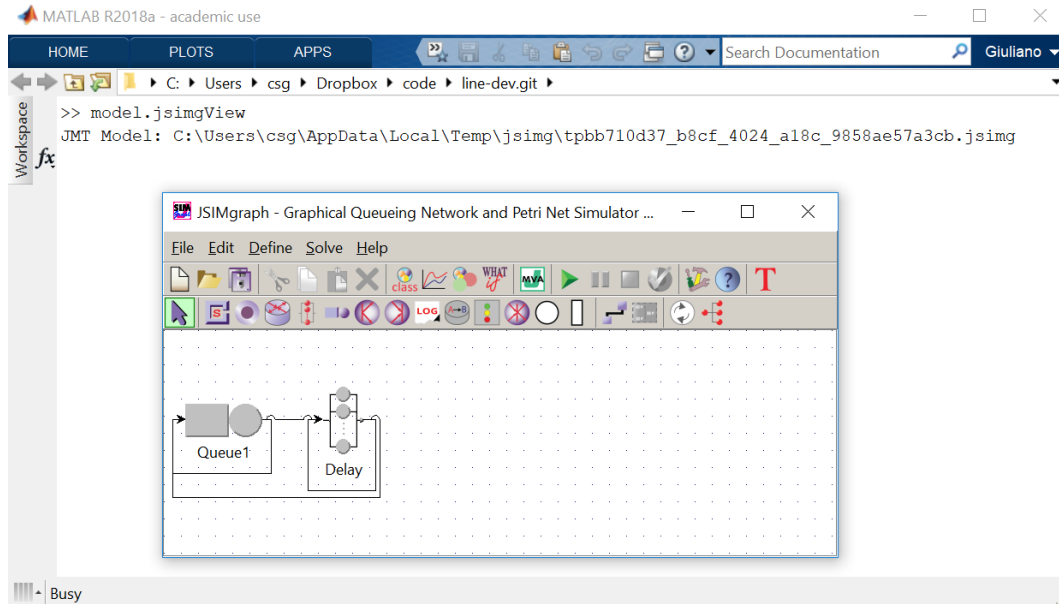
```
TODO: under development
```

plots a graph of the network topology in term of stations only. In a similar manner, the following variant of the same command shows the model in terms of nodes, which corresponds to the internal representation within LINE.

```
TODO: under development
```

The latter example is invoked by default if we type

```
TODO: under development
```

Figure 3.1: `jsimgView` function

which also adds automatic node coloring to highlight the class-switch nodes.

Figure 3.2 shows the difference between the two commands for an open queueing network with two classes and class-switching. Weights on the edges correspond to routing probabilities. In the station topology on the left, note that since the Sink node is not a station, departures to the Sink are drawn as returns to the Source. The node topology on the right, illustrates all nodes, including ClassSwitch nodes that are automatically added by LINE to apply the class-switching routing strategy. Double arcs between nodes indicate that both classes are routed to the destination.

Furthermore, the graph properties concisely summarize the key features of the network

TODO: under development

Here, `Edge.Weight` is the routing probability between the nodes, whereas `Edge.Rate` is the service rate of the node appearing in the first column under `EndNodes`.

3.4 Model import and export

LINE offers a number of scripts to import external models into `Network` object instances that can be analyzed through its solvers. The available scripts are as follows:

- `JMT2LINE` imports a JMT simulation model (`.jsimg` or `.jsimw` file) instance.




Figure 3.2: `getGraph` function: station topology (left) and node topology (right) for a 2-class tandem queueing network with class-switching.

- `PMIF2LINE` imports an XML file containing a PMIF 1.0 model.

Both scripts require in input the filename and desired model name, and return a single output, e.g.,

```
TODO: under development
```

where `sn` is an instance of the `Network` class.

```
TODO: under development
```

```
TODO: under development
```

3.4.1 Creating a LINE model using JMT

Using the features presented in the previous section, one can create a model in JMT and automatically derive a corresponding LINE script from it. For instance, the following command performs the import and translation into a script, e.g.,

```
TODO: under development
```

transforms and save the given `JSIMgraph` model into a corresponding `LINE` model.

`LINE` also provides two static functions to inspect `jsimg` and `jsimw` files before conversion, called `SolverJMT.jsimgOpen` and `SolverJMT.jsimwOpen` require as an input parameter only the JMT file name, e.g., `myModel.jsimg`.

It is also possible to automate the editing and import of JMT models from MATLAB using the `jsimgEdit` command. This will open an empty JMT model and upon saving this will be automatically reimported into MATLAB.

3.4.2 Supported JMT features

Table 3.4 lists the JSIMgraph/JSIMwiz model features supported by the JMT2LINE transformation. We indicate as “Fully” supported a feature that is supported in the import and such that the resulting model can be solved in LINE using at least SolverJMT. A feature with “Partial” support implies that some core aspects of this feature available in JSIM are not available in LINE.

A few notes are needed to clarify the entries with partial support:

- Fork and Join are supported with their default policies. Advanced policies, such as partial joins or setting a distribution for the forked tasks on each output link, are not supported yet.
- a single Sink and a single Source can be instantiated in a LINE model, whereas there is no such constraint in JMT.

Table 3.4: Supported JSIM features for automated model import and analysis

JMT Feature	Support	Notes
Distributions	Full	Phase-Type, Burst (MAP), Burst (MMPP2), Deterministic, Disabled, Exponential, Erlang, Gamma, Hyperexponential, Coxian, Logistic, Pareto, Uniform, Zero Service Time, Replayer, Weibull
Classes	Full	Open class, Closed class, Class priorities
Metrics	Full	Number of customers, Residence Time, Throughput, Response Time, Throughput per sink, Utilization, Arrival Rate
Nodes	Full	Finite Capacity Region, ClassSwitch, Place, Delay, Logger, Queue, Router, Transition
Routing	Full	Random, Probabilities, Round Robin, Join the Shortest Queue
Mechanisms	Full	N/A
Scheduling	Full	FCFS, HOL, LCFS, LCFS-PR, SIRO (Random), SJF, SEPT, LJF, LEPT, PS, DPS, GPS
Nodes	Partial	Fork, Join, Source, Sink
Distributions	No	Burst (General), Normal
Nodes	No	Scaler, Semaphore
Routing	No	Shortest Response Time, Least Utilization, Fastest Service, Load Dependent, Class Switch Routing
Metrics	No	Drop rate, Response time per sink, Power
Scheduling	No	LPS, EDD, EDF, TBS, SRPT
Mechanisms	No	Load Dependence, Retrial, Impatience, Soft deadlines, Parallelism, , Heterogeneous servers, Server Compatibilities, Setup times, Polling, Switchover times

Chapter 4

Analysis methods

4.1 Performance metrics

As discussed earlier, LINE supports a set of steady-state and transient performance metrics. Table 4.1 summarizes the definition of the associated random variables. For each metric, one or more analysis types may be available, which are extensively discussed in the next sections.

Table 4.1: Performance metrics

Metric	Acronym	Description
Queue-length	QLen	Number of jobs of class r (or chain- c) residing at a node i
Utilization	Util	Utilization of class- r (or chain- c) jobs at node i , scaled in $[0,1]$ for multi-server nodes, equal to QLen at infinite server nodes
Response time	RespT	Time that a class- r (or chain- c) jobs spends for a single visit at node i
Residence time	ResidT	Cumulative time that a class- r (or chain- c) jobs spends across all visits at node i
Arrival rate	ArvR	Arrival rate of class- r (or chain- c) jobs at node i
Throughput	Tput	Throughput of class- r (or chain- c) jobs at node i
System Response time	SysRespT	For a open chain c , this is the time from leaving the source to arriving at the sink for <i>any</i> class in the chain. For a closed chain c , this is the interval of time between two successive visits to the reference station in any two <i>completing classes</i> within the chain.
System Throughput	SysTput	For a open chain c , this is the departure rate towards the sink for <i>any</i> class in the chain. For a closed chain c , this is the rate of arrival of <i>completing classes</i> in the chain at the reference station.

4.2 Steady-state analysis

4.2.1 Station average performance

LINE decouples network specification from its solution, allowing to evaluate the same model with multiple solvers. Model analysis is carried out in LINE according to the following general steps:

Step 1: Definition of the model. This proceeds as explained in the previous chapters.

Step 2: Instantiation of the solver(s). A solver is an instance of the `Solver` class. LINE offers multiple solvers, which can be configured through a set of common and individual solver options. For example,

```
solver = SolverJMT(model)
```

returns a handle to a simulation-based solver based on JMT, configured with default options.

Step 3: Solution. Finally, this step solves the network and retrieves the concrete values for the performance indexes of interest. This may be done as follows, e.g.,

```
# QN(i,r): mean queue-length of class r at station i
QN = solver.getAvgQLen()
# UN(i,r): utilization of class r at station i
UN = solver.getAvgUtil()
# RN(i,r): mean response time of class r at station i (per visit)
RN = solver.getAvgRespT()
# TN(i,r): mean throughput of class r at station i
TN = solver.getAvgTput()
# AN(i,r): mean arrival rate of class r at station i
AN = solver.getAvgArvR()
# WN(i,r): mean residence time of class r at station i (summed on visits)
WN = solver.getAvgResidT()
```

Alternatively, all the above metrics may be obtained in a single method call as

```
TODO: under development
```

In the methods above, LINE assigns station and class indexes (e.g., i , r) in order of creation in order of creation of the corresponding station and class objects. However, large models may be easier to debug by checking results using class and station names, as opposed to indexes. This can be done either by requesting LINE to build a table with the result

```
TODO: under development
```

which however tends to be a rather slow data structure to use in case of repeated invocations of the solver, or by indexing the matrices returned by `getAvg` using the model objects. That is, if the first instantiated node is `queue` with name `MyQueue` and the second instantiated class is `cclass` with name `MyClass`, then the following commands are equivalent

```
TODO: under development
```

Similar methods are defined to obtain aggregate performance metrics at chain level at each station, namely `getAvgQLenChain` for queue-lengths, `getAvgUtilChain` for utilizations, `getAvgRespTChain` for response times, `getAvgTputChain` for throughputs, and the `getAvgChain` method to obtain all the previous metrics.

4.2.2 Station response time distribution

`SolverFluid` supports the computation of response time distributions for individual classes through the `getCdfRespT` function. The function returns the response time distribution for every station and class. For example, the following code plots the cumulative distribution function at steady-state for class 1 jobs when they visit station 2:

```
TODO: under development
```

4.2.3 System average performance

LINE also allows users to analyze models for end-to-end performance indexes such a system throughput or system response time. However, in models with class switching the notion of system-wide metrics can be ambiguous. For example, consider a job that enters the network in one class and departs the network in another class. In this situation one may attribute system response time to either the arriving class or the departing one, or attempt to partition it proportionally to the time spent by the job within each class. In general, the right semantics depends on the aim of the study.

LINE tackles this issue by supporting only the computation of system performance indexes *by chain*, instead than by class. In this way, since a job switching from a class to another remains by definition in the same chain, there is no ambiguity in attributing the system metrics to the chain. The solver functions `getAvgSys` and `getAvgSysTable` return system response time and system throughput per chain as observed: (i) upon arrival to the sink, for open classes; (ii) upon arrival to the reference station, for closed classes.

In some cases, it is possible that a chain visits multiple times the reference station before the job completes. This also affects the definition of the system averages, since one may want to avoid counting each visit as a completion of the visit to the system. In such cases, LINE allows users to specify which classes of the chain can complete at the reference station. For example, in the code below we require that a job visits reference station 1 twice, in classes 1 and 2, but completes at the reference station only when arriving in class 2. Therefore, the system response time will be counted between successive passages in class 2.

```
TODO: under development
```

Note that the `completes` property of a class always refers to the reference station for the chain.

4.3 Specifying states

In some analyses it is important to specify the state of the network, for example to assign the initial position of the jobs in a transient analysis. We thus discuss the support in LINE for state modeling.

4.3.1 Station states

We begin by explaining how to specify a state s_0 for a station. For example, it is not supported for shortest job first (`SchedStrategy.SJF`) scheduling, in which state must include the service time samples for the jobs and it is therefore a continuous quantity.

Suppose that the network has R classes and that service distributions are phase-type, i.e., that they inherit from `PhaseType`. Let K_r be the number of phases for the service distribution in class r at a given station. Then, we define three types of state variables:

- c_j : class of the job waiting in position $j \leq b$ of the buffer, out of the b currently occupied positions. If $b = 0$, then the state vector is indicated with a single empty element $c_1 = 0$.
- k_j : service phase of the job waiting in position $j \leq b$ of the buffer, out of the b currently occupied positions.
- n_r : total number of jobs of class r in the station
- b_r : total number of jobs of class r in the station's buffer
- s_{rk} : total number of jobs of class r running in phase k in the server

Here, by phase we mean the number of states of a distribution of class `PhaseType`. If the distribution is not Markovian, then there is a single phase. With these definitions, the table below illustrates how to specify in `LINE` a valid state for a station depending on its scheduling strategy. There, S is the number of servers of the queueing station. All state variables are non-negative integers. The `SchedStrategy.EXT` policy is used for the `Source` node, which may be seen as a special station with an infinite pool of jobs sitting in the buffer and a dedicated server for each class $r = 1, \dots, R$.

Table 4.2: State descriptors for Markovian scheduling policies

Sched. strategy	Station state vector	State condition
EXT	$[Inf, s_{11}, \dots, s_{1K_1}, \dots, s_{R1}, \dots, s_{RK_R}]$	$\sum_k s_{rk} = 1, \forall r$
FCFS, HOL, LCFS	$[c_b, \dots, c_1, s_{11}, \dots, s_{1K_1}, \dots, s_{R1}, \dots, s_{RK_R}]$	$\sum_r \sum_k s_{rk} = S$
LCFSPR	$[c_b, k_b, \dots, c_1, k_1, s_{11}, \dots, s_{1K_1}, \dots, s_{R1}, \dots, s_{RK_R}]$	$\sum_r \sum_k s_{rk} = S$
SEPT, SIRO	$[b_1, \dots, b_R, s_{11}, \dots, s_{1K_1}, \dots, s_{R1}, \dots, s_{RK_R}]$	$\sum_r \sum_k s_{rk} = S$
PS, DPS, GPS, INF	$[s_{11}, \dots, s_{1K_1}, \dots, s_{R1}, \dots, s_{RK_R}]$	None

States can be manually specified or enumerated automatically. `LINE` library functions for handling and generating states are as follows:

- `State.fromMarginal`: enumerates all states that have the same marginal state $[n_1, n_2, \dots, n_R]$.

- `State.fromMarginalAndRunning`: restricts the output of `State.fromMarginal` to states with given number of running jobs, irrespectively of the service phase in which they currently run.
- `State.fromMarginalAndStarted`: restricts the output of `State.fromMarginal` to states with given number of running jobs, all assumed to be in service phase $k = 1$.
- `State.fromMarginalBounds`: similar to `State.fromMarginal`, but produces valid states between given minimum and maximum value of the number of resident jobs.
- `State.toMarginal`: extracts marginal statistics from a state, such as the total number of jobs in a given class that are running at the station in a certain phase.

Note that if a function call returns an empty state (`[]`), this should be interpreted as an indication that no valid state exists that meets the required criteria. Often, this is because the state supplied in input is invalid.

Example

We consider the example network in `TODO`. We look at the state of station 3, which is a multi-server FCFS station. There are 4 classes all having exponential service times except class 2 that has Erlang-2 service times. We are interested to states with 2 running jobs in class 1 and 1 in class 2, and with 2 jobs, respectively of classes 3 and 4, waiting in the buffer. We can automatically generate this state space, which we store in the `space` variable, as:

```
TODO: under development
```

Here, each row of `space` corresponds to a valid state. The argument `TODO` gives the number of jobs in the node for the 4 classes, while `TODO` gives the number of running jobs in each class. This station has four valid states, differing on whether the class-2 job runs in the first or in the second phase of the Erlang-2 and on the relative position of the jobs of class 3 and 4 in the waiting buffer.

To obtain states where the jobs have just started running, we can instead use

```
TODO: under development
```

We see that the above state space restricted the one obtained with `State.fromMarginalAndRunning` to states where the job in class 1 is always in the first phase.

If we instead remove the specification of the running jobs, we can use `State.fromMarginal` to generate all possible combinations of states depending on the class and phase of the running jobs. In the example, this returns a space of 20 possible states.

```
TODO: under development
```


Assigning a prior to an initial state

It is possible to assign the initial state to a station using the `setState` function on that station's object. LINE offers the possibility to specify a prior probability on the initial states, so that if multiple states have a non-zero prior, then the solver will need to solve an independent model using each one of those initial states, and then carry out a weighting of the results according to the prior probabilities. The default is to assign a probability 1.0 to the *first* specified state. The functions `setStatePrior` and `getStatePrior` can be used to check and change the prior probabilities for the initial states specified for a station or stateful node.

4.3.2 Network states

A collection of states that are valid for each station is not necessarily valid for the network as a whole. For example, if the sum of jobs of a closed class exceeds the population of the class, then the network state would be invalid. To identify these situations, LINE requires to specify the initial state of a network using functions supplied by the `Network` class. These functions are `initFromMarginal`, `initFromMarginalAndRunning`, and `initFromMarginalAndStarted`. They require a matrix `n` with elements (i, r) specifying the total number of resident class- r jobs at node i and the latter two require a matrix `s` with elements (i, r) with the number of running (or started) class- r jobs at node i . The user can also manually verify if the supplied network state is going to be valid using `State.IsValid`.

It is also possible to request LINE to automatically identify a valid initial state, which is done using the `initDefault` function available in the `Network` class. This is going to select a state where:

- no jobs in open classes are present in the network;
- jobs in closed classes all start at their reference stations;
- the servers at each reference station are occupied by jobs of in class order, i.e., jobs in the firstly created class are assigned to the server, then spare server are allocated to jobs in the second class, and so forth;
- service or arrival processes are initialized in phase 1 for each job;
- if the scheduling strategy requires it, jobs are ordered in the buffer by class, with the firstly created class at the head and the lastly created class at the tail of the buffer.

The `initFromAvgQLen` method is a wrapper for `initFromMarginal` to initialize the system as close as possible to the average steady-state distribution of the network. Since averages are typically not integer-valued, this function rounds the average values to the nearest integer and adjusts the result to ensure feasibility of the initialization.

4.3.3 Initialization of transient classes

Because of class-switching, it is possible that a class r with a non-empty population at time $t = 0$ becomes empty at some position time $t' > t$ without ever being visited again by any job. LINE allows users to place jobs in transient classes and therefore it will not trigger an error in the presence of this situation. If a user wishes to prohibit the use of a class at a station, it is sufficient to specify that the corresponding service process uses the `Disabled` distribution.

Certain solvers may incur problems in identifying that a class is transient and in setting to zero its steady-state measures. For example, the JMT solver uses an heuristic whereby a class is considered transient if it has fewer samples than jobs initially placed in the corresponding chain the class belongs to. For such classes, JMT will set the values of steady-state performance indexes to zero.

4.3.4 State space generation

As discussed in Example 3, the state space of a model can be obtained by either invoking `model.getStateSpace()` or `solver.getStateSpace()` on an instant of the CTMC solver, where the latter returns the state space cached during the solution of the CTMC.

LINE supports two state space generation method, which is configurable using the option `options.config.state-s` of the CTMC solver. Details may be found in Table 5.2.

4.4 Transient analysis

So far, we have seen how to compute steady-state average performance indexes, which are given by

$$E[n] = \lim_{t \rightarrow +\infty} E[n(t)]$$

where $n(t)$ is an arbitrary performance index, e.g., the queue-length of a given class at time t .

We now consider instead the computation of the quantity $E[n(t)|s_0]$, which is the *transient average* of the performance index, conditional on a given initial system state s_0 . Compared to $n(t)$, this quantity averages the system state at time t across all possible evolutions of the system from state s_0 during the t time units, weighted by their probability. In other words, we observe all possible stochastic evolutions of the system from state s_0 for t time units, recording the final values of $n(t)$ in each trajectory, and finally average the recorded values at time t to obtain $E[n(t)|s_0]$.

4.4.1 Computing transient averages

The computation of transient metrics proceeds similarly to the steady-state case. We first obtain the handles for transient averages:

TODO: under development

After solving the model, we will be able to retrieve *both* steady-state and transient averages as follows

```
TODO: under development
```

The transient average queue-length at node i for class r is stored within `QNT` in row i and column r .

Note that the above code specifies a maximum time t for the output time series. This can be done using the `timespan` solver option. This applies also to average metrics. In the following example, the first model is solved at steady-state, while the second model reports averages at time $t = 1$ after initialization

```
TODO: under development
```

4.4.2 First passage times into stations

When the model is in a transient, the average state seen upon arrival to a station changes over time. That is, in a transient, successive visits by a job may experience different response time distributions. The function `getTranCdfRespT`, implemented by `SolverJMT` offers the possibility to obtain this distribution given the initial state specified for the model. As time passes, this distribution will converge to the steady-state one computed by solvers equipped with the function `getCdfRespT`.

However, in some cases one prefers to replace the notion of response time distribution in transient by the one of *first passage time*, i.e., the distribution of the time to complete the *first visit* to the station under consideration. The function `getTranCdfFirstPassT` provides this distribution, assuming as initial state the one specified for the model, e.g., using `setState` or `initDefault`. This function is available only in `SolverFluid` and has a similar syntax as `getCdfRespT`.

4.5 Sample path analysis

With `LINE` is also possible to obtain a particular sample path from the stochastic process underlying the queueing network. The following functions are available for this purpose:

- `sample`: returns a data structure including the time-varying state of a given stateful node, labelled with information about the events that changed the node state.
- `sampleAggr`: returns a data structure similar to the one provided by `sample`, but where the state is aggregate to count the number of jobs in each class at the node.
- `sampleSys`: similar to the `sample` function, but returns the state of every stateful node in the model.
- `sampleSysAggr`: similar to the `sampleAggr` function, but returns the aggregated state of every stateful node in the model.

It is worth noting that the JMT solver only supports `sampleAggr` since the simulator does not offer a simple way to extra detailed data such as phase change information in the service process. This information is instead available with the SSA solver.

For example, the following command extract a sample path consisting of 10 samples for a $APH(2)/M/1$ queue:

```
TODO: under development
```

In the example, TODO refers to the time since initialization at which the node 2 (here the $APH(2)/M/1$ queueing station) enters the state shown in the second column.

If we repeat the same experiment with the SSA solver and using the `sampleSys` function, we now have the full state space of the model, including both the source and the queueing station:

```
TODO: under development
```

4.6 Sensitivity analysis and numerical optimization

Frequently, performance and reliability analysis requires to change one or more model parameters to see the sensitivity of the results or to optimize some goal function. In order to do this efficiently, we have discussed before the internal representation of the `Network` objects used within the LINE solvers. By applying changes directly to this internal representation it is possible to considerably speed-up the sequential evaluation of models as discussed next.

4.6.1 Fast parameter update

Successive invocations of `getStruct()` will return a cached copy of the `NetworkStruct` representation, unless the user has called `model.refreshStruct()` or `model.reset()` in-between the invocations. The `refreshStruct` function regenerates the internal representation, while `reset` destroys it, together with all other representations and cached results stored in the `Network` object. In the case of `reset`, the internal data structure will be regenerated at the next `refreshStruct()` or `getStruct()` call.

The performance cost of updating the representation can be significant, as some of the structure array field require a dedicated algorithm to compute. For example, finding the chains in the model requires an analysis of the weakly connected components of the network routing matrix. For this reason, the `Network` class provides several functions to selectively refresh only part of the `NetworkStruct` representation, once the modification has been applied to the objects (e.g., stations, classes, ...) used to define the network. These functions are as follows:

- `refreshArrival`: this function should be called after updating the inter-arrival distribution at a `Source`.

- `refreshCapacity`: this function should be called after changing buffer capacities, as it updates the `capacity` and `classcapacity` fields.
- `refreshChains`: this function should be used after changing the routing topology, as it refreshes the `rt`, `chains`, `nchains`, `nchainjobs`, and `visits` fields.
- `refreshPriorities`: this function updates class priorities in the `classprio` field.
- `refreshScheduling`: updates the `sched`, `schedid`, and `schedparam` fields.
- `refreshProcesses`: updates the `mu`, `phi`, `phases`, `rates` and `scv` fields.

For example, suppose we wish to update the service time distribution for class-1 at node 1 to be exponential with unit rate. This can be done efficiently as follows:

```
TODO: under development
```

4.6.2 Refreshing a network topology with non-probabilistic routing

The `resetNetwork` function should be used before changing a network topology with non-probabilistic routing. It will destroy by default all class switching nodes. This can be avoided if the function is called as, e.g., `model.resetNetwork(False)`. The default behavior is though shown in the next example

```
TODO: under development
```

As shown, `resetNetwork` updates the station indexes and the revised list of nodes that compose the topology is obtained as a return parameter. To avoid stations to change index, one may simply create `ClassSwitch` nodes as last before solving the model. This node list can be employed as usual to reinstantiate new stations or `ClassSwitch` nodes. The `addLink`, `setRouting`, and possibly the `setProbRouting` functions will also need to be re-applied as described in the previous sections.

4.6.3 Saving a network object before a change

The `Network` object, and its inner objects that describe the network elements, are always passed by reference. The `copy` function should be used to clone `LINE` objects, for example before modifying a parameter for a sensitivity analysis. This function recursively clones all objects in the model, therefore creating an independent copy of the network. For example, consider the following code

```
TODO: under development
```

Using the `getName` function it is then possible to verify that `model` has now name `myModel1`, since the first assignment was by reference. Conversely, `modelByCopy.setName` did not affect the original `model` since this is a clone of the original network.

Chapter 5

Network solvers

5.1 Overview

Solvers analyze objects of class `Network` to return average, transient, distributions, or state probability metrics. A solver can implement one or more *methods*, which although featuring a similar overall solution strategy, they can differ significantly from each other in the way this strategy is actually implemented and on whether the final solution is exact or approximate.

A ‘method’ flag can be passed upon invoking a solver to specify the solution method that should be used. For example, the following invocations are identical:

```
TODO: under development
```

In what follows, we describe the general characteristics and supported model features for each solver available in LINE and their methods.

Available solvers

The following `Network` solvers are available within LINE 2.0.x:

- **LINE:** This solver uses an algorithm to select the best solution method for the model under consideration, among those offered by the other solvers. Analytical solvers are always preferred to simulation-based solvers. This solver is implemented by the `LINE` class.
- **CTMC:** This is a solver that returns the exact values of the performance metrics by explicit generation of the continuous-time Markov chain (CTMC) underpinning the model. As the CTMC typically incurs state-space explosion, this solver can successfully analyze only small models. The CTMC solver is the only method offered within LINE that can return an exact solution on all Markovian models, all other solvers are either approximate or are simulators. This solver is implemented by the `SolverCTMC` class.

- **FLUID:** This solver analyzes the model by means of an approximate fluid model, leveraging a representation of the queueing network as a system of ordinary differential equations (ODEs). The fluid model is approximate, but if the servers are all PS or INF, it can be shown to become exact in the limit where the number of users and the number of servers in each node grow to infinity [37]. This solver is implemented by the `SolverFluid` class.
- **JMT:** This is a solver that uses a model-to-model transformation to export the LINE representation into a JMT simulation (JSIM) or analytical (JMVA) models [2]. The JSIM simulation solver can analyze also non-Markovian models, in particular those involving deterministic or Pareto distributions, or empirical traces. This solver is implemented by the `SolverJMT` class.
- **MAM:** This is a matrix-analytic method solver, which relies on quasi-birth death (QBD) processes to analyze open queueing systems. This solver is implemented by the `SolverMAM` class.
- **MVA:** This is a solver based on approximate and exact mean-value analysis. This solver is typically the fastest and offers very good accuracy in a number of situations, in particular models where stations have a single-server. This solver is implemented by the `SolverMVA` class.
- **NC:** This solver uses a combination of methods based on the normalizing constant of state probability to solve a model. The underpinning algorithm are particularly useful to compute marginal and joint state probabilities in queueing network models. This solver is implemented by the `SolverNC` class.
- **SSA:** This is a discrete-event simulator based on the CTMC representation of the model. Contrary to the JMT simulator, which has online estimators for all the performance metrics, SSA estimates only the probability distribution of the system states, indirectly deriving the metrics after the simulation is completed. Moreover, the SSA execution can more efficiently parallelized on multi-core machines. Moreover, it is possible to retrieve the evolution over time of each node state, including quantities that are not loggable in JMT, e.g., the active phase of a service or arrival distribution. This solver is implemented by the `SolverSSA` class.

5.2 Solution methods

We now describe the solution methods available within the `Network` solvers. Table 5.1 provides a global summary. Some of the listed methods (e.g., `mg1`) are not associated to a specific solver, as they do not fall in one of the reference formalisms. A solver that runs these methods can be instantiated as follows, e.g.:

```
TODO: under development
```

Note that the `LINE.load` notation can also be used to instantiate a custom solver pre-configured with the specified method. For example

```
TODO: under development
```

runs the CTMC solver with default options. Solver-specific methods can be specified by appending their name to the method option, e.g. this command creates the CTMC solver with `gpu` method enabled:

TODO: under development

Table 5.1: Solution methods for Network solvers.

Solver	Method	Description	Refs.
CTMC	default	Solution based on global balance	[5, §2.1.2]
CTMC	gpu	Solution based on global balance run on GPU	–
FLUID	default	mean field fluid ODEs	[38,44]
FLUID	matrix	Alias for the default method	[38,44]
FLUID	closing	Fluid with closing method for open classes	[5][p. 507]
JMT	default	Alias for the <code>jsim</code> method	–
JMT	jmva	Alias for the <code>jmva.mva</code> method	–
JMT	jmva.mva	Exact MVA in JMVA	[41]
JMT	jmva.recal	Exact RECAL algorithm in JMVA	[19]
JMT	jmva.comom	Exact CoMoM algorithm in JMVA	[8]
JMT	jmva.amva	Approximate MVA, alias for <code>jmva.bs</code> .	–
JMT	jmva.aql	AQL algorithm in JMVA	[51]
JMT	jmva.bs	Bard-Schweitzer algorithm in JMVA	[5, §9.1.1]
JMT	jmva.chow	Chow algorithm in JMVA	[17]
JMT	jmva.dmlin	De Souza-Muntz Linearizer in JMVA	[20]
JMT	jmva.lin	Linearizer algorithm in JMVA	[16]
JMT	jmva.ls	Logistic sampling in JMVA	[10]
JMT	jsim	Exact discrete-event simulation in JSIM	[2]
MAM	default	Matrix-analytic solution of structured QBDs	[31]
MAM	dec.source	Decomposition with arrivals as from the source	–
MAM	dec.poisson	Decomposition based on Poisson arrival flows	–
MVA	default	Approximate MVA, same as <code>qd</code> option	–
MVA	amva	Approximate MVA, same as <code>qd</code> option	–
MVA	aql	Aggregate queue length approximate MVA	[51]
MVA	bs	Bard-Schweitzer approximate MVA	[5, §9.1.1]
MVA	fli	Wang-Sevcik fraction line approximate MVA	[48]
MVA	lin	Linearizer approximate MVA	[16]
MVA	qd	Queue-dependent approximate MVA	[14]

Continued on next page

Table 5.1 – Solution methods for `Network` solvers. *Continued from previous page*

Solver	Method	Description	Refs.
MVA	qdaq1	Queue-dependent Aggregate queue length (AQL) approximate MVA	–
MVA	qdlin	Queue-dependent Linearizer approximate MVA	–
MVA	qli	Wang-Sevcik queue line approximate MVA	[48]
MVA	exact	Exact solution, method depends on model	–
MVA	mva	Alias for the <code>mva.amva</code> method	[41], [7]
MVA	aba.upper	Asymptotic bound analysis (upper bounds)	[5, §9.4]
MVA	aba.lower	Asymptotic bound analysis (lower bounds)	[5, §9.4]
MVA	bjb.upper	Balanced job bounds (upper bounds)	[13, Table 3]
MVA	bjb.lower	Balanced job bounds (lower bounds)	[13, Table 3]
MVA	gb.upper	Geometric square-root bounds (upper bounds)	[13]
MVA	gb.lower	Geometric square-root bounds (lower bounds)	[13]
MVA	pb.upper	Proportional bounds (upper bounds)	[13, Table3]
MVA	pb.lower	Proportional bounds (lower bounds)	[13, Table3]
MVA	sb.upper	Simple bounds (upper bounds, Thm. 3.2, $n = 3$)	[28, Table3]
MVA	sb.lower	Simple bounds (lower bounds, Eq. 1.6)	[28, Table3]
MVA	gig1.allen	Allen-Cunneen formula - GI/G/1	[5, §6.3.4]
MVA	gig1.heyman	Heyman formula - GI/G/1	–
MVA	gig1.kingman	Kingman upper bound- GI/G/1	[5, §6.3.6]
MVA	gig1.klb	Kramer-Langenbach-Belz formula - GI/G/1	[5, §6.3.4]
MVA	gig1.kobayashi	Kobayashi diffusion approximation - GI/G/1	[5, §10.1.1]
MVA	gig1.marchal	Marchal formula - GI/G/1	[5, §10.1.3]
MVA	gigk	Kingman approximation - GI/G/k	
MVA	mg1	Pollaczek–Khinchine formula - M/G/1	[5, §3.3.1]
MVA	mml	Exact formula - M/M/1	[5, §6.2.1]
MVA	mmk	Exact formula - M/M/k (Erlang-C)	
NC	default	Alias for the adaptive method	–
NC	adaptive	Automated choice of deterministic method	–
NC	exact	Automated choice of exact solution method.	–
NC	brute	Brute force summation over all the state space	–
NC	ca	Multiclass convolution algorithm (exact)	–
NC	comom	Class-oriented method of moments (exact)	[8]
NC	cub	Grundmann-Moeller cubature rules	[10]
NC	mva	Product of throughputs on MVA lattice (exact)	[40, Eq. (47)]
NC	imci	Improved Monte carlo integration sampler	[49]
NC	kt	Knessl-Tier asymptotic expansion	[32]
NC	le	Logistic asymptotic expansion	[10]

Continued on next page

Table 5.1 – Solution methods for `Network` solvers. *Continued from previous page*

Solver	Method	Description	Refs.
NC	<code>ls</code>	Logistic sampling	[10]
NC	<code>mom</code>	Method of moments (exact)	[9]
NC	<code>mmint2</code>	McKenna-Mitra integral (2-station models only)	[36]
NC	<code>nr.logit</code>	Norlund-Rice integral with logit transformation	[12]
NC	<code>nr.probit</code>	Norlund-Rice integral with probit transformation	[12]
NC	<code>pana</code>	Panacea asymptotic expansion	[36], [42]
NC	<code>propfair</code>	Product of proportionally fair throughputs	[45]
NC	<code>rd</code>	Reduction heuristic	[12]
NC	<code>rgf</code>	Recursive generating function	[29]
NC	<code>sampling</code>	Automated choice of sampling method	–
SSA	<code>default</code>	Alias for the <code>serial</code> method	–
SSA	<code>serial</code>	CTMC stochastic simulation on a single core	[27]
SSA	<code>serial.hash</code>	<code>serial</code> with state hashing (slower, less memory)	–
SSA	<code>para</code>	Parallel simulations (independent replicas)	–
SSA	<code>para.hash</code>	<code>para</code> with state hashing (slower, less memory)	–

5.2.1 LINE

The `LINE` class, also callable with the alias `SolverAuto`, provides interfaces to the core solution functions (e.g., `getAvg`, ...) that dynamically bind to one of the other solvers implemented in `LINE` (`CTMC`, `NC`, ...). It is often difficult to identify the best solver without some performance results on the model, for example to determine if it operates in light, moderate, or heavy-load regime.

Therefore, heuristics are used to identify a solver based on structural properties of the model, such as based on the scheduling strategies used at the stations as well as the number of jobs, chains, and classes. Such heuristics, though, are independent of the core function called, thus it is possible that the optimal solver does not support the specific function called (e.g., `getTranAvg`). In such cases the `LINE` solver determines what other solvers would be feasible and prioritizes them in execution time order, with the fastest one on average having the higher priority. Eventually, the solver will be always able to identify a solution strategy, through at least simulation-based solvers such as `JMT` or `SSA`.

5.2.2 CTMC

The `SolverCTMC` class solves the model by first generating the infinitesimal generator of the `Network` and then calling an appropriate solver. Steady-state analysis is carried out by solving the global balance

equations defined by the infinitesimal generator. If the `keep` option is set to true, the solver will save the infinitesimal generator in a temporary file and its location will be shown to the user.

Transient analysis is carried out by numerically solving Kolmogorov's forward equations using MATLAB's ODE solvers. The range of integration is controlled by the `timespan` option. The ODE solver choice is the same as for `SolverFluid`.

The CTMC solver heuristically limits the solution to models with no more than 6000 states. The `force` option needs to be set to true to bypass this control. In models with infinite states, such as networks with open classes, the `cutoff` option should be used to reduce the CTMC to a finite process. If specified as a scalar value, `cutoff` is the maximum number of jobs that a class can place at an arbitrary station. More generally, a matrix assignment of `cutoff` indicates to LINE that `cutoff` has in row i and column r the maximum number of jobs of class r that can be placed at station i .

Details on the additional configuration options of the CTMC solver is given in the next table.

Table 5.2: `SolverCTMC` configuration options

Option	Value	Description
<code>options.config.hide_immediate</code>	boolean	If true, immediate transitions are hidden from the CTMC.
<code>options.config.state_space_gen</code>	reachable	Direct state space enumeration from initial state.
<code>options.config.state_space_gen</code>	full	Direct state space enumeration from all possible initial states.

5.2.3 FLUID

This solver is based on the system of fluid ordinary differential equations for INF-PS queueing networks presented in [38]. The latter is based on Kurtz's mean-field approximation theory. The fluid ODEs are normally solved with a Java port of the LSODA algorithm for stiff and non-stiff ordinary differential equations. More details about the port are available at: <https://github.com/imperial-qore/lsoda-java>.

ODE variables corresponding to an infinite number of jobs, as in the job pool of a source station, or to jobs in a disabled class are not included in the solution vector. These rules apply also to the `options.init_sol` vector.

The solution of models with FCFS stations maps these stations into corresponding PS stations where the service rates across classes are set identical to each other with a service distribution given by a mixture of the service processes of the service classes. The mixture weights are determined iteratively by solving a sequence of PS models until convergence. Upon initializing FCFS queues, jobs in the buffer are all initialized in the first phase of the service.

5.2.4 JMT

The class is a wrapper for the JMT and consists of a model-to-model transformation from the `Network` data structure into the JMT's input XML formats (either `.jsimg` or `.jmva`) and a corresponding parser for JMT's results. Upon first invocation, the JMT JAR archive will be searched in the MATLAB path and if unavailable automatically downloaded.

This solver offers two main methods. The default method is the JSIM solver (`'jsim'` method), which runs JMT's discrete-event simulator. The alternative method is the JMVA analytical solver (`'jmva'` method), which is applicable only to queueing network models that admit a product-form solution. This can be verified calling `model.hasProductFormSolution` prior to running the JMVA solver.

In the transformation to JSIM, artificial nodes will be automatically added to the routing table to represent class-switching nodes used in the simulator to specify the switching rules. One such class-switching node is defined for every ordered pair of stations (i, j) such that jobs change class in transit from i to j .

5.2.5 MAM

This is a basic solver for some Markovian open queueing systems that can be analyzed using matrix analytic methods. The core solver is based on the BU tools library for matrix-analytic methods [31]. The solution of open queueing networks is based on traffic decomposition methods that compute the arrival process at each queue resulting from the superposition of multiple source streams.

5.2.6 MVA

The solver offers approximate mean value analysis (AMVA) (`options.method='default'`), but also exact MVA algorithms (`options.method='exact'`). The default AMVA solver is based on Linearizer [16], unless there are two or less jobs in total within closed classes, in which case the solver runs the Bard-Schweitzer algorithm [45]. Extended queueing models are handled as follows:

- Non-exponential service times in FCFS nodes are handled only in the single-server case via the method selected in the `options.config.highvar` setting. By default high variance is ignored, as the FCFS solver tends to produce good result in closed models also without specialized corrections. It is alternatively possible to handle high variance either using the Diffusion-M/G/k interpolation from [11], casted with weights $a_i = b_i = 10^{-8}$, or using the high-variance MVA (HV-MVA) corrections proposed in [6, 39]. The multi-server extension is ongoing; we point to the NC solver for a version already available.
- Multi-servers are dealt with using the methods listed in Table 5.3 for the `options.config.multiserver` option. These are coupled with a modification of the Rolia-Sevcik correction [43], where in light-load the Rolia-Sevcik correction is treated as if there was a single server.
- Non-preemptive are dealt with using the methods listed in Table 5.3 for the configuration option `options.config.np_priority`. The solver feature in particular AMVA-CL and the shadow server methods [22].
- DPS queues are analyzed with a standard method similar to the biased processor sharing approximation reported in [34, §11.4]. Here, an arriving job of class r sees a queue-length in class $s \neq r$ scaled by the correction factor w_s/w_r , where w_s is the weight of class s .

- Limited load-dependence (intended here as other than multi-server) and class-dependence are handled through the correction factors proposed in [14]. If a station is both limited load-dependent and multi-server, then if the `softmin` method is chosen the solver will suitably combine the `softmin` term and the limited load-dependent correcting factors. Moreover, iterative queue-length corrections such as those applied by the AQL and Linearizer methods are also applied to these terms.
- Fork-join networks are assumed to feature a direct acyclic graph (DAG) in-between forks and joins. They are analyzed by iteratively transforming the sibling tasks into jobs belonging to independent classes, using the algorithm specified in `options.config.fork_join`. If a fork has fan out f (i.e., the fork out-degree), in the implementation of the Heidelberg-Trivedi [30] method, one artificial open class is created for each of $f - 1$ sibling task, while also retaining a task in the original class. The residence times along a branch are then treated as exponential random variables and their maximum, corresponding to the response time of the fork-join section, is computed using specialized results for this distribution. LINE supports this method, but uses as a default a custom variant whereby in which the original and artificial classes can take with probability $1/f$ any of the outgoing branches. While the latter can result in states that do not exist in the original model, since two sibling tasks may take the same branch, it is correct in expectation and it does not treat differently the artificial classes than the original class, which can be beneficial when the original class is *closed* and thus differs significantly from an *open* artificial class.

Solver-specific configuration options are reported in Table 5.3.

Table 5.3: SolverMVA configuration options

Option	Value	Description
<code>options.config.multiserver</code>	default	Equals <code>softmin</code> at PS queues and <code>seidmann</code> at FCFS queues.
<code>options.config.multiserver</code>	<code>seidmann</code>	Seidmann's decomposition [46].
<code>options.config.multiserver</code>	<code>softmin</code>	QD-AMVA's <code>softmin</code> approximation [14].
<code>options.config.np_priority</code>	default	Non-preemptive priority handling. Equals <code>cl</code> .
<code>options.config.np_priority</code>	<code>cl</code>	Chandy-Lakshmi [22].
<code>options.config.np_priority</code>	<code>shadow</code>	Sevcik's shadow server [47].
<code>options.config.highvar</code>	default	Ignored - no correction applied.
<code>options.config.highvar</code>	<code>interp</code>	Diffusion-M/G/k interpolation from [11].
<code>options.config.highvar</code>	<code>hvmva</code>	High-variance MVA as in [6], extended to multiclass as [23, Eq. 3.21].
<code>options.config.fork_join</code>	default	Equals <code>mmt</code> .
<code>options.config.fork_join</code>	<code>mmt</code>	LINE custom fork-join transformation [21]
<code>options.config.fork_join</code>	<code>ht</code>	Heidelberg-Trivedi [30]

5.2.7 NC

The `SolverNC` class implements a family of solution algorithms based on the normalizing constant of state probability of product-form queueing networks. Contrary to the other solvers, this method typically

maps the problem to certain multidimensional integrals, allowing the use of numerical methods such as MonteCarlo sampling and asymptotic expansions in their approximation.

5.2.8 SSA

The `SolverSSA` class is a basic stochastic simulator for continuous-time Markov chains. It reuses some of the methods that underpin `SolverCTMC` to generate the network state space and subsequently simulates the state dynamics by probabilistically choosing one among the possible events that can incur in the system, according to the state spaces of each of node in the network. For efficiency reasons, states are tracked at the level of individual stations, and hashed. The state space is not generated upfront, but rather stored during the simulation, starting from the initial state. If the initialization of a station generates multiple possible initial states, SSA initializes the model using the first state found. The list of initial states for each station can be obtained using the `getInitState` functions of the `Network` class.

The SSA solver offers two methods: `'serial'` and `'para'` (default). The serial methods run on a single core, while the parallel methods run on multicore

5.3 Supported language features and options

5.3.1 Solver features

Once a model is specified, it is possible to use the `getUsedLangFeatures` function to obtain a list of the features of a model. For example, the following conditional statement checks if the model contains a FCFS node

```
TODO: under development
```

Every LINE solver implements the `support` to check if it supports all language features used in a certain model

```
TODO: under development
```

5.3.2 Class functions

The table below lists the steady-state and transient analysis functions implemented by the `Network` solvers. Since the features of the LINE solver are the union of the features of the other solvers, in what follows it will be omitted from the description.

The functions listed above with the `Table` suffix (e.g., `getAvgTable`) provide results in tabular format corresponding to the corresponding core function (e.g., `getAvg`). The features of the core functions are as follows:

Table 5.4: Solver support for average performance metrics

Function	Regime	Network Solver						
		CTMC	FLUID	JMT	MAM	MVA	NC	SSA
getAvg	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgTable	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgChain	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgChainTable	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgNode	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgNodeTable	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgSys	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgSysTable	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgArvR	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgArvRChain	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgQLen	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgQLenChain	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgRespT	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgRespTChain	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgSysRespT	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgTput	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgTputChain	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgSysTput	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgUtil	Steady-state	✓	✓	✓	✓	✓	✓	✓
getAvgUtilChain	Steady-state	✓	✓	✓	✓	✓	✓	✓
getTranAvg	Transient	✓	✓	✓				

- `getAvg`: returns the mean queue-length, utilization, mean response time (for one visit), and throughput for each station and class.
- `getAvgChain`: returns the mean queue-length, utilization, mean response time (for one visit), and throughput for every station and chain.
- `getAvgSys`: returns the system response time and system throughput, as seen as the reference node, by chain.
- `getCdfRespT`: returns the distribution of response times (for one visit) for the stations at steady-state.
- `getAvgNode`: behaves similarly to `getAvg`, but returns performance metrics for each node and class. For example, throughputs at the sinks can be obtained with this method.

Table 5.5: Solver support for advanced metrics

Function	Regime	Network Solver						
		CTMC	FLUID	JMT	MAM	MVA	NC	SSA
getCdfRespT	Steady-state	✓	✓	✓	✓			
getProb	Steady-state	✓						
getProbAggr	Steady-state	✓	✓	✓		✓	✓	
getProbSys	Steady-state	✓						✓
getProbSysAggr	Steady-state	✓		✓			✓	
getProbNormConstAggr	Steady-state			✓		✓	✓	
getTranCdfPassT	Transient		✓					
getTranCdfRespT	Transient			✓				
getTranProb	Transient	✓						
getTranProbAggr	Transient	✓						
getTranProbSys	Transient	✓						
getTranProbSysAggr	Transient	✓						
sample	Transient							✓
sampleAggr	Transient			✓				✓
sampleSys	Transient							✓
sampleSysAggr	Transient			✓				✓

- `getProb`: returns state probabilities at equilibrium at a given station.
- `getProbAggr`: returns marginal state probabilities for jobs of different classes at a given station.
- `getProbSys`: returns joint probabilities for a given system state.
- `getProbSysAggr`: returns joint probabilities for jobs of different classes at all stations.
- `getProbNormConstAggr`: returns the normalizing constant of the state probabilities for the model.
- `getTranAvg`: returns transient mean queue length, utilization and throughput for every station and chain from a given initial state.
- `getTranCdfPassT`: returns the distribution of first passage times in transient regime.
- `getTranCdfRespT`: returns the distribution of response times in transient regime.
- `sample`: returns the transient marginal state for a station from a given initial state.
- `sampleAggr`: returns the transient marginal state for jobs of different classes at a given station from a given initial state.

- `sampleSys`: returns the transient marginal system state for a station from a given initial state.
- `sampleSysAggr`: returns the transient marginal system state for jobs of different classes at a given station from a given initial state.

5.3.3 Node types

The table below shows the node types supported by the different solvers. It should be noted that the `FLUID` solver is capable of handling `Sink` and `Source` nodes, but due to low accuracy when run on open models this feature is disabled in the current release.

Table 5.6: Solver support for Network nodes

Strategy	Network Solver						
	CTMC	FLUID	JMT	MAM	MVA	NC	SSA
Cache	✓				✓	✓	✓
ClassSwitch	✓	✓	✓	✓	✓	✓	✓
Delay	✓	✓	✓	✓	✓	✓	✓
Fork			✓		✓		
Join			✓		✓		
Queue	✓	✓	✓	✓	✓	✓	✓
Sink	✓		✓	✓	✓	✓	✓
Source	✓		✓	✓	✓	✓	✓

5.3.4 Scheduling strategies

The table below shows the supported scheduling strategies within `LINE` queueing stations. Each strategy belongs to a policy class:

- preemptive resume (`SchedStrategyType.PR`)
- non-preemptive (`SchedStrategyType.NP`)
- non-preemptive priority (`SchedStrategyType.NPPrio`).

The table primarily refers to invocation of the `getAvg` methods. Specialized methods, such as transient or response time distribution analysis, may be available only for a subset of the scheduling strategies supported by a solver.

Table 5.7: Solver support for scheduling strategies

Strategy	Class	Network Solver						
		CTMC	FLUID	JMT	MAM	MVA	NC	SSA
FCFS	NP	✓	✓	✓	✓	✓	✓	✓
INF	NP	✓	✓	✓	✓	✓	✓	✓
SIRO	NP	✓		✓		✓	✓	✓
SEPT	NP	✓		✓				✓
SJF	NP			✓				
HOL	NPPrio	✓		✓		✓		✓
PS	PR	✓	✓	✓	✓	✓	✓	✓
DPS	PR	✓	✓	✓		✓		✓
GPS	PR	✓		✓				✓

5.3.5 Statistical distributions

The table below summarizes the current level of support for arrival and service distributions within each solver. `Replayer` represents an empirical trace read from a file, which will be either replayed as-is by the JMT solver, or fitted automatically to a `Cox` by the other solvers. Note that JMT requires that the last row of the trace must be a number, *not* an empty row.

Table 5.8: Solver support for statistical distributions

Distribution	Network Solver						
	CTMC	FLUID	JMT	MAM	MVA	NC	SSA
APH	✓		✓	✓	✓	✓	✓
Coxian	✓	✓	✓	✓	✓	✓	✓
Exp	✓	✓	✓	✓	✓	✓	✓
Erlang	✓	✓	✓	✓	✓	✓	✓
HyperExp	✓	✓	✓	✓	✓	✓	✓
Disabled	✓	✓	✓	✓	✓	✓	✓
Det			✓				
Gamma			✓				
Lognormal			✓				
Pareto			✓				
Replayer			✓				
Uniform			✓				
Weibull			✓				

5.3.6 Solver options

Solver options are encoded in LINE in a structure array that is internally passed to the solution algorithms.

This can be specified as an argument to the constructor of the solver. For example, the following two constructor invocations are identical

```
TODO: under development
```

Modifiers to the default options can either be specified directly in the `options` data structure, or alternatively be specified as argument pairs to the constructor, i.e., the following two invocations are equivalent

```
TODO: under development
```

Available solver options are as follows:

- `cache (logical)` if set to true the solver after the first invocation will return the same result upon subsequent calls, without solving again the model. This option is true by default. Caching can be bypassed using the `refresh` methods (see Section 4.6).
- `config (struct)` this is data structure to pass solver-specific configuration options to customize the execution of particular methods.
- `cutoff (integer ≥ 1)` requires to ignore states where stations have more than the specified number of jobs. This is a mandatory option to analyze open classes using the CTMC solver.
- `force (logical)` requires the solver to proceed with analyzing the model. This bypasses checks and therefore can result in the solver either failing or requiring an excessive amount of resources from the system.
- `iter_max (integer ≥ 1)` controls the maximum number of iterations that a solver can use, where applicable. If `iter_max = n`, this option forces the FLUID solver to compute the ODEs over the timespan $t \in [0, 10n/\mu^{\min}]$, where μ^{\min} is the slowest service rate in the model. For the MVA solver this option instead regulates the number of successive substitutions allowed in the fixed-point iteration.
- `iter_tol (double)` controls the numerical tolerance used to convergence of iterative methods. In the FLUID solver this option regulates both the absolute and relative tolerance of the ODE solver.
- `init_sol (solver dependent)` re-initializes iterative solvers with the given configuration of the solution variables. In the case of MVA, this is a matrix where element (i, j) is the mean queue-length at station i in class j . In the case of FLUID, this is a model-dependent vector with the values of all the variables used within the ODE system that underpins the fluid approximation.
- `keep (logical)` determines if the model-to-model transformations store on file their intermediate outputs. In particular, if `verbose ≥ 1` then the location of the `.jsimg` models sent to JMT will be printed on screen.

- `method` (string) configures the internal algorithm used to solve the model.
- `samples` (integer ≥ 1) controls the number of samples collected *for each* performance index by simulation-based solvers. JMT requires a minimum number of samples of $5 \cdot 10^3$ samples.
- `seed` (integer ≥ 1) controls the seed used by the pseudo-random number generators. For example, simulation-based solvers will give identical results across invocations only if called with the same seed.
- `stiff` (logical) requires the solver to use a stiff ODE solver.
- `timestamp` (real interval) requires the transient solver to produce a solution in the specified temporal range. If the value is set to `[Inf, Inf]` the solver will only return a steady-state solution. In the case of the FLUID solver and in simulation, `[Inf, Inf]` has the same computational cost of `[0, Inf]` therefore the latter is used as default.
- `tol` default numerical tolerance for all uses other than the ones where `iter_tol` is used.
- `verbose` controls the verbosity level of the solver. Supported levels are 0 for silent, 1 for standard verbosity, 2 for debugging.

Table 5.9: Default values of the LINE solver options and their default assignments

Option	Solver default						
	MVA	CTMC	FLUID	JMT	MAM	NC	SSA
cache	true	true	true	true	true	true	true
config							
cutoff		(no default)					
force	false	false	false	false	false	false	false
keep				false			
init_sol	[]		[]				
iter_max	10^3		10				
iter_tol	10^{-6}		10^{-4}		10^{-4}		
method	'default'	'default'	'default'	'default'	'default'	'default'	'default'
lang	'matlab'	'matlab'	'matlab'	'matlab'	'matlab'	'matlab'	'matlab'
samples				10^4			10^4
seed	rand	rand	rand	rand	rand	rand	rand
stiff			true				
timespan		[Inf, Inf]	[0, Inf]	[0, Inf]		[Inf, Inf]	[0, Inf]
tol		10^{-4}	10^{-4}				
verbose	1	1	1	1	1	1	1

5.4 Solver maintenance

The following best practices can be helpful in maintaining the LINE installation:

- To install a new release of JMT, it is necessary to delete (or overwrite) the `JMT.jar` file under the `'SolverJMT'` folder. This forces LINE to download the latest version of the JMT executable.
- To remove temporary by-products of the JMT solver it is recommended to periodically run the `jmtCleanTempDir` script. This is more important when using the `'keep'` option, which stores on disk the temporary `.jsimg` and `.jsimw` models sent to JMT.

Chapter 6

Layered network models

In this chapter, we present the definition of the `LayeredNetwork` class, which encodes the support in LINE for a class of generalized layered stochastic networks. In their basic form, these models are called layered queueing networks (LQNs) and differ from regular queueing networks as servers, in order to process jobs, can issue synchronous and asynchronous calls among each others. We point to [24] and to the LQNS user manual for an introduction [25]. Contrary to the original LQNs, layered networks in LINE can also include non-queueing servers, such as caches, hence they may be conceptualized as more general layered stochastic networks.

The topology of call dependencies in a layered network makes it possible to partition the model into a set of layers, each consisting of a subset of the servers. Each of these layers is then solved in isolation, updating with an iterative procedure its parameters and performance metrics until the layers solutions jointly converge to a consistent solution.

6.1 Basics about layered networks

Layered network models describe a collection of resources called *tasks*, each representing for example a software server, that run on resources called *host processors*. Classes of service exposed by a task are called *entries*. Each entry is an endpoint at which a task can be invoked; for example, if a task represents a web server then its web pages may be described as different entries.

A special task, called the *reference task* is used to represent a group of system users. In this case, the host processor for a reference task can either be real, as in the case of users that are themselves software systems, or fictitious, as in the case of human users.

Each entry can be specified by a workflow of operations called *activities*, typically organized as a directed acyclic graph. The time demand that each activity places at the underpinning host processor is called its *host demand* and it is a random variable with a user-specified distribution.

Activity graphs may include *calls* to entries exposed by other tasks. This is an abstraction of the calls that distributed system components have among themselves. Calls can either be *synchronous*, *asynchronous*, or

forwarding. At present, LINE supports only the first two kinds of activities. Synchronous calls are requests that block the sender until a reply is received, while asynchronous calls are non-blocking and the sender execution can continue after issuing the call. Calls can either be repeated either *deterministic* or *stochastic*, meaning in the latter case that the number of calls issued is a random variable, e.g. geometrically distributed.

Contrary to ordinary layered queueing networks, a layered network in LINE can also feature *cache tasks*, *item entries*, and *cache-access* precedence relations.

- Cache tasks have the basic properties of tasks, but add three specific properties for caching: the total number of items, the cache capacity and the cache replacement policy. Cached items can be either contents or services. Cache capacity indicates the storage constraints of the cache.
- An item-entry provides instead access to a group of entries of a cache, Item-entries have the basic properties of entries, but add the property of the popularity of the items they give access to.
- A precedence relationship called *cache-access* is defined for the cache hit and miss activities under each item-entry. That is, it is possible to proceed to a different activity depending on whether the cache access produced a cache hit or cache miss. For example, a cache miss can produce a call to a remote entry to retrieve the missing content.

Note that the above extensions are not queueing-based and this explains why these models are referred to in LINE as layered networks and not as layered queueing networks. Similar to the latter, the analysis of a layered networks uses a decomposition of the model into a set of submodels, each being a `Network` object, which are then iterative analyzed using different solution methods.

6.2 LayeredNetwork object definition

6.2.1 Creating a layered network topology

A layered queueing network consists of four types of elements: processors, tasks, entries and activities. An entry is a class of service specified through a finite sequence of activities, and hosted by a task running on a (physical) processor. A task is typically a software queue that models access to the capacity of the underpinning processor. Activities model either demands required at the underpinning processor, or calls to entries exposed by some remote tasks.

In the `LayeredNetwork` class, the terms *host* and *processor* are entirely interchangeable.

To create our first layered network, we instantiate a new model as

```
TODO: under development
```

We now proceed to instantiate the static topology of processors, tasks and entries:

```
TODO: under development
```

An equivalent way to specify the above example is to use the `Host` class instead than the `Processor` class, with identical parameters.

In the above code, the `on` method specifies the associations between the elements, e.g., task `T1` runs on processor `P1`, and accepts calls to entry `E1`. Furthermore, the multiplicity of `T1` is 5, meaning that up to 5 calls can be simultaneously served by this element (i.e., 5 is the multiplicity of servers in the underpinning queueing system for `T1`).

Both processors and tasks can be associated to the standard `LINE` scheduling strategies. For instance, `T2` will process incoming requests in parallel according as an infinite server node, since we selected the `SchedStrategy.INF` scheduling policy. An exception is that `SchedStrategy.REF` should be used to denote the reference task (e.g. a node representing the clients of the models), which has a similar meaning to the reference node in the `Network` object.

6.2.2 Describing host demands of entries

The demands placed by an entry on the underpinning host (also called in layered queueing networks the *host demand*) is described in terms of execution of one or more activities. Although in tools such as `LQNS` activities can be associated to either entries or tasks, `LINE` supports only the more general of the two options, i.e., the definition of activities at the level of tasks. In this case:

- Every task defines a collection of activities.
- Every entry needs to specify an initial activity where the execution of the entry starts (the activity is said to be “bound to the entry”) and a replying activity, which upon completion terminates the execution of the entry.

For example, in our running example, we may now associate an activity to each entry as follows:

`TODO: under development`

Here, `A1` is a task activity for `T1`, acts as initial activity for `E1`, consumes an exponential distributed time on the processor underpinning `T1`, and requires on average 3.5 synchronous calls to `E2` to complete. Each call to entry `E2` is served by the activity `A2`, with a demand on the processor hosting `T2` given by an exponential distribution with rate $\lambda = 2.0$.

Activity graphs

Often, it is useful to structure the sequence of activities carried out by an entry in a graph. Activity graphs can be characterized by precedence relationships of the following kinds:

- *sequence*: two activities are executed sequentially, one after each other. This is implemented through the `ActivityPrecedence.Serial` construct.
- *loop*: an activity is repeated a number of times. This is implemented in `ActivityPrecedence.Loop`.

- *and-fork*: a serial execution is forked into concurrent activities. This can be materialized using the `ActivityPrecedence.AndFork` construct.
- *or-fork*: the server chooses probabilistically which activity to execute next among a set of alternatives. This is implemented in `ActivityPrecedence.OrFork`.
- *and-join*: concurrent activities are joined into a single serial execution. This is implemented in `ActivityPrecedence.AndJoin`.
- *or-join*: merge point for alternative activities that may execute in parallel after a *or-fork*. This is implemented in `ActivityPrecedence.OrFork`.
- *cache-access*: split point for cache hit/cache miss results in an activity graph. This is implemented in `ActivityPrecedence.CacheAccess`. For usage examples, see `example.cacheModel_3` and `example.cacheModel_4` in the `examples/` folder.

A composite example showing fork/join precedences and loops is given in `example.layeredModel_7` in the `examples/` folder.

For instance, we may replace in the running example the specification of the activities underpinning a call to E2 as

```
TODO: under development
```

such that a call to E2 serially executes A20, A21, and A22 prior to replying. Here, A21 is chosen to be an Erlang distribution with given mean (1.0) and number of phases (2).

6.2.3 Debugging and visualization

The structure of a `LayeredNetwork` object can be graphically visualized as follows

```
TODO: under development
```

An example of the result is shown in the next figure. The figure shows two processors (P1 and P2), two tasks (T1 and T2), and three entries (E1, E2, and E3) with their associated activities. Both dependencies and calls are both shown as directed arcs, with the edge weight on call arcs corresponding to the average number of calls to the target entry. For example, A1 calls E3 on average 2.0 times. In this figure, by clicking with the mouse on a node MATLAB will display, using a data tip, some relevant properties of the node such as scheduling or multiplicity.

When invoked as

```
TODO: under development
```

the plot is also accompanied by a task graph that illustrates the client-server dependencies between the layers.

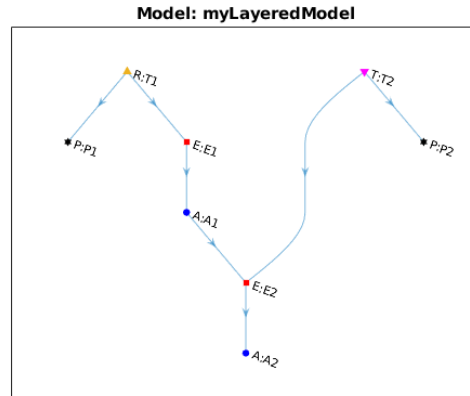


Figure 6.1: LayeredNetwork.plot method

Lastly, the `jsimView` and `jsimwView` methods can be used to visualize in JMT each layer. This can be done by first calling the `getLayers` method to obtain a cell array consisting of the `Network` objects, each one corresponding to a layer, and then invoking the `jsimView` and `jsimwView` methods on the desired layer. This is discussed in more details in the next section.

6.3 Internals

6.3.1 Representation of the model structure

It is possible to access the internal representation of a `LayeredNetwork` model in a similar way as for `Network` objects, i.e.:

TODO: under development

The return `lqn` structure, of class `LayeredNetworkStruct`, contains all the information about the specified model. It relies on relative and absolute indexing for the elements of the `LayeredNetwork`.

- A *relative* index is a number between 1 and the number of similar elements in the model, e.g., for a model with 3 tasks, the relative index t of a task would be a number in $[1, 3]$.
- An *absolute* index is a number between 1 and the total number of elements (of any kind, except calls) in the model, e.g., for a model with 2 hosts, 3 tasks, 5 entries, and 8 activities, the total number of elements is `nidx= 18` and last activity a may have an absolute index `aidx= 18` and a relative index `a= 8`.

- The difference between the relative and the absolute index of an element is referred to as *shift*, e.g., in the previous example $a_{\text{shift}} = 18 - 8 = 10$.
- Absolute and relative indexing for calls and hosts are identical, call index `cidx` ranges in $[1, n_{\text{calls}}]$ and host index `hidx` ranges in $[1, n_{\text{hosts}}]$.

Using the above convention, the internal representation of the model is described in Table 6.3.1. As in the examples above, relative and absolute indexes are differentiated by using the suffix `idx` in the latter (e.g., `a` vs. `aidx`). This indexing style is used throughout the codebase as well.

6.3.2 Decomposition into layers

Layers are a form of decomposition where we model the performance of one or more servers. The activity of clients not detailed in that layer is taken into account through an artificial delay station, placed in a closed loop to the servers [43]. This artificial delay is used to model the inter-arrival time between calls issued by that client.

The current version of LINE adopts SRVN-type layering [25], whereby a layer corresponds to one and only one resource, either a processor or a task. The `getLayers` method returns a cell array consisting of the `Network` objects corresponding to each layer

TODO: under development

The decomposition is performed through the LN solver described later.

Within each layer, classes are used to model the time a job spends in a given activity or call, with synchronous calls being modeled by classed with label including an arrow, e.g., `'AS1=>E3'` is a closed class used represent synchronous calls from activity `AS1` to entry `E3`, whereas `'AS1->E3'` denotes an asynchronous call. Artificial delays and reference nodes are modelled as a delay station named `'Clients'`, whereas the task or processor assigned to the layer is modelled as the other node in the layer.

6.4 Solvers

LINE offers two solvers for the solution of a `LayeredNetwork` model consisting in its own native solver (LN) and a wrapper (LQNS) to the LQNS solver [25]. The latter requires a distribution of LQNS to be available on the operating system command line.

The solution methods available for `LayeredNetwork` models are similar to those for `Network` objects. For example, the `getAvgTable` can be used to obtain a full set of mean performance indexes for the model, e.g.,

TODO: under development

Table 6.1: LayeredNetworkStruct static properties

Field	Type	Description
nidx	integer	Total number of LayeredNetwork elements
nhosts	integer	Number of Hosts or Processor elements
ntasks	integer	Number of Tasks elements
nentries	integer	Number of Entry elements
nacts	integer	Number of Activity elements
ncalls	integer	Number of calls issued by Activity elements
hshift	integer	For host h , the value $h+hshift$ returns its absolute index in $1\dots nidx$
tshift	integer	For task t , the value $t+tshift$ returns its absolute index in $1\dots nidx$
eshift	integer	For task e , the value $e+eshift$ returns its absolute index in $1\dots nidx$
ashift	integer	For activity a , the value $a+ashift$ returns its absolute index in $1\dots nidx$
cshift	integer	For call c , the value $c+cshift$ returns its absolute index in $1\dots ncalls$
tasksof{ $hidx$ }	array	Array of absolute indexes for tasks on host with absolute index $hidx$
entriesof{ $tidx$ }	array	Array of absolute indexes for entries on task with absolute index $tidx$
actsof{ $tidx$ }	array	Array of absolute indexes for activities on task with absolute index $tidx$
callsof{ $aidx$ }	array	Array of absolute indexes for calls on activity with absolute index $aidx$
hostdem{ idx }	object	Object for the host demand distribution for element idx
think{ idx }	object	Object for the think time distribution for element idx
sched{ idx }	char	SchedStrategy scheduling policy for a host or task with absolute index idx
schedid{ idx }	integer	SchedStrategy scheduling policy id for a host or task with absolute index idx
names{ idx }	char	Name of the element with absolute index idx
hashnames{ idx }	char	Same as names{ idx } with type prefix (“H:” for host, “R:” for reference task, “T:” for task, “C:” for cache task, “E:” for entry, “I:” for item entry, “A:” for activity)
mult{ idx }	integer	Multiplicity for a host or task with absolute index idx
repl{ idx }	integer	Replication factor for a host or task with absolute index idx
type{ idx }	integer	LayeredNetworkElement id for an element with absolute index idx
nitems{ idx }	integer	Total number of items in CacheTask or ItemEntry idx
itemcap{ idx }(l)	integer	CacheTask idx capacity for cache list l
replacement{ $tidx$ }	integer	ReplacementPolicy id for cache task with absolute index $tidx$
itemproc{ idx }	object	DiscreteDistrib object for for the item popularity in ItemEntry idx
calltype{ $cidx$ }	integer	CallType id for call $cidx$
callpair{ $cidx$, j }	integer	$j = 1$ is the activity $aidx$ issuing call $cidx$, $j = 2$ is the entry $eidx$ called by $cidx$
callproc{ $cidx$ }	logical	DiscreteDistrib object for the number of calls with absolute index $cidx$
callnames{ $cidx$ }	char	Automatically assigned name to call with absolute index $cidx$
callhashnames{ $cidx$ }	char	Same as callnames{ $cidx$ } but with type prefix for source and destination
actpretype{ $aidx$ }	integer	ActivityPrecedenceType id preceding the activity with absolute index $aidx$
actposttype{ $aidx$ }	integer	ActivityPrecedenceType id after the activity with absolute index $aidx$
graph{ idx_i , idx_j }	float	$\neq 0$ if idx_i “runs on”, “calls” or “precedes” idx_j . Values are by default 1.0 unless they specify a mean number of calls, a probability, fork fan-out or join fan-in degrees
parent{ idx }	integer	Return absolute index value of the parent node for the node whose absolute index value is idx
replygraph{ $aidx$, $eidx$ }	logical	True if activity with absolute index $aidx$ on entry $eidx$ replies ending the entry call
taskgraph{ $tidx_i$, $tidx_j$ }	logical	True if host or task idx_i is a caller of host or task idx_j
iscache{ $tidx$ }	logical	True if task $tidx$ is a CacheTask
iscaller{ idx_i , idx_j }	logical	True if task or activity idx_i calls task or entry idx_j
issynccaller{ idx_i , idx_j }	logical	True if task or activity idx_i issues a synchronous call to task or entry idx_j
isasynccaller{ idx_i , idx_j }	logical	True if task or activity idx_i issues an asynchronous call to task or entry idx_j
isref{ $tidx$ }	logical	True if task $tidx$ is a reference task

Note that in the above table, some performance indexes are marked as NaN because they are not defined in a layered queueing network. Further, compared to the `getAvgTable` method in `Network` objects, `LayeredNetwork` do not have an explicit differentiation between stations and classes, since in a layer a task may either act as a server station or a client class.

The main challenge in solving layered queueing networks through analytical methods is that the parameterization of the artificial delays depends on the steady-state performance of the other layers, thus causing a cyclic dependence between input parameters and solutions across the layers. Depending on the solver in use, such issue can be addressed in a different way, but in general a decomposition into layers will remain parametric on a set of response times, throughputs and utilizations.

This issue can be resolved through solvers that, starting from an initial guess, cyclically analyze the layers and update their artificial delays on the basis of the results of these analyses. Both `LN` and `LQNS` implement this solution method. Normally, after a number of iterations the model converges to a steady-state solution, where the parameterization of the artificial delays does not change after additional iterations.

6.4.1 LQNS

The `LQNS` wrapper operates by first transforming the specification into a valid `LQNS` XML file. Subsequently, `LQNS` calls the solver and parses the results from disks in order to present them to the user in the appropriate `LINE` tables or vectors. The `options.method` can be used to configure the `LQNS` execution as follows:

- `options.method='std'` or `'lqns'`: `LQNS` analytical solver with default settings.
- `options.method='exact'`: the solver will execute the standard `LQNS` analytical solver with the exact MVA method.
- `options.method='srvn'`: `LQNS` analytical solver with `SRVN` layering.
- `options.method='srvnexact'`: the solver will execute the standard `LQNS` analytical solver with `SRVN` layering and the exact MVA method.
- `options.method='lqsim'`: `LQSIM` simulator, with simulation length specified via the `samples` field (i.e., with parameter `-A options.samples, 0.95`).

Upon invocation, the `lqns` or `lqsim` commands will be searched for in the system path. If they are unavailable, the termination of `SolverLQNS` will interrupt.

6.4.2 QNS

`LINE` also includes a dedicated wrapper `Network` solver for the `qnsolver` utility distributed within `LQNS`, called `SolverQNS`. This allows users to evaluate product-form models using the MVA algorithms implemented within `LQNS`. The available options specify the multiserver handling algorithm:

- `options.method='conway'`: Conway's multiserver approximation within the Linearizer algorithm proposed in [18].
- `options.method='rolia'`: Rolia's multiserver in the Methods of Layers paper [43].
- `options.method='reiser'`: load-dependent mean-value analysis as described originally by Reiser-Lavenberg in [41].
- `options.method='zhou'`: Zhou-Woodside's multiserver approximation in [52].

6.4.3 LN

The native LN solver iteratively applies the layer updates until convergence of the steady-state measures. Since updates are parametric on the solution of each layer, LN can apply any of the Network solvers described in the solvers chapter to the analysis of individual layers, as illustrated in the following example for the MVA solver

```
TODO: under development
```

Options parameters may also be omitted. The LN method converges when the maximum relative change of mean response times across layers from the last iteration is less than `options.iter_tol`.

Methods supported by the LN solver include:

- `options.method='default'`: default recursive solution based on mean values
- `options.method='moment3'`: solution by recursive 3-moment approximation of response time distributions.

6.5 Model import and export

A `LayeredNetwork` can be easily read from, or written to, a XML file based on the LQNS meta-model format¹. The read operation can be done using a static method of the `LayeredNetwork` class, i.e.,

```
TODO: under development
```

Conversely, the write operation is invoked directly on the model object

```
TODO: under development
```

In both examples, `filename` is a string including both file name and its path.

Finally, we point out that it is possible to export a LQN in the legacy SRVN file format² by means of the `writeSRVN(filename)` function.

¹<https://raw.githubusercontent.com/layeredqueueing/V5/master/xml/lqn.xsd>

²<http://www.sce.carleton.ca/rads/lqns/lqn-documentation/format.pdf>

Chapter 7

Random environments

Systems modeled with LINE can be described as operating in an environment with a state that affects the way the system dynamics. To distinguish the states of the environment from the ones of the system within it, we shall refer to the former as the environment *stages*. In particular, LINE 2.0.0 supports the definition of a class of random environments subject to three assumptions:

- The stage of the environment evolves independently of the state of the system.
- The dynamics of the environment stage can be described by a continuous-time Markov chain.
- The topology of the system is independent of the environment stage.

The above definitions are in particular appropriate to describe systems specified by input parameters (e.g., service rates, scheduling weights, etc) that change with the environment stage. For example, an environment with two stages, say normal load and peak load, may differ for the number of servers that are available in a queueing station, i.e., the system controller may add more servers during peak load. Upon a stage change in the environment, the model parameters will instantaneously change, and the system state reached during the previous stage will be used to initialize the system in the new stage.

Although in a number of cases the system performance may be similar to a weighted combination of the average performance in each stage, this is not true in general, especially if the system dynamic (i.e., the rate at which jobs arrive and get served) and the environment dynamic (i.e., the rate at which the environment changes active stage) have a similar magnitude [15].

7.1 Environment object definition

7.1.1 Specifying the environment

In LINE, an environment is internally described by a Markov renewal process (MRP) with transition times belonging to the `PhaseType` class. A MRP is similar to a Markov chain, but state transitions are not

restricted to be exponential. Although the time spent in each state of the MRP is not exponential, the MRP with phase-type transitions can be easily transformed into an equivalent continuous-time Markov chain (CTMC) to enable analysis, a task that `LINE` performs automatically.

To specify an environment, we first create an `Env` object with the environment name

```
TODO: under development
```

We then add two stages

```
TODO: under development
```

where the constructor specifies the stage name, an arbitrary string to classify the stage (here taken from a taxonomy in the `Semantics` class), follows by a `Network` object describing the system model conditional on the environment being in the corresponding stage.

We now describe that the transitions between stages are both exponential, with different rates

```
TODO: under development
```

We can also add a self-loop on the online stage as follows

```
TODO: under development
```

which would cause a race condition between two distributions in stage two: the exponential transition back to the offline stage, and the Erlang-2 distributed transition with unit rate that remains in the online stage. The underpinning CTMC will therefore consider the distribution of the minimum between the exponential and the Erlang-2 distribution, in order to decide the next stage transition. State space explosion may occur in the definition of an environment if the user specifies a large number of non-exponential transition. For example, a race condition among n Erlang-2 distribution translates at the level of the CTMC into a state space with 2^n states. In such situations, it is recommended to replace some of the distributions with exponential ones.

To summarize the properties of the environment defined above we may use the `getStageTable` method

```
TODO: under development
```

In the table, the `State` column gives a numerical identifier for each stage, followed by its stage probability at equilibrium, a Markovian representation of the time spent in it before a transition, and by a pointer to the sub-model associated to that stage.

7.1.2 Specifying a reset policy

When the environment transitions, the default policy is that the associated model is re-initialized using the marginal queue-length values observed at departure instants. This means in practice that jobs in execution at a server are required all to restart execution at that server upon occurrence of a transition. This may not

be possible in some models, for example when a station is removed from the model. In that case, one can define a custom reset policy by instantiating transitions as, e.g.,

```
TODO: under development
```

7.1.3 Specifying system models for each stage

LINE places loose assumptions in the way the system should be described in each stage. It is just expected that the user supplies a model object, either a `Network` or a `LayeredNetwork`, in each stage, and that a transient analysis method is available in the chosen solver, a requirement fulfilled for example by `SolverFluid`.

However, we note that the model definition can be somewhat simplified if the user describes the system model in a separate MATLAB function, accepting the stage-specific parameters in input to the function. This enables reuse of the system topology across stages, while creating independent model objects.

7.2 Solvers

The steady-state analysis of a system in a random environment is carried out in LINE using the blending method [15], which is an iterative algorithm leveraging the transient solution of the model. In essence, the model looks at the *average* state of the system at the instant of each stage transition, and upon restarting the system in the new stage re-initializes it from this average value. This algorithm is implemented in LINE by the `SolverEnv` class, which is described next.

7.2.1 ENV

The `SolverEnv` class applies the blending algorithm by iteratively carrying out a transient analysis of each system model in each environment stage, and probabilistically weighting the solution to extract the steady-state behavior of the system.

As in the transient analysis of `Network` objects, LINE does not supply a method to obtain mean response times, since Little's law does not hold in the transient regime. To obtain the mean queue-length, utilization and throughput of the system one can call as usual the `getAvg` method on the `SolverEnv` object, e.g.,

```
TODO: under development
```

Note that as model complexity grows, the number of iterations required by the blending algorithm to converge may grow large. In such cases, the `options.iter_max` option may be used to bound the maximum analysis time.

Bibliography

- [1] S. Balsamo. Product form queueing networks. In Günter Haring, Christoph Lindemann, and Martin Reiser, editors, *Performance Evaluation: Origins and Directions*, volume 1769 of *Lecture Notes in Computer Science*, pages 377–401. Springer, 2000.
- [2] M. Bertoli, G. Casale, and G. Serazzi. The JMT simulator for performance evaluation of non-product-form queueing networks. In *Proc. of the 40th Annual Simulation Symposium (ANSS)*, pages 3–10, 2007.
- [3] D. Bini, B. Meini, S. Steffé, J. F. Pérez, and B. Van Houdt. Smcsolver and q-mam: tools for matrix-analytic methods. *SIGMETRICS Performance Evaluation Review*, 39(4):46, 2012.
- [4] A. Bobbio, A. Horváth, M. Scarpa, and M Telek. Acyclic discrete phase type distributions: properties and a parameter estimation algorithm. *Perform. Eval.*, 54(1):1–32, 2003.
- [5] G. Bolch, S. Greiner, H. de Meer, and K. S. Trivedi. *Queueing Networks and Markov Chains*. Wiley, 2006.
- [6] A. B. Bondi and W. Whitt. The influence of service-time variability in a closed network of queues. *Perform. Eval.*, 6:219–234, 1986.
- [7] S. C. Bruell, G. Balbo, and P. V. Afshari. Mean value analysis of mixed, multiple class BCMP networks with load dependent service stations. *Performance Evaluation*, 4:241–260, 1984.
- [8] G. Casale. CoMoM: Efficient class-oriented evaluation of multiclass performance models. *IEEE Trans. on Software Engineering*, 35(2):162–177, 2009.
- [9] G. Casale. Exact analysis of performance models by the method of moments. *Perform. Eval.*, 68(6):487–506, 2011.
- [10] G. Casale. Accelerating performance inference over closed systems by asymptotic methods. In *Proc. of ACM SIGMETRICS*. ACM Press, 2017.
- [11] G. Casale. Integrated Performance Evaluation of Extended Queueing Network Models with Line. In *2020 Winter Simulation Conference (WSC)*, pages 2377–2388. IEEE, dec 2020.

- [12] G. Casale, P.G. Harrison, and O.W. Hong. Facilitating load-dependent queueing analysis through factorization. *Perform. Eval.*, 2021.
- [13] G. Casale, Richard R. Muntz, and Giuseppe Serazzi. Geometric bounds: A noniterative analysis technique for closed queueing networks. *IEEE Trans. Computers*, 57(6):780–794, 2008.
- [14] G. Casale, J. F. Pérez, and W. Wang. QD-AMVA: Evaluating systems with queue-dependent service requirements. In *Proceedings of IFIP PERFORMANCE*, 2015.
- [15] G. Casale, M. Tribastone, and P. G. Harrison. Blending randomness in closed queueing network models. *Perform. Eval.*, 82:15–38, 2014.
- [16] K. M. Chandy and D. Neuse. Linearizer: A heuristic algorithm for queueing network models of computing systems. *Comm. of the ACM*, 25(2):126–134, 1982.
- [17] W.-M. Chow. Approximations for large scale closed queueing networks. *Perform. Eval*, 3(1):1–12, 1983.
- [18] A. E. Conway. *Fast Approximate Solution of Queueing Networks with Multi-Server Chain-Dependent FCFS Queues*, pages 385–396. Springer US, Boston, MA, 1989.
- [19] A. E. Conway and N. D. Georganas. RECAL - A new efficient algorithm for the exact analysis of multiple-chain closed queueing networks. *JACM*, 33(4):768–791, 1986.
- [20] E. de Souza e Silva and R. R. Muntz. A note on the computational cost of the linearizer algorithm for queueing networks. *IEEE Trans. Computers*, 39(6):840–842, 1990.
- [21] Rares-Andrei Dobre, Zifeng Niu, and Giuliano Casale. Approximating fork-join systems via mixed model transformations. In *Companion of the 15th ACM/SPEC International Conference on Performance Engineering, ICPE '24 Companion*, page 273–280, New York, NY, USA, 2024. Association for Computing Machinery.
- [22] D. L. Eager and J. N. Lipscomb. The AMVA priority approximation. *Performance Evaluation*, 8(3):173–193, 1988.
- [23] G. Franks. *Performance Analysis of Distributed Server Systems*. PhD thesis, Carleton, 1996.
- [24] G. Franks, T. Al-Omari, M. Woodside, O. Das, and S. Derisavi. Enhanced modeling and solution of layered queueing networks. *Software Engineering, IEEE Transactions on*, 35(2):148–161, 2009.
- [25] G. Franks, P. Maly, C. M. Woodside, D. C. Petriu, A. Hubbard, and M. Mroz. *Layered Queueing Network Solver and Simulator User Manual*, 2012.
- [26] N. Gast and B. Van Houdt. Transient and steady-state regime of a family of list-based cache replacement algorithms. *Queueing Syst*, 83(3-4):293–328, 2016.

- [27] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81(25):2340–2361, 1977.
- [28] A. Harel, S. Namn, and J. Sturm. Simple bounds for closed queueing networks. *Queueing Systems*, 31(1-2):125–135, 1999.
- [29] P. G. Harrison and T. T. Lee. A new recursive algorithm for computing generating functions in closed queueing networks. In *Proc. of IEEE MASCOTS Symposium*, pages 223–230. IEEE Press, 2004.
- [30] P. Heidelberger and K. Trivedi. Queueing network models for parallel processing with asynchronous tasks. *IEEE Transactions on Computers*, 100(11):1099–1109, 1982.
- [31] G. Horváth and M. Telek. Butools 2: A rich toolbox for markovian performance evaluation. In *Proc. of VALUETOOLS*, pages 137–142, ICST, Brussels, Belgium, Belgium, 2017. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).
- [32] C. Knessl and C. Tier. Asymptotic expansions for large closed queueing networks with multiple job classes. *IEEE Trans. Computers*, 41(4):480–488, 1992.
- [33] S. S. Lavenberg. A perspective on queueing models of computer performance. *Perform. Eval.*, 10(1):53–76, 1989.
- [34] E. D. Lazowska, J. Zahorjan, G. S. Graham, and K. C. Sevcik. *Quantitative System Performance*. Prentice-Hall, 1984.
- [35] KT Marshall. Some relationships between the distributions of waiting time, idle time and interoutput time in the gi/g/1 queue. *SIAM Journal on Applied Mathematics*, 16(2):324–327, 1968.
- [36] J. McKenna and D. Mitra. Asymptotic expansions and integral representations of moments of queue lengths in closed markovian networks. *JACM*, 31(2):346–360, April 1984.
- [37] J. F. Pérez and G. Casale. Assessing SLA compliance from Palladio component models. In *Proceedings of the 2nd MICAS*, 2013.
- [38] J. F. Pérez and G. Casale. Line: Evaluating software applications in unreliable environments. *IEEE Transactions on Reliability*, 66(3):837–853, Sept 2017.
- [39] M. Reiser. A queueing network analysis of computer communication networks with window flow control. *Communications, IEEE Transactions on*, 27(8):1199–1209, 1979.
- [40] M. Reiser. Mean-value analysis and convolution method for queue-dependent servers in closed queueing networks. *Perform. Eval.*, 1:7–18, 1981.
- [41] M. Reiser and S. Lavenberg. Mean-value analysis of closed multichain queueing networks. *Journal of the ACM*, 27:313–322, 1980.

- [42] T. G. Robertazzi. *Computer Networks and Systems*. Springer, 2000.
- [43] J. A. Rolia and K. C. Sevcik. The method of layers. *IEEE Transactions on Software Engineering*, 21(8):689–700, August 1995.
- [44] J. Ruuskanen, T. Berner, K.-E. Årzén, and A. Cervin. Improving the mean-field fluid model of processor sharing queueing networks for dynamic performance models in cloud computing. *Perform. Evaluation*, 151:102231, 2021.
- [45] P. J. Schweitzer. Approximate analysis of multiclass closed networks of queues. In *Proc. of the Int’l Conf. on Stoch. Control and Optim.*, pages 25–29, Amsterdam, 1979.
- [46] A. Seidmann, P. J. Schweitzer, and S. Shalev-Oren. Computerized closed queueing network models of flexible manufacturing systems: A comparative evaluation. *Large Scale Systems*, 12:91–107, 1987.
- [47] K. Sevcik. Priority scheduling disciplines in queueing network models of computer systems. In *IFIP Congress*, 1977.
- [48] H. Wang and K. C. Sevcik. Experiments with improved approximate mean value analysis algorithms. *Perform. Eval.*, 39(1-4):189–206, 2000.
- [49] W. Wang, G. Casale, and C. A. Sutton. A bayesian approach to parameter inference in queueing networks. *ACM Trans. Model. Comput. Simul.*, 27(1):2:1–2:26, 2016.
- [50] M. Woodside. *Tutorial Introduction to Layered Modeling of Software Performance*. Carleton University, February 2013.
- [51] J. Zahorjan, D. L. Eager, and H. M. Sweillam. Accuracy, speed, and convergence of approximate mean value analysis. *Perform. Eval.*, 8(4):255–270, 1988.
- [52] S. Zhou and M. Woodside. A multiserver approximation for cloud scaling analysis. In *Companion of the 2022 ACM/SPEC International Conference on Performance Engineering, ICPE ’22*, page 129–136, New York, NY, USA, 2022. Association for Computing Machinery.

Appendix A

Examples

The table below lists the Jupyter notebooks available under the `examples` folder.

Table A.1: Examples

Example	Problem
<code>example_cacheModel.1</code>	A small cache model with an open arrival process
<code>example_cacheModel.2</code>	A small cache model with a closed job population
<code>example_cacheModel.3</code>	A layered network with a caching layer
<code>example_cacheModel.4</code>	A layered network with a caching layer having a multi-level cache
<code>example_cacheModel.5</code>	A caching model with state-dependent output routing
<code>example_cdfRespT.1</code>	Station response time distribution in a single-class single-job closed network
<code>example_cdfRespT.2</code>	Station response time distribution in a multi-chain closed network
<code>example_cdfRespT.3</code>	Station response time distribution in a multi-chain open network
<code>example_cdfRespT.4</code>	Simulation-based station response time distribution analysis
<code>example_cdfRespT.5</code>	Station response time distribution under increasing job populations
<code>example_closedModel.1</code>	Solving a single-class exponential closed queueing network
<code>example_closedModel.2</code>	Solving a closed queueing network with a multi-class FCFS station
<code>example_closedModel.3</code>	Solving exactly a multi-chain product-form closed queueing network
<code>example_closedModel.4</code>	Local state space generation for a station in a closed network
<code>example_closedModel.5</code>	1-line exact MVA solution of a cyclic network of PS and INF stations
<code>example_closedModel.6</code>	Closed network with round robin scheduling
<code>example_closedModel.7</code>	Comparison of different scheduling policies that preserve the product-form solution
<code>example_forkJoin.1</code>	A simple single class open fork-join network
<code>example_forkJoin.2</code>	A multiclass open fork-join network
<code>example_forkJoin.3</code>	A closed model with nested forks and joins
<code>example_forkJoin.4</code>	An open model with a fork but without a join
<code>example_forkJoin.5</code>	A simple single class closed fork-join network
<code>example_forkJoin.6</code>	Two open fork-joins subsystems in tandem
<code>example_forkJoin.7</code>	Two-class fork-join with a class that switches into the other after the fork
<code>example_forkJoin.8</code>	Two fork-joins loops within the same chain
<code>example_initState.1</code>	Specifying an initial state and prior in a single class model.
<code>example_initState.2</code>	Specifying an initial state and prior in a multiclass model.

Continued on next page

Table A.1 – Examples. *Continued from previous page*

Example	Problem
example_initState_3	Specifying an initial state and prior in a model with class-switching.
example_layeredModel_1	Analyze a layered network specified in a LQNS XML file
example_layeredModel_2	Specifying and solving a basic layered network
example_loadDependent_1	Solving a single-class load-depedent closed model
example_loadDependent_2	Solving a two-node multiclass load-depedent closed model
example_loadDependent_3	Solving a three-node multiclass load-depedent closed model
example_loadDependent_4	Solving a load-independent closed model specified as a load-dependent model
example_misc_1	Use of performance indexes handles
example_misc_2	Update and refresh of service times
example_misc_3	Parameterization of a discriminatory processor sharing (DPS) station
example_misc_4	Automatic detection of solvers that cannot analyze the model
example_mixedModel_1	Solving a queueing network model with both closed and open classes
example_mixedModel_2	A difficult mixed model with sparse routing among multi-server nodes
example_openModel_1	Solving a queueing network model with open classes, scalar cutoff options
example_openModel_2	1-line solution of a tandem network of PS and INF stations
example_openModel_3	Solving a queueing network model with open classes, matrix cutoff options
example_openModel_4	Trace-driven simulation of an M/M/1 queue
example_openModel_5	A model illustrating the emulation of multiple sinks
example_openModel_6	A large multiclass example with PS and FCFS
example_prio_1	A multiclass example with PS, SIRO, FCFS, HOL priority
example_prio_2	A high-load multiclass example with PS, SIRO, FCFS, HOL priority
example_randomEnvironment_1	Solving a model in a 2-stage random environment with exponential rates
example_randomEnvironment_2	Solving a model in a 4-stage random environment with Coxian rates
example_randomEnvironment_3	Solving a model in a 3-stage random environment with Erlang rates
example_stateDependentRouting_1	A model with round-robin routing
example_stateDependentRouting_2	A model with round-robin routing after multi-class PH and MAP service
example_stateDependentRouting_2	A load-balancer modeled as a router
example_stateProbabilities_1	Computing marginal state probabilities for a node
example_stateProbabilities_2	Computing marginal state probabilities for a node under class-switching
example_stateProbabilities_3	Computing joint state probabilities for a system with two nodes under class-switching
example_stateProbabilities_4	Computing joint state probabilities under class-switching and with delay nodes
example_stateProbabilities_5	Computing probabilities under PS class-switching and with delay nodes
example_stateProbabilities_6	Computing probabilities under PS and FCFS class-switching and with delay nodes
example_stochPetriNet_1	JMT simulation of a simple stochastic Petri net model
example_stochPetriNet_2	JMT simulation of a complex stochastic Petri net model
example_scvEstimation_1	Demand estimation in a single class model using the UBR method
example_scvEstimation_2	Demand estimation in a multiclass model using the ERPS method
example_scvEstimation_3	Demand estimation in a multiclass model using the UBR method
example_scvEstimation_4	Demand estimation in a multiclass model using the UBO method