

---

# **Serverless Performance Simulator**

***Release 0.1.2***

**Nima Mahmoudi**

**Jun 24, 2020**



**CONTENTS:**

**1   Installation** **3**

**2   Serverless Simulator** **5**

**3   Temporal Simulator** **11**

    3.1   Serverless Temporal Simulator . . . . . 11

    3.2   Exponential Temporal Simulator . . . . . 11

**4   Simulated Process** **13**

**5   API Reference** **17**

    5.1   ServerlessSimulator . . . . . 17

    5.2   ServerlessTemporalSimulator . . . . . 21

    5.3   FunctionInstance . . . . . 22

    5.4   SimProcess . . . . . 23

    5.5   Utility . . . . . 25

**6   Indices and tables** **27**

**Python Module Index** **29**

**Index** **31**



This is a project done in [PACS Lab](#) aiming to develop a performance simulator for serverless computing platforms. Using this simulator, we can calculate Quality of Service (QoS) metrics like average response time, average probability of cold start, average running servers (directly reflecting average cost), histogram of different events, distribution of number of servers throughout time, and many other characteristics.

The developed performance model can be used to debug/improve analytical performance models, try new and improved management schema, or dig up a whole lot of properties of a common modern scale-per-request serverless platforms.

You can check out the source code in our [Github Repository](#).



## INSTALLATION

Install using pip:

```
pip install pacssim
```

Upgrading using pip:

```
pip install pacssim --upgrade
```

For installation in development mode:

```
git clone https://github.com/pacslab/serverless-performance-simulator
cd serverless-performance-simulator
pip install -e .
```

And in case you want to be able to execute the examples:

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



## SERVERLESS SIMULATOR

```
class pacssim.ServerlessSimulator.ServerlessSimulator (arrival_process=None,
                                                    warm_service_process=None,
                                                    cold_service_process=None,
                                                    expiration_threshold=600,
                                                    max_time=86400,      maxi-
                                                    mum_concurrency=1000,
                                                    **kwargs)
```

Bases: object

ServerlessSimulator is responsible for executing simulations of a sample serverless computing platform, mainly for the performance analysis and performance model evaluation purposes.

### Parameters

- **arrival\_process** (pacssim.SimProcess.SimProcess, optional) – The process used for generating inter-arrival samples, if absent, *arrival\_rate* should be passed to signal exponential distribution, by default None
- **warm\_service\_process** (pacssim.SimProcess.SimProcess, optional) – The process which will be used to calculate service times, if absent, *warm\_service\_rate* should be passed to signal exponential distribution, by default None
- **cold\_service\_process** (pacssim.SimProcess.SimProcess, optional) – The process which will be used to calculate service times, if absent, *cold\_service\_rate* should be passed to signal exponential distribution, by default None
- **expiration\_threshold** (float, optional) – The period of time after which the instance will be expired and the capacity release for use by others, by default 600
- **max\_time** (float, optional) – The maximum amount of time for which the simulation should continue, by default 24\*60\*60 (24 hours)
- **maximum\_concurrency** (int, optional) – The maximum number of concurrently executing function instances allowed on the system This will be used to determine when a rejection of request should happen due to lack of capacity, by default 1000

### Raises

- **Exception** – Raises if neither *arrival\_process* nor *arrival\_rate* are present
- **Exception** – Raises if neither *warm\_service\_process* nor *warm\_service\_rate* are present
- **Exception** – Raises if neither *cold\_service\_process* nor *cold\_service\_rate* are present
- **ValueError** – Raises if *warm\_service\_rate* is smaller than *cold\_service\_rate*

**analyze\_custom\_states** (*hist\_states*, *skip\_init\_time=None*, *skip\_init\_index=None*)

Analyses a custom states list and calculates the amount of time spent in each state each time we entered that state, and the times at which transitions have happened.

**Parameters**

- **hist\_states** (*list[object]*) – The states calculated, should have the same dimensions as the *hist\_\** arrays.
- **skip\_init\_time** (*float*, *optional*) – The amount of time skipped in the beginning, by default None
- **skip\_init\_index** (*int*, *optional*) – The number of indices skipped in the beginning, by default None

**Returns** (*residence\_times*, *transition\_times*) where *residence\_times* is an array of the amount of times we spent in each state, and *transition\_times* are the moments of time at which each transition has occurred

**Return type** *list[float]*, *list[float]*

**calculate\_time\_average** (*values*, *skip\_init\_time=None*, *skip\_init\_index=None*)

*calculate\_time\_average* calculates the time-averaged of the values passed in with optional skipping a specific number of time steps (*skip\_init\_index*) and a specific amount of time (*skip\_init\_time*).

**Parameters**

- **values** (*list*) – A list of values with the same dimensions as history array (number of transitions)
- **skip\_init\_time** (*Float*, *optional*) – Amount of time skipped in the beginning to let the transient part of the solution pass, by default None
- **skip\_init\_index** (*[type]*, *optional*) – Number of steps skipped in the beginning to let the transient behaviour of system pass, by default None

**Returns** returns (*uniq\_vals*, *val\_times*) where *uniq\_vals* is the unique values inside the values list and *val\_times* is the portion of the time that is spent in that value.

**Return type** (*list*, *list*)

**calculate\_time\_lengths** ()

Calculate the time length for each step between two event transitions. Records the values in *self.time\_lengths*.

**cold\_start\_arrival** (*t*)

Goes through the process necessary for a cold start arrival which includes generation of a new function instance in the *COLD* state and adding it to the cluster.

**Parameters** *t* (*float*) – The time at which the arrival has happened. This is used to record the creation time for the server and schedule the expiration of the instance if necessary.

**generate\_trace** (*debug\_print=False*, *progress=False*)

Generate a sample trace.

**Parameters**

- **debug\_print** (*bool*, *optional*) – If True, will print each transition occurring during the simulation, by default False
- **progress** (*bool*, *optional*) – Whether or not the progress should be outputted using the *tqdm* library, by default False

**Raises Exception** – Raises if FunctionInstance enters an unknown state (other than *IDLE* for idle or *TERM* for terminated) after making an internal transition

**get\_average\_lifespan()**

Get the average lifespan of each instance, calculated by the amount of time from creation of instance, until its expiration.

**Returns** The average lifespan

**Return type** float

**get\_average\_residence\_times**(*hist\_states*, *skip\_init\_time=None*, *skip\_init\_index=None*)

Get the average residence time for each state in custom state encoding.

**Parameters**

- **hist\_states** (*list[object]*) – The states calculated, should have the same dimensions as the *hist\_\** arrays.
- **skip\_init\_time** (*float*, *optional*) – The amount of time skipped in the beginning, by default None
- **skip\_init\_index** (*int*, *optional*) – The number of indices skipped in the beginning, by default None

**Returns** The average residence time for each state, averaged over the times we transitioned into that state

**Return type** float

**get\_average\_server\_count()**

Get the time-average server count.

**Returns** Average server count

**Return type** float

**get\_average\_server\_idle\_count()**

Get the time-averaged idle server count.

**Returns** Average idle server count

**Return type** float

**get\_average\_server\_running\_count()**

Get the time-averaged running server count.

**Returns** Average running server count

**Return type** float

**get\_cold\_start\_prob()**

Get the probability of cold start for the simulated trace.

**Returns** The probability of cold start calculated by dividing the number of cold start requests, over all requests

**Return type** float

**get\_index\_after\_time**(*t*)

Get the first historical array index (for all arrays storing historical events) that is after the time *t*.

**Parameters** *t* (*float*) – The time in the beginning we want to skip

**Returns** The calculated index in *self.hist\_times*

**Return type** int

**get\_request\_custom\_states** (*hist\_states*, *skip\_init\_time=None*, *skip\_init\_index=None*)

Get request statistics for an array of custom states.

**Parameters**

- **hist\_states** (*list[object]*) – An array of custom states calculated by the user for which the statistics should be calculated, should be the same size as *hist\_\** objects, these values will be used as the keys for the returned dataframe.
- **skip\_init\_time** (*float, optional*) – The amount of time skipped in the beginning, by default None
- **skip\_init\_index** (*int, optional*) – The number of indices that should be skipped in the beginning to calculate steady-state results, by default None

**Returns** A pandas dataframe including different statistics like *p\_cold* (probability of cold start)

**Return type** pandas.DataFrame

**get\_result\_dict** ()

Get the results of the simulation as a dict, which can easily be integrated into web services.

**Returns** A dictionary of different characteristics.

**Return type** dict

**get\_skip\_init** (*skip\_init\_time=None*, *skip\_init\_index=None*)

Get the minimum index which satisfies both the time and index count we want to skip in the beginning of the simulation, which is used to reduce the transient effect for calculating the steady-state values.

**Parameters**

- **skip\_init\_time** (*float, optional*) – The amount of time skipped in the beginning, by default None
- **skip\_init\_index** (*[type], optional*) – The number of indices we want to skip in the historical events, by default None

**Returns** The number of indices after which both index and time requirements are satisfied

**Return type** int

**get\_trace\_end** ()

Get the time at which the trace (one iteration of the simulation) has ended. This mainly due to the fact that we keep on simulating until the trace time goes beyond *max\_time*, but the time is incremented until the next event.

**Returns** The time at which the trace has ended

**Return type** float

**has\_server** ()

Returns True if there are still instances (servers) in the simulated platform, False otherwise.

**Returns** Whether or not the platform has instances (servers)

**Return type** bool

**static print\_time\_average** (*vals*, *probs*, *column\_width=15*)

Print the time average of states.

**Parameters**

- **vals** (*list[object]*) – The values for which the time average is to be printed
- **probs** (*list[float]*) – The probability of each of the members of the values array

- **column\_width** (*int*, *optional*) – The width of the printed result for *vals*, by default 15

**print\_trace\_results** ()

Print a brief summary of the results of the trace.

**req** ()

Generate a request inter-arrival from *self.arrival\_process*

**Returns** The generated inter-arrival sample

**Return type** float

**reset\_trace** ()

resets all the historical data to prepare the class for a new simulation

**schedule\_warm\_instance** (*t*)

Goes through a process to determine which warm instance should process the incoming request.

**Parameters** *t* (*float*) – The time at which the scheduling is happening

**Returns** The function instances that the scheduler has selected for the incoming request.

**Return type** *pacssim.FunctionInstance.FunctionInstance*

**trace\_condition** (*t*)

The condition for resulting the trace, we continue the simulation until this function returns false.

**Parameters** *t* (*float*) – current time in the simulation since the start of simulation

**Returns** True if we should continue the simulation, false otherwise

**Return type** bool

**warm\_start\_arrival** (*t*)

Goes through the process necessary for a warm start arrival which includes selecting a warm instance for processing and recording the request information.

**Parameters** *t* (*float*) – The time at which the arrival has happened. This is used to record the creation time for the server and schedule the expiration of the instance if necessary.



## TEMPORAL SIMULATOR

In this family of classes, we want to extract temporal characteristics using execution of simulations. We can extract average estimates by average over several executions of the simulation (sample average). All of these classes extend the functionality provided by *ServerlessSimulator*, you can use the same arguments and call the same methods, with some extended functionality provided below.

### 3.1 Serverless Temporal Simulator

```
class pacssim.ServerlessTemporalSimulator.ServerlessTemporalSimulator (running_function_instances,  
                                                                    idle_function_instances,  
                                                                    *args,  
                                                                    **kwargs)
```

Bases: *pacssim.ServerlessSimulator.ServerlessSimulator*

ServerlessTemporalSimulator extends ServerlessSimulator to enable extraction of temporal characteristics. Also gets all of the arguments accepted by *ServerlessSimulator*

#### Parameters

- **running\_function\_instances** (*list [FunctionInstance]*) – A list containing the running function instances
- **idle\_function\_instances** (*list [FunctionInstance]*) – A list containing the idle function instances

### 3.2 Exponential Temporal Simulator

The exponential temporal simulator assume exponential inter-event distribution for both arrival and departure from each function instance.

```
class pacssim.ServerlessTemporalSimulator.ExponentialServerlessTemporalSimulator (running_function_instances,  
                                                                    idle_function_instances,  
                                                                    *args,  
                                                                    **kwargs)
```

Bases: *pacssim.ServerlessTemporalSimulator.ServerlessTemporalSimulator*

ExponentialServerlessTemporalSimulator is a simulator assuming exponential distribution for processing times which means each process is state-less and we can generate a service time and use that from now on. This class extends ServerlessTemporalSimulator which has functionality for other processes as well.

#### Parameters

- **running\_function\_instance\_count** (*integer*) – running\_function\_instance\_count is the number of instances currently processing a request
- **idle\_function\_instance\_next\_terminations** (*list[float]*) – idle\_function\_instance\_next\_terminations is an array of next termination scheduled for idle functions if they receive no new requests.

## SIMULATED PROCESS

```
class pacssim.SimProcess.ConstSimProcess (rate)
```

Bases: *pacssim.SimProcess.SimProcess*

ConstSimProcess extends the functionality of *SimProcess* for constant processes, meaning this is a deterministic process and fires exactly every  $1/\text{rate}$  seconds. This class does not implement the *pdf* and *cdf* functions.

**rate** [float] The rate at which the process should fire off

**generate\_trace** ()

generate\_trace function is supposed to be replaced with the override function of each of the child classes.

**NotImplementedError** By default, this function raises NotImplementedError unless overridden by a child class.

```
class pacssim.SimProcess.ExpSimProcess (rate)
```

Bases: *pacssim.SimProcess.SimProcess*

ExpSimProcess extends the functionality of *SimProcess* for exponentially distributed processes. This class also implements the *pdf* and *cdf* functions which can be used for visualization purposes.

**rate** [float] The rate at which the process should fire off

**cdf** (*x*)

cdf function is called for visualization for classes with *self.has\_cdf = True*.

**Parameters** *x* (*float*) – The time for which the cdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises NotImplementedError unless overridden by a child class.

**generate\_trace** ()

generate\_trace function is supposed to be replaced with the override function of each of the child classes.

**NotImplementedError** By default, this function raises NotImplementedError unless overridden by a child class.

**pdf** (*x*)

pdf function is called for visualization for classes with *self.has\_pdf = True*.

**Parameters** *x* (*float*) – The time for which the pdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises NotImplementedError unless overridden by a child class.

```
class pacssim.SimProcess.GaussianSimProcess (rate, std)
```

Bases: *pacssim.SimProcess.SimProcess*

GaussianSimProcess extends the functionality of *SimProcess* for gaussian processes. This class also implements the *pdf* and *cdf* functions which can be used for visualization purposes.

**rate** [float] The rate at which the process should fire off

**std** [float] The standard deviation of the simulated process

**cdf** (*x*)

cdf function is called for visualization for classes with *self.has\_cdf = True*.

**Parameters** *x* (*float*) – The time for which the cdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises **NotImplementedError** unless overridden by a child class.

**generate\_trace** ()

generate\_trace function is supposed to be replaced with the override function of each of the child classes.

**NotImplementedError** By default, this function raises **NotImplementedError** unless overridden by a child class.

**pdf** (*x*)

pdf function is called for visualization for classes with *self.has\_pdf = True*.

**Parameters** *x* (*float*) – The time for which the pdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises **NotImplementedError** unless overridden by a child class.

**class** `pacssim.SimProcess.SimProcess`

Bases: `object`

`SimProcess` gives us a single interface to simulate different processes. This will later on be used to simulated different processes and compare them against a custom analytical model. In the child class, after performing `super().__init__()`, properties *self.has\_pdf* and *self.has\_cdf* by default value of *False* will be created. In case your class has the proposed PDF and CDF functions available, you need to override these values in order for your model PDF to show up in the output plot.

**cdf** (*x*)

cdf function is called for visualization for classes with *self.has\_cdf = True*.

**Parameters** *x* (*float*) – The time for which the cdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises **NotImplementedError** unless overridden by a child class.

**generate\_trace** ()

generate\_trace function is supposed to be replaced with the override function of each of the child classes.

**NotImplementedError** By default, this function raises **NotImplementedError** unless overridden by a child class.

**pdf** (*x*)

pdf function is called for visualization for classes with *self.has\_pdf = True*.

**Parameters** *x* (*float*) – The time for which the pdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises **NotImplementedError** unless overridden by a child class.

**visualize** (*num\_traces=10000*, *num\_bins=100*)

visualize function visualizes the PDF and CDF of the simulated process by generating traces from your function using `generate_trace()` and converting the resulting histogram values (event counts) to densities to be comparable with PDF and CDF functions calculated analytically.

**num\_traces** [int, optional] Number of traces we want to generate for calculating the histogram, by default 10000

**num\_bins** [int, optional] Number of bins for the histogram which created the density probabilities, by default 100



## API REFERENCE

You can also checkout [Module Index](#) for a list of all modules implemented in this package.

### 5.1 ServerlessSimulator

```
class pacssim.ServerlessSimulator.ServerlessSimulator (arrival_process=None,  
                                                    warm_service_process=None,  
                                                    cold_service_process=None,  
                                                    expiration_threshold=600,  
                                                    max_time=86400,      maxi-  
                                                    mum_concurrency=1000,  
                                                    **kwargs)
```

Bases: `object`

ServerlessSimulator is responsible for executing simulations of a sample serverless computing platform, mainly for the performance analysis and performance model evaluation purposes.

#### Parameters

- **arrival\_process** (`pacssim.SimProcess.SimProcess`, *optional*) – The process used for generating inter-arrival samples, if absent, *arrival\_rate* should be passed to signal exponential distribution, by default `None`
- **warm\_service\_process** (`pacssim.SimProcess.SimProcess`, *optional*) – The process which will be used to calculate service times, if absent, *warm\_service\_rate* should be passed to signal exponential distribution, by default `None`
- **cold\_service\_process** (`pacssim.SimProcess.SimProcess`, *optional*) – The process which will be used to calculate service times, if absent, *cold\_service\_rate* should be passed to signal exponential distribution, by default `None`
- **expiration\_threshold** (*float*, *optional*) – The period of time after which the instance will be expired and the capacity release for use by others, by default `600`
- **max\_time** (*float*, *optional*) – The maximum amount of time for which the simulation should continue, by default `24*60*60` (24 hours)
- **maximum\_concurrency** (*int*, *optional*) – The maximum number of concurrently executing function instances allowed on the system This will be used to determine when a rejection of request should happen due to lack of capacity, by default `1000`

#### Raises

- **Exception** – Raises if neither *arrival\_process* nor *arrival\_rate* are present
- **Exception** – Raises if neither *warm\_service\_process* nor *warm\_service\_rate* are present

- **Exception** – Raises if neither `cold_service_process` nor `cold_service_rate` are present
- **ValueError** – Raises if `warm_service_rate` is smaller than `cold_service_rate`

**analyze\_custom\_states** (*hist\_states*, *skip\_init\_time=None*, *skip\_init\_index=None*)

Analyses a custom states list and calculates the amount of time spent in each state each time we entered that state, and the times at which transitions have happened.

**Parameters**

- **hist\_states** (*list[object]*) – The states calculated, should have the same dimensions as the *hist\_\** arrays.
- **skip\_init\_time** (*float, optional*) – The amount of time skipped in the beginning, by default `None`
- **skip\_init\_index** (*int, optional*) – The number of indices skipped in the beginning, by default `None`

**Returns** (*residence\_times*, *transition\_times*) where *residence\_times* is an array of the amount of times we spent in each state, and *transition\_times* are the moments of time at which each transition has occurred

**Return type** *list[float], list[float]*

**calculate\_time\_average** (*values*, *skip\_init\_time=None*, *skip\_init\_index=None*)

`calculate_time_average` calculates the time-averaged of the values passed in with optional skipping a specific number of time steps (*skip\_init\_index*) and a specific amount of time (*skip\_init\_time*).

**Parameters**

- **values** (*list*) – A list of values with the same dimensions as history array (number of transitions)
- **skip\_init\_time** (*Float, optional*) – Amount of time skipped in the beginning to let the transient part of the solution pass, by default `None`
- **skip\_init\_index** (*[type], optional*) – Number of steps skipped in the beginning to let the transient behaviour of system pass, by default `None`

**Returns** returns (*uniq\_vals*, *val\_times*) where *uniq\_vals* is the unique values inside the values list and *val\_times* is the portion of the time that is spent in that value.

**Return type** (*list*, *list*)

**calculate\_time\_lengths** ()

Calculate the time length for each step between two event transitions. Records the values in *self.time\_lengths*.

**cold\_start\_arrival** (*t*)

Goes through the process necessary for a cold start arrival which includes generation of a new function instance in the *COLD* state and adding it to the cluster.

**Parameters** *t* (*float*) – The time at which the arrival has happened. This is used to record the creation time for the server and schedule the expiration of the instance if necessary.

**generate\_trace** (*debug\_print=False*, *progress=False*)

Generate a sample trace.

**Parameters**

- **debug\_print** (*bool, optional*) – If `True`, will print each transition occurring during the simulation, by default `False`

- **progress** (*bool, optional*) – Whether or not the progress should be outputted using the *tqdm* library, by default False

**Raises Exception** – Raises if FunctionInstance enters an unknown state (other than *IDLE* for idle or *TERM* for terminated) after making an internal transition

**get\_average\_lifespan()**

Get the average lifespan of each instance, calculated by the amount of time from creation of instance, until its expiration.

**Returns** The average lifespan

**Return type** float

**get\_average\_residence\_times** (*hist\_states, skip\_init\_time=None, skip\_init\_index=None*)

Get the average residence time for each state in custom state encoding.

**Parameters**

- **hist\_states** (*list[object]*) – The states calculated, should have the same dimensions as the *hist\_\** arrays.
- **skip\_init\_time** (*float, optional*) – The amount of time skipped in the beginning, by default None
- **skip\_init\_index** (*int, optional*) – The number of indices skipped in the beginning, by default None

**Returns** The average residence time for each state, averaged over the times we transitioned into that state

**Return type** float

**get\_average\_server\_count()**

Get the time-average server count.

**Returns** Average server count

**Return type** float

**get\_average\_server\_idle\_count()**

Get the time-averaged idle server count.

**Returns** Average idle server count

**Return type** float

**get\_average\_server\_running\_count()**

Get the time-averaged running server count.

**Returns** Average running server count

**Return type** float

**get\_cold\_start\_prob()**

Get the probability of cold start for the simulated trace.

**Returns** The probability of cold start calculated by dividing the number of cold start requests, over all requests

**Return type** float

**get\_index\_after\_time** (*t*)

Get the first historical array index (for all arrays storing historical events) that is after the time *t*.

**Parameters** *t* (*float*) – The time in the beginning we want to skip

**Returns** The calculated index in *self.hist\_times*

**Return type** int

**get\_request\_custom\_states** (*hist\_states*, *skip\_init\_time=None*, *skip\_init\_index=None*)

Get request statistics for an array of custom states.

**Parameters**

- **hist\_states** (*list[object]*) – An array of custom states calculated by the user for which the statistics should be calculated, should be the same size as *hist\_\** objects, these values will be used as the keys for the returned dataframe.
- **skip\_init\_time** (*float, optional*) – The amount of time skipped in the beginning, by default None
- **skip\_init\_index** (*int, optional*) – The number of indices that should be skipped in the beginning to calculate steady-state results, by default None

**Returns** A pandas dataframe including different statistics like *p\_cold* (probability of cold start)

**Return type** pandas.DataFrame

**get\_result\_dict** ()

Get the results of the simulation as a dict, which can easily be integrated into web services.

**Returns** A dictionary of different characteristics.

**Return type** dict

**get\_skip\_init** (*skip\_init\_time=None*, *skip\_init\_index=None*)

Get the minimum index which satisfies both the time and index count we want to skip in the beginning of the simulation, which is used to reduce the transient effect for calculating the steady-state values.

**Parameters**

- **skip\_init\_time** (*float, optional*) – The amount of time skipped in the beginning, by default None
- **skip\_init\_index** (*[type], optional*) – The number of indices we want to skip in the historical events, by default None

**Returns** The number of indices after which both index and time requirements are satisfied

**Return type** int

**get\_trace\_end** ()

Get the time at which the trace (one iteration of the simulation) has ended. This mainly due to the fact that we keep on simulating until the trace time goes beyond *max\_time*, but the time is incremented until the next event.

**Returns** The time at which the trace has ended

**Return type** float

**has\_server** ()

Returns True if there are still instances (servers) in the simulated platform, False otherwise.

**Returns** Whether or not the platform has instances (servers)

**Return type** bool

**static print\_time\_average** (*vals*, *probs*, *column\_width=15*)

Print the time average of states.

**Parameters**

- **vals** (*list[object]*) – The values for which the time average is to be printed
- **probs** (*list[float]*) – The probability of each of the members of the values array
- **column\_width** (*int, optional*) – The width of the printed result for *vals*, by default 15

**print\_trace\_results()**

Print a brief summary of the results of the trace.

**req()**

Generate a request inter-arrival from *self.arrival\_process*

**Returns** The generated inter-arrival sample

**Return type** float

**reset\_trace()**

resets all the historical data to prepare the class for a new simulation

**schedule\_warm\_instance(t)**

Goes through a process to determine which warm instance should process the incoming request.

**Parameters** *t* (*float*) – The time at which the scheduling is happening

**Returns** The function instances that the scheduler has selected for the incoming request.

**Return type** *pacssim.FunctionInstance.FunctionInstance*

**trace\_condition(t)**

The condition for resulting the trace, we continue the simulation until this function returns false.

**Parameters** *t* (*float*) – current time in the simulation since the start of simulation

**Returns** True if we should continue the simulation, false otherwise

**Return type** bool

**warm\_start\_arrival(t)**

Goes through the process necessary for a warm start arrival which includes selecting a warm instance for processing and recording the request information.

**Parameters** *t* (*float*) – The time at which the arrival has happened. This is used to record the creation time for the server and schedule the expiration of the instance if necessary.

## 5.2 ServerlessTemporalSimulator

**class** *pacssim.ServerlessTemporalSimulator.ExponentialServerlessTemporalSimulator* (*running\_function*, *idle\_function*, *\*args*, *\*\*kwargs*)

Bases: *pacssim.ServerlessTemporalSimulator.ServerlessTemporalSimulator*

ExponentialServerlessTemporalSimulator is a simulator assuming exponential distribution for processing times which means each process is state-less and we can generate a service time and use that from now on. This class extends ServerlessTemporalSimulator which has functionality for other processes as well.

**Parameters**

- **running\_function\_instance\_count** (*integer*) – running\_function\_instance\_count is the number of instances currently processing a request

- **idle\_function\_instance\_next\_terminations** (*list[float]*) – idle\_function\_instance\_next\_terminations is an array of next termination scheduled for idle functions if they receive no new requests.

```
class pacssim.ServerlessTemporalSimulator.ServerlessTemporalSimulator(running_function_instances,  
                                                                    idle_function_instances,  
                                                                    *args,  
                                                                    **kwargs)
```

Bases: *pacssim.ServerlessSimulator.ServerlessSimulator*

ServerlessTemporalSimulator extends ServerlessSimulator to enable extraction of temporal characteristics. Also gets all of the arguments accepted by *ServerlessSimulator*

#### Parameters

- **running\_function\_instances** (*list[FunctionInstance]*) – A list containing the running function instances
- **idle\_function\_instances** (*list[FunctionInstance]*) – A list containing the idle function instances

## 5.3 FunctionInstance

```
class pacssim.FunctionInstance.FunctionInstance(t,                                cold_service_process,  
                                                warm_service_process,          expiration_threshold)
```

Bases: *object*

FunctionInstance aims to simulate the behaviour of a function instance in a serverless platform, with all the internal transitions necessary.

#### Parameters

- **t** (*float*) – The time at which the instance is being created
- **cold\_service\_process** (*pacssim.SimProcess.SimProcess*) – The process used to sample cold start response times
- **warm\_service\_process** (*pacssim.SimProcess.SimProcess*) – The process used to sample warm start response times
- **expiration\_threshold** (*float*) – The amount of time it takes for an instance to get expired and the resources consumed by it released after processing the last request

**arrival\_transition** (*t*)

Make an arrival transition, which causes the instance to go from IDLE to WARM

**Parameters** **t** (*float*) – The time at which the transition has occurred, this also updates the next termination.

**Raises** **Exception** – Raises if currently process a request by being in *COLD* or *WARM* states

**get\_life\_span** ()

**get\_next\_departure** (*t*)

Get the time until the next departure

**Parameters** **t** (*float*) – Current time

**Returns** Amount of time until the next departure

**Return type** float

**Raises Exception** – Raises if called after the departure

**get\_next\_termination** (*t*)

Get the time until the next termination

**Parameters** *t* (*float*) – Current time

**Returns** Amount of time until the next termination

**Return type** *float*

**Raises Exception** – Raises if called after the termination

**get\_next\_transition\_time** (*t=0*)

Get how long until the next transition.

**Parameters** *t* (*float, optional*) – The current time, by default 0

**Returns** The seconds remaining until the next transition

**Return type** *float*

**get\_state** ()

**is\_idle** ()

Whether or not the instance is currently idle, and thus can accept new requests.

**Returns** True if idle, false otherwise

**Return type** *bool*

**make\_transition** ()

Make the next internal transition, either transition into *IDLE* of already processing a request, or *TERM* if scheduled termination has arrived.

**Returns** The state after making the internal transition

**Return type** *str*

**Raises Exception** – Raises if already in *TERM* state, since no other internal transitions are possible

**update\_next\_termination** ()

Update the next scheduled termination if no other requests are made to the instance.

## 5.4 SimProcess

**class** `pacssim.SimProcess.ConstSimProcess` (*rate*)

Bases: `pacssim.SimProcess.SimProcess`

`ConstSimProcess` extends the functionality of `SimProcess` for constant processes, meaning this is a deterministic process and fires exactly every  $1/\text{rate}$  seconds. This class does not implement the *pdf* and *cdf* functions.

**rate** [*float*] The rate at which the process should fire off

**generate\_trace** ()

`generate_trace` function is supposed to be replaced with the override function of each of the child classes.

**NotImplementedError** By default, this function raises `NotImplementedError` unless overridden by a child class.

```
class pacssim.SimProcess.ExpSimProcess (rate)
```

```
    Bases: pacssim.SimProcess.SimProcess
```

ExpSimProcess extends the functionality of *SimProcess* for exponentially distributed processes. This class also implements the *pdf* and *cdf* functions which can be used for visualization purposes.

**rate** [float] The rate at which the process should fire off

**cdf** (*x*)

cdf function is called for visualization for classes with *self.has\_cdf = True*.

**Parameters** *x* (*float*) – The time for which the cdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises **NotImplementedError** unless overridden by a child class.

**generate\_trace** ()

generate\_trace function is supposed to be replaced with the override function of each of the child classes.

**NotImplementedError** By default, this function raises **NotImplementedError** unless overridden by a child class.

**pdf** (*x*)

pdf function is called for visualization for classes with *self.has\_pdf = True*.

**Parameters** *x* (*float*) – The time for which the pdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises **NotImplementedError** unless overridden by a child class.

```
class pacssim.SimProcess.GaussianSimProcess (rate, std)
```

```
    Bases: pacssim.SimProcess.SimProcess
```

GaussianSimProcess extends the functionality of *SimProcess* for gaussian processes. This class also implements the *pdf* and *cdf* functions which can be used for visualization purposes.

**rate** [float] The rate at which the process should fire off

**std** [float] The standard deviation of the simulated process

**cdf** (*x*)

cdf function is called for visualization for classes with *self.has\_cdf = True*.

**Parameters** *x* (*float*) – The time for which the cdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises **NotImplementedError** unless overridden by a child class.

**generate\_trace** ()

generate\_trace function is supposed to be replaced with the override function of each of the child classes.

**NotImplementedError** By default, this function raises **NotImplementedError** unless overridden by a child class.

**pdf** (*x*)

pdf function is called for visualization for classes with *self.has\_pdf = True*.

**Parameters** *x* (*float*) – The time for which the pdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises **NotImplementedError** unless overridden by a child class.

```
class pacssim.SimProcess.SimProcess
```

```
    Bases: object
```

SimProcess gives us a single interface to simulate different processes. This will later on be used to simulated different processes and compare them against a custom analytical model. In the child class, after performing `super().__init__()`, properties `self.has_pdf` and `self.has_cdf` by default value of `False` will be created. In case your class has the proposed PDF and CDF functions available, you need to override these values in order for your model PDF to show up in the output plot.

**cdf** (*x*)

cdf function is called for visualization for classes with `self.has_cdf = True`.

**Parameters** *x* (*float*) – The time for which the cdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises `NotImplementedError` unless overridden by a child class.

**generate\_trace** ()

generate\_trace function is supposed to be replaced with the override function of each of the child classes.

**NotImplementedError** By default, this function raises `NotImplementedError` unless overridden by a child class.

**pdf** (*x*)

pdf function is called for visualization for classes with `self.has_pdf = True`.

**Parameters** *x* (*float*) – The time for which the pdf value (density) should be returned

**Raises** **NotImplementedError** – By default, this function raises `NotImplementedError` unless overridden by a child class.

**visualize** (*num\_traces=10000*, *num\_bins=100*)

visualize function visualizes the PDF and CDF of the simulated process by generating traces from your function using `generate_trace()` and converting the resulting histogram values (event counts) to densities to be comparable with PDF and CDF functions calculated analytically.

**num\_traces** [int, optional] Number of traces we want to generate for calculating the histogram, by default 10000

**num\_bins** [int, optional] Number of bins for the histogram which created the density probabilities, by default 100

## 5.5 Utility

`pacssim.Utility.convert_hist_pdf(_values, num_bins)`

convert\_hist\_pdf converts the histogram resulting from `_values` and `num_bins` to a density plot by dividing the probability of falling into a bin by the bin size, converting the values to density. The resulting values could be plotted and compared with the analytical pdf and cdf functions.

**\_values** [list[float]] A list of values that we want to analyze and calculate the histogram for

**num\_bins** [int] Number of bins used for generating the histogram

**list[float], list[float], list[float]** base, values, cumulative are returned which are the histogram bases, density values, and cumulative densities which can be compared with the analytical cdf function



## INDICES AND TABLES

- `genindex`
- `modindex`
- `search`



## PYTHON MODULE INDEX

### p

- `pacssim.FunctionInstance`, [22](#)
- `pacssim.ServerlessSimulator`, [17](#)
- `pacssim.ServerlessTemporalSimulator`, [21](#)
- `pacssim.SimProcess`, [13](#)
- `pacssim.Utility`, [25](#)



## A

`analyze_custom_states()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 5, 18

`arrival_transition()` (*pacssim.FunctionInstance.FunctionInstance method*), 22

## C

`calculate_time_average()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 6, 18

`calculate_time_lengths()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 6, 18

`cdf()` (*pacssim.SimProcess.ExpSimProcess method*), 13, 24

`cdf()` (*pacssim.SimProcess.GaussianSimProcess method*), 14, 24

`cdf()` (*pacssim.SimProcess.SimProcess method*), 14, 25

`cold_start_arrival()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 6, 18

`ConstSimProcess` (class in *pacssim.SimProcess*), 13, 23

`convert_hist_pdf()` (in module *pacssim.Utility*), 25

## E

`ExponentialServerlessTemporalSimulator` (class in *pacssim.ServerlessTemporalSimulator*), 11, 21

`ExpSimProcess` (class in *pacssim.SimProcess*), 13, 23

## F

`FunctionInstance` (class in *pacssim.FunctionInstance*), 22

## G

`GaussianSimProcess` (class in *pacssim.SimProcess*), 13, 24

`generate_trace()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 6, 18

`generate_trace()` (*pacssim.SimProcess.ConstSimProcess method*), 13, 23

`generate_trace()` (*pacssim.SimProcess.ExpSimProcess method*), 13, 24

`generate_trace()` (*pacssim.SimProcess.GaussianSimProcess method*), 14, 24

`generate_trace()` (*pacssim.SimProcess.SimProcess method*), 14, 25

`get_average_lifespan()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 7, 19

`get_average_residence_times()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 7, 19

`get_average_server_count()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 7, 19

`get_average_server_idle_count()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 7, 19

`get_average_server_running_count()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 7, 19

`get_cold_start_prob()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 7, 19

`get_index_after_time()` (*pacssim.ServerlessSimulator.ServerlessSimulator method*), 7, 19

`get_life_span()` (*pacssim.FunctionInstance.FunctionInstance method*), 22

`get_next_departure()` (*pacssim.FunctionInstance.FunctionInstance method*), 22

```

get_next_termination() (pac- pdf() (pacssim.SimProcess.ExpSimProcess method),
                        ssim.FunctionInstance.FunctionInstance 13, 24
                        method), 23
get_next_transition_time() (pac- pdf() (pacssim.SimProcess.GaussianSimProcess
                        ssim.FunctionInstance.FunctionInstance method), 14, 24
                        method), 23
get_request_custom_states() (pac- pdf() (pacssim.SimProcess.SimProcess method), 14, 25
                        ssim.ServerlessSimulator.ServerlessSimulator
                        method), 7, 20
                        print_time_average() (pac-
get_result_dict() (pac- ssim.ServerlessSimulator.ServerlessSimulator
                        ssim.ServerlessSimulator.ServerlessSimulator static method), 8, 20
                        method), 8, 20
                        print_trace_results() (pac-
get_skip_init() (pac- ssim.ServerlessSimulator.ServerlessSimulator
                        ssim.ServerlessSimulator.ServerlessSimulator method), 9, 21
                        method), 8, 20
                        R
get_state() (pacssim.FunctionInstance.FunctionInstance req() (pacssim.ServerlessSimulator.ServerlessSimulator
                        method), 23
                        reset_trace() (pac-
get_trace_end() (pac- ssim.ServerlessSimulator.ServerlessSimulator
                        ssim.ServerlessSimulator.ServerlessSimulator method), 9, 21
                        method), 8, 20
                        S
H
schedule_warm_instance() (pac-
has_server() (pac- ssim.ServerlessSimulator.ServerlessSimulator
                        ssim.ServerlessSimulator.ServerlessSimulator method), 9, 21
                        method), 8, 20
                        ServerlessSimulator (class in pac-
I
                        ssim.ServerlessSimulator), 5, 17
is_idle() (pacssim.FunctionInstance.FunctionInstance ServerlessTemporalSimulator (class in pac-
                        method), 23
                        ssim.ServerlessTemporalSimulator), 11, 22
M
                        SimProcess (class in pacssim.SimProcess), 14, 24
make_transition() (pac-
                        T
                        trace_condition() (pac-
module
                        ssim.ServerlessSimulator.ServerlessSimulator
                        pacssim.FunctionInstance, 22
                        method), 9, 21
                        pacssim.ServerlessSimulator, 17
                        U
                        update_next_termination() (pac-
                        pacssim.ServerlessTemporalSimulator,
                        ssim.FunctionInstance.FunctionInstance
                        21
                        method), 23
                        pacssim.SimProcess, 13, 23
                        V
                        visualize() (pacssim.SimProcess.SimProcess
                        pacssim.Utility, 25
                        method), 14, 25
                        W
P
warm_start_arrival() (pac-
pacssim.FunctionInstance
                        ssim.ServerlessSimulator.ServerlessSimulator
                        module, 22
                        method), 9, 21
pacssim.ServerlessSimulator
                        module, 17
pacssim.ServerlessTemporalSimulator
                        module, 21
pacssim.SimProcess
                        module, 13, 23
pacssim.Utility
                        module, 25

```