
Neurokernel Documentation

Release 0.1.0

Lev Givon

Apr 11, 2018

Contents

1	Contents	3
1.1	Introduction	3
1.2	Installation	3
1.3	Reference	6
1.4	Authors & Acknowledgements	32
1.5	Licenses	32
2	Index	35

[Project Website](#) | [GitHub Repository](#) | [Mailing List](#) | [Forum](#)

1.1 Introduction

Neurokernel is an open software platform written in Python for emulation of the brain of the fruit fly (*Drosophila melanogaster*) on multiple Graphics Processing Units (GPUs). It provides a programming model based upon the organization of the fly's brain into fewer than 50 modular subdivisions called *local processing units* (LPUs) that are each characterized by unique populations of local neurons¹. Using Neurokernel's API, researchers can develop models of individual LPUs and combine them with other independently developed LPU models to collaboratively construct models of entire subsystems of the fly brain. Neurokernel's support for LPU model integration also enables exploration of brain functions that cannot be exclusively attributed to individual LPUs or brain subsystems.

Examples of Neurokernel's use are available on the [project website](#).

1.2 Installation

1.2.1 Prerequisites

Neurokernel requires

- Linux (other operating systems may work, but have not been tested);
- Python 2.7 (Python 3.0 is not guaranteed to work);
- at least one NVIDIA GPU with [Fermi](#) architecture or later;
- NVIDIA's [GPU drivers](#);
- [CUDA 5.0](#) or later;
- [OpenMPI 1.8.4](#) or later compiled with CUDA support.

To check what GPUs are in your system, you can use the [inxi](#) command available on most Linux distributions:

¹ Chiang, A.-S., Lin, C.-Y., Chuang, C.-C., Chang, H.-M., Hsieh, C.-H., Yeh, C.-W., et al. (2011), Three-dimensional reconstruction of brain-wide wiring networks in *Drosophila* at single-cell resolution, *Current Biology*, 21, 1, 1–11, doi:10.1016/j.cub.2010.11.056

```
inxi -G
```

You can verify that the drivers are loaded as follows:

```
lsmod | grep nvidia
```

If no drivers are present, you may have to manually load them by running something like:

```
modprobe nvidia
```

as root.

Although some Linux distributions do include CUDA in their stock package repositories, you are encouraged to use those distributed by NVIDIA because they often are more up-to-date and include more recent releases of the GPU drivers. See [this page](#) for download information.

If you install Neurokernel in a virtualenv environment, you will need to install OpenMPI. See [this page](#) for OpenMPI installation information. *Note that OpenMPI 1.8 cannot run on Windows.*

Some of Neurokernel's demos require either `ffmpeg` or `libav` installed to generate visualizations (see [Examples](#)).

1.2.2 Installation

Download the latest Neurokernel code as follows:

```
git clone https://github.com/neurokernel/neurokernel.git
```

Since Neurokernel requires a fair number of additional Python packages to run, it is recommended that it either be installed in a [virtualenv](#) or [conda](#) environment. Follow the relevant instructions below.

Virtualenv

See [this page](#) for virtualenv installation information.

Create a new virtualenv environment and install several required dependencies:

```
cd ~/
virtualenv NK
~/NK/bin/pip install numpy cython numexpr pycuda
```

If installation of PyCUDA fails because some of the CUDA development files or libraries are not found, you may need to specify where they are explicitly. For example, if CUDA is installed in `/usr/local/cuda/`, try installing PyCUDA as follows:

```
CUDA_ROOT=/usr/local/cuda/ CFLAGS=-I${CUDA_ROOT}/include \
LDLFLAGS=-L${CUDA_ROOT}/lib64 ~/NK/bin/pip install pycuda
```

Replace `${CUDA_ROOT}/lib` with `${CUDA_ROOT}/lib64` if your system is running 64-bit Linux. If you continue to encounter installation problems, see the [PyCUDA Wiki](#) for more information.

Run the following to install the remaining Python package dependencies listed in `setup.py`:

```
cd ~/neurokernel
~/NK/bin/python setup.py develop
```

Conda

Note that conda packages are currently only available for 64-bit Ubuntu Linux 14.04. If you would like packages for another distribution, please submit a request to the [Neurokernel developers](#).

First, install the following Ubuntu packages:

- libibverbs1
- libnuma1
- libpmi0
- libslurm26
- libtorque2

These are required by the conda OpenMPI packages prepared for Neurokernel. Ensure that the stock Ubuntu OpenMPI packages are not installed because they may interfere with the ones that will be installed by conda. You also need to ensure that CUDA has been installed in `/usr/local/cuda`.

Install conda by either installing [Anaconda](#) or [Miniconda](#). Make sure that the following lines appear in your `~/.condarc` file so that conda can find the packages required by Neurokernel:

```
channels:
- https://conda.binstar.org/neurokernel/channel/ubuntu1404
- defaults
```

Create a new conda environment containing the packages required by Neurokernel by running the following command:

```
conda create -n NK neurokernel_deps
```

PyCUDA packages compiled against several versions of CUDA are available. If you need one compiled against a specific version that differs from the one automatically installed by the above command, you will need to manually install it afterwards as follows (replace `cuda75` with the appropriate version):

```
source activate NK
conda install pycuda=2015.1.3=np110py27_cuda75_0
source deactivate
```

Activate the new environment and install Neurokernel in it as follows:

```
source activate NK
cd ~/neurokernel
python setup.py develop
```

1.2.3 Examples

Introductory examples of how to use Neurokernel to build and integrate models of different parts of the fly brain are available in the [Neurodriver](#) package. To install it run the following:

```
git clone https://github.com/neurokernel/neurodriver
cd ~/neurodriver
python setup.py develop
```

Other models built using Neurokernel are available on [GitHub](#).

1.2.4 Building the Documentation

To build Neurokernel’s HTML documentation locally, you will need to install

- `mock` 1.0 or later.
- `sphinx` 1.3 or later.
- `sphinx_rtd_theme` 0.1.6 or later.

Once these are installed, run the following:

```
cd ~/neurokernel/docs
make html
```

1.3 Reference

1.3.1 Model Development API

Local Processing Units

<code>neurokernel.core.Module</code>	Processing module.
<code>neurokernel.core_gpu.Module</code>	Processing module.

neurokernel.core.Module

class `neurokernel.core.Module` (*sel, sel_in, sel_out, sel_gpot, sel_spike, data_gpot, data_spike, columns=['interface', 'io', 'type'], ctrl_tag=1, gpot_tag=2, spike_tag=3, id=None, device=None, routing_table=None, rank_to_id=None, debug=False, time_sync=False*)

Processing module.

This class repeatedly executes a work method until it receives a quit message via its control network port.

Parameters

- **sel** (*str, unicode, or sequence*) – Path-like selector describing the module’s interface of exposed ports.
- **sel_out, sel_gpot, sel_spike** (*sel_in,*) – Selectors respectively describing all input, output, graded potential, and spiking ports in the module’s interface.
- **data_spike** (*data_gpot,*) – Data arrays associated with the graded potential and spiking ports in the . Array length must equal the number of ports in a module’s interface.
- **columns** (*list of str*) – Interface port attributes. Network port for controlling the module instance.
- **gpot_tag, spike_tag** (*ctrl_tag,*) – MPI tags that respectively identify messages containing control data, graded potential port values, and spiking port values transmitted to worker nodes.
- **id** (*str*) – Module identifier. If no identifier is specified, a unique identifier is automatically generated.
- **device** (*int*) – GPU device to use. May be set to None if the module does not perform GPU processing.

- **routing_table** (`neurokernel.routing_table.RoutingTable`) – Routing table describing data connections between modules. If no routing table is specified, the module will be executed in isolation.
- **rank_to_id** (`bidict.bidict`) – Mapping between MPI ranks and module object IDs.
- **debug** (`bool`) – Debug flag. When True, exceptions raised during the work method are not be suppressed.
- **time_sync** (`bool`) – Time synchronization flag. When True, debug messages are not emitted during module synchronization and the time taken to receive all incoming data is computed.

interface

Interface – Object containing information about a module’s ports.

pm

dict – `pm[‘gpot’]` and `pm[‘spike’]` are instances of `neurokernel.pm.PortMapper` that map a module’s ports to the contents of the values in *data*.

data

dict – `data[‘gpot’]` and `data[‘spike’]` are arrays of data associated with a module’s graded potential and spiking ports.

`__init__` (*sel*, *sel_in*, *sel_out*, *sel_gpot*, *sel_spike*, *data_gpot*, *data_spike*, *columns*=[‘interface’, ‘io’, ‘type’], *ctrl_tag*=1, *gpot_tag*=2, *spike_tag*=3, *id*=None, *device*=None, *routing_table*=None, *rank_to_id*=None, *debug*=False, *time_sync*=False)

Methods

<code>__init__</code> (<i>sel</i> , <i>sel_in</i> , <i>sel_out</i> , <i>sel_gpot</i> , ...)	
<code>do_work</code> ()	Work method.
<code>post_run</code> ()	Code to run after main loop.
<code>pre_run</code> ()	Code to run before main loop.
<code>recv_parent</code> ([<i>tag</i>])	Receive data from parent process.
<code>recv_peer</code> ([<i>source</i> , <i>tag</i>])	
<code>run</code> ()	Body of process.
<code>run_step</code> ()	Module work method.
<code>send_parent</code> (<i>data</i> [, <i>tag</i>])	Send data to parent process.
<code>send_peer</code> (<i>data</i> , <i>dest</i> [, <i>tag</i>])	Send data to peer process.

neurokernel.core_gpu.Module

class `neurokernel.core_gpu.Module` (*sel*, *sel_in*, *sel_out*, *sel_gpot*, *sel_spike*, *data_gpot*, *data_spike*, *columns*=[‘interface’, ‘io’, ‘type’], *ctrl_tag*=1, *gpot_tag*=2, *spike_tag*=3, *id*=None, *device*=None, *routing_table*=None, *rank_to_id*=None, *debug*=False, *time_sync*=False)

Processing module.

This class repeatedly executes a work method until it receives a quit message via its control network port.

Parameters

- **sel** (*str*, *unicode*, or *sequence*) – Path-like selector describing the module’s interface of exposed ports.

- **sel_out, sel_gpot, sel_spike** (*sel_in*,) – Selectors respectively describing all input, output, graded potential, and spiking ports in the module’s interface.
- **data_spike** (*data_gpot*,) – Data arrays associated with the graded potential and spiking ports in the . Array length must equal the number of ports in a module’s interface.
- **columns** (*list of str*) – Interface port attributes. Network port for controlling the module instance.
- **gpot_tag, spike_tag** (*ctrl_tag*,) – MPI tags that respectively identify messages containing control data, graded potential port values, and spiking port values transmitted to worker nodes.
- **id** (*str*) – Module identifier. If no identifier is specified, a unique identifier is automatically generated.
- **device** (*int*) – GPU device to use. May be set to None if the module does not perform GPU processing.
- **routing_table** (`neurokernel.routing_table.RoutingTable`) – Routing table describing data connections between modules. If no routing table is specified, the module will be executed in isolation.
- **rank_to_id** (*bidict.bidict*) – Mapping between MPI ranks and module object IDs.
- **debug** (*bool*) – Debug flag. When True, exceptions raised during the work method are not be suppressed.
- **time_sync** (*bool*) – Time synchronization flag. When True, debug messages are not emitted during module synchronization and the time taken to receive all incoming data is computed.

interface

Interface – Object containing information about a module’s ports.

pm

dict – *pm*['gpot'] and *pm*['spike'] are instances of `neurokernel.pm_gpu.PortMapper` that map a module’s ports to the contents of the values in *data*.

data

dict – *data*['gpot'] and *data*['spike'] are arrays of data associated with a module’s graded potential and spiking ports.

`__init__` (*sel, sel_in, sel_out, sel_gpot, sel_spike, data_gpot, data_spike, columns=['interface', 'io', 'type'], ctrl_tag=1, gpot_tag=2, spike_tag=3, id=None, device=None, routing_table=None, rank_to_id=None, debug=False, time_sync=False*)

Methods

<code>__init__</code> (<i>sel, sel_in, sel_out, sel_gpot, ...</i>)	
<code>do_work</code> ()	Work method.
<code>post_run</code> ()	Code to run after main loop.
<code>pre_run</code> ()	Code to run before main loop.
<code>recv_parent</code> ([<i>tag</i>])	Receive data from parent process.
<code>recv_peer</code> ([<i>source, tag</i>])	
<code>run</code> ()	Body of process.
<code>run_step</code> ()	Module work method.

Continued on next page

Table 1.3 – continued from previous page

<code>send_parent(data[, tag])</code>	Send data to parent process.
<code>send_peer(data, dest[, tag])</code>	Send data to peer process.

Inter-LPU Connectivity

<code>neurokernel.pattern.Interface</code>	Container for set of interface comprising ports.
<code>neurokernel.pattern.Pattern</code>	Connectivity pattern linking sets of interface ports.

neurokernel.pattern.Interface

class `neurokernel.pattern.Interface` (*selector=""*, *columns=['interface', 'io', 'type']*)

Container for set of interface comprising ports.

This class contains information about a set of interfaces comprising path-like identifiers and the attributes associated with them. By default, each port must have at least the following attributes; other attributes may be added:

- `interface` - indicates which interface a port is associated with.
- `io` - indicates whether the port receives input ('in') or emits output ('out').
- `type` - indicates whether the port emits/receives spikes or graded potentials.

All port identifiers in an interface must be unique. For two interfaces to be deemed compatible, they must contain the same port identifiers and their identifiers' 'io' attributes must be the inverse of each other (i.e., every 'in' port in one interface must be mirrored by an 'out' port in the other interface).

Examples

```
>>> i = Interface('/foo[0:4],/bar[0:3]')
>>> i['/foo[0:2]', 'interface', 'io', 'type'] = [0, 'in', 'spike']
>>> i['/foo[2:4]', 'interface', 'io', 'type'] = [1, 'out', 'spike']
```

data

`pandas.DataFrame` – Port attribute data.

index

`pandas.MultiIndex` – Index of port identifiers.

Parameters

- **selector** (*str*, *unicode*, or *sequence*) – Selector string (e.g., 'foo[0:2]') or sequence of token sequences (e.g., [['foo', (0, 2)]]) describing the port identifiers comprised by the interface.
- **columns** (*list*, *default* = ['interface', 'io', 'type']) – Data column names.

See also:

`plsel.SelectorMethods`

`__init__` (*selector=""*, *columns=['interface', 'io', 'type']*)

Methods

<code>__init__([selector, columns])</code>	
<code>clear()</code>	Clear all ports in class instance.
<code>copy()</code>	Make a copy of this object.
<code>data_select(f[, inplace])</code>	Restrict Interface data with a selection function.
<code>equals(other)</code>	Check whether this interface is equivalent to another interface.
<code>from_csv(file_name, **kwargs)</code>	Create an Interface from a properly formatted CSV file.
<code>from_df(df)</code>	Create an Interface from a properly formatted DataFrame.
<code>from_dict(d)</code>	Create an Interface from a dictionary of selectors and data values.
<code>from_graph(g)</code>	Create an Interface from a NetworkX graph.
<code>from_selectors(sel[, sel_in, sel_out, ...])</code>	Create an Interface instance from selectors.
<code>get_common_ports(a, i, b[, t])</code>	Get port identifiers common to this and another Interface instance.
<code>gpot_ports([i, tuples])</code>	Restrict Interface ports to graded potential ports.
<code>in_ports([i, tuples])</code>	Restrict Interface ports to input ports.
<code>interface_ports([i, tuples])</code>	Restrict Interface ports to specific interface.
<code>is_compatible(a, i, b[, allow_subsets])</code>	Check whether two interfaces can be connected.
<code>is_in_interfaces(s)</code>	Check whether ports comprised by a selector are in the stored interfaces.
<code>out_ports([i, tuples])</code>	Restrict Interface ports to output ports.
<code>port_select(f[, inplace])</code>	Restrict Interface ports with a selection function.
<code>set_pm(t, pm)</code>	Set port mapper associated with a specific port type.
<code>spike_ports([i, tuples])</code>	Restrict Interface ports to spiking ports.
<code>to_selectors([i])</code>	Retrieve Interface's port identifiers as list of path-like selectors.
<code>to_tuples([i])</code>	Retrieve Interface's port identifiers as list of tuples.
<code>which_int(s)</code>	Return the interface containing the identifiers comprised by a selector.

neurokernel.pattern.Pattern

class `neurokernel.pattern.Pattern` (**selectors, **kwargs*)

Connectivity pattern linking sets of interface ports.

This class represents connection mappings between interfaces comprising sets of ports. Ports are represented using path-like identifiers; the presence of a row linking the two identifiers in the class' internal index indicates the presence of a connection. A single data attribute ('conn') associated with defined connections is created by default. Specific attributes may be accessed by specifying their names after the port identifiers; if a nonexistent attribute is specified when a sequential value is assigned, a new column for that attribute is automatically created:

```
p['/x[0]', '/y[0]', 'conn', 'x'] = [1, 'foo']
```

The direction of connections between ports in a class instance determines whether they are input or output ports. Ports may not both receive input or emit output. Patterns may contain fan-out connections, i.e., one source port connected to multiple destination ports, but not fan-in connections, i.e., multiple source ports connected to a single destination port.

Examples

```
>>> p = Pattern('/x[0:3]', '/y[0:4]')
>>> p['/x[0]', '/y[0:2]'] = 1
>>> p['/y[2]', '/x[1]'] = 1
>>> p['/y[3]', '/x[2]'] = 1
```

data

pandas.DataFrame – Connection attribute data.

index

pandas.MultiIndex – Index of connections.

interface

Interface – Interfaces containing port identifiers and attributes.

Parameters

- **sel1**, .. (*sel0*,) – Selectors defining the sets of ports potentially connected by the pattern. These selectors must be disjoint, i.e., no identifier comprised by one selector may be in any other selector.
- **columns** (*sequence of str*) – Data column names.

See also:

`plsel.SelectorMethods`

`__init__` (**selectors*, ***kwargs*)

Methods

<code>__init__</code> (* <i>selectors</i> , ** <i>kwargs</i>)	
<code>clear()</code>	Clear all connections in class instance.
<code>connected_port_pairs</code> ([<i>as_str</i>])	Return connections as pairs of port identifiers.
<code>connected_ports</code> ([<i>i</i> , <i>tuples</i>])	Return ports that are connected by the pattern.
<code>dest_idx</code> (<i>src_int</i> , <i>dest_int</i> [, <i>src_type</i> , ...])	Retrieve destination ports connected to the specified source ports.
<code>from_concat</code> (* <i>selectors</i> , ** <i>kwargs</i>)	Create pattern from the concatenation of identifiers in two selectors.
<code>from_csv</code> (<i>file_name</i> , ** <i>kwargs</i>)	Read connectivity data from CSV file.
<code>from_df</code> (<i>df_int</i> , <i>df_pat</i>)	Create a Pattern from properly formatted DataFrames.
<code>from_graph</code> (<i>g</i>)	Convert a NetworkX directed graph into a Pattern instance.
<code>from_product</code> (* <i>selectors</i> , ** <i>kwargs</i>)	Create pattern from the product of identifiers comprised by two selectors.
<code>gpot_ports</code> ([<i>i</i> , <i>tuples</i>])	Restrict Interface ports to graded potential ports.
<code>in_ports</code> ([<i>i</i> , <i>tuples</i>])	Restrict Interface ports to input ports.
<code>interface_ports</code> ([<i>i</i> , <i>tuples</i>])	Restrict Interface ports to specific interface.
<code>is_connected</code> (<i>from_int</i> , <i>to_int</i>)	Check whether the specified interfaces are connected.
<code>is_in_interfaces</code> (<i>selector</i>)	Check whether a selector is supported by any stored interface.
<code>out_ports</code> ([<i>i</i> , <i>tuples</i>])	Restrict Interface ports to output ports.

Continued on next page

Table 1.6 – continued from previous page

<code>spike_ports([i, tuples])</code>	Restrict Interface ports to spiking ports.
<code>split_multiindex(idx, a, b)</code>	Split a single MultiIndex into two instances.
<code>src_idx(src_int, dest_int[, src_type, ...])</code>	Retrieve source ports connected to the specified destination ports.
<code>to_graph()</code>	Convert the pattern to a networkx directed graph.
<code>which_int(s)</code>	Return the interface containing the identifiers comprised by a selector.

1.3.2 Emulation Management

Construction and Execution

<code>neurokernel.core.Manager</code>	Module manager.
<code>neurokernel.core_gpu.Manager</code>	Module manager.

neurokernel.core.Manager

class `neurokernel.core.Manager` (*required_args=['sel', 'sel_in', 'sel_out', 'sel_gpot', 'sel_spike'], ctrl_tag=1*)

Module manager.

Instantiates, connects, starts, and stops modules comprised by an emulation. All modules and connections must be added to a module manager instance before they can be run.

ctrl_tag

int – MPI tag to identify control messages.

modules

dict – Module instances. Keyed by module object ID.

routing_table

routing_table.RoutingTable – Table of data transmission connections between modules.

rank_to_id

bidict.bidict – Mapping between MPI ranks and module object IDs.

__init__ (*required_args=['sel', 'sel_in', 'sel_out', 'sel_gpot', 'sel_spike'], ctrl_tag=1*)

Methods

<code>__init__([required_args, ctrl_tag])</code>	
<code>add(target, id, *args, **kwargs)</code>	Add a module class to the emulation.
<code>connect(id_0, id_1, pat[, int_0, int_1])</code>	Specify connection between two module instances with a Pattern instance.
<code>process_worker_msg(msg)</code>	
<code>quit()</code>	Tell the workers to quit.
<code>recv([tag])</code>	Receive data from child process.
<code>send(data, dest[, tag])</code>	Send data to child process.
<code>spawn()</code>	Spawn MPI processes for and execute each of the managed targets.
<code>start([steps])</code>	Tell the workers to start processing data.

Continued on next page

Table 1.8 – continued from previous page

<code>stop()</code>	Tell the workers to stop processing data.
<code>validate_args(target)</code>	Check whether a class' constructor has specific arguments.
<code>wait()</code>	

neurokernel.core_gpu.Manager

class `neurokernel.core_gpu.Manager` (*required_args=['sel', 'sel_in', 'sel_out', 'sel_gpot', 'sel_spike'], ctrl_tag=1*)

Module manager.

Instantiates, connects, starts, and stops modules comprised by an emulation. All modules and connections must be added to a module manager instance before they can be run.

ctrl_tag

int – MPI tag to identify control messages.

modules

dict – Module instances. Keyed by module object ID.

routing_table

routing_table.RoutingTable – Table of data transmission connections between modules.

rank_to_id

bidict.bidict – Mapping between MPI ranks and module object IDs.

`__init__` (*required_args=['sel', 'sel_in', 'sel_out', 'sel_gpot', 'sel_spike'], ctrl_tag=1*)

Methods

<code>__init__</code> ([<i>required_args</i> , <i>ctrl_tag</i>])	
<code>add</code> (<i>target</i> , <i>id</i> , <i>*args</i> , <i>**kwargs</i>)	Add a module class to the emulation.
<code>connect</code> (<i>id_0</i> , <i>id_1</i> , <i>pat</i> [, <i>int_0</i> , <i>int_1</i>])	Specify connection between two module instances with a Pattern instance.
<code>process_worker_msg</code> (<i>msg</i>)	
<code>quit</code> ()	Tell the workers to quit.
<code>recv</code> ([<i>tag</i>])	Receive data from child process.
<code>send</code> (<i>data</i> , <i>dest</i> [, <i>tag</i>])	Send data to child process.
<code>spawn</code> ()	Spawn MPI processes for and execute each of the managed targets.
<code>start</code> ([<i>steps</i>])	Tell the workers to start processing data.
<code>stop</code> ()	Tell the workers to stop processing data.
<code>validate_args</code> (<i>target</i>)	Check whether a class' constructor has specific arguments.
<code>wait</code> ()	

Support Classes

<code>neurokernel.routing_table.RoutingTable</code>	Routing table class.
<code>neurokernel.mpi.Worker</code>	MPI worker class.
<code>neurokernel.mpi.WorkerManager</code>	Self-launching MPI worker manager.

neurokernel.routing_table.RoutingTable

class neurokernel.routing_table.**RoutingTable** (*g=None*)

Routing table class.

Simple class that stores pairs of strings that can signify one-hop routes between entities in a graph. Assigning a value to a pair that isn't in the class instance will result in the pair and value being added to the class instance. All data associated with a specific connection is stored as a dict; if a non-dict is assigned to a connection, it is stored in a dict with key 'data'. Specific values in the dict can be retrieved by passing the desired key directly to the [] operator.

Examples

```
>>> r = RoutingTable()
>>> r['a', 'b'] = 1
>>> r['a', 'b']
1
>>> r['a', 'c'] = [1, 2]
>>> r['a', 'c']
[1, 2]
>>> r['a', 'c'] = {'x': 1}
>>> r['a', 'c', 'x']
1
>>> r['a', 'c']['x']
1
```

Parameters *g* (*networkx.DiGraph*) – Directed graph that describes the routes between entities.

connections

list – List of directed connections between identifiers.

ids

list – Identifiers currently in routing table.

copy()

Return a copy of the routing table.

dest_ids (*src_id*)

Destination identifiers connected to the specified source identifier.

has_node (*n*)

Check whether the routing table contains the specified identifier.

ids()

IDs currently in routing

src_ids (*dest_id*)

Source identifiers connected to the specified destination identifier.

subtable (*ids*)

Return subtable containing only those connections between specified identifiers.

to_df()

Return a pandas DataFrame listing all of the connections.

__init__ (*g=None*)

Methods

<code>__init__([g])</code>	
<code>copy()</code>	
<code>dest_ids(src_id)</code>	Destination identifiers connected to the specified source identifier.
<code>has_node(n)</code>	Check whether the routing table contains the specified identifier.
<code>src_ids(dest_id)</code>	Source identifiers connected to the specified destination identifier.
<code>subtable(ids)</code>	Return subtable containing only those connections between specified identifiers.
<code>to_df()</code>	Return a pandas DataFrame listing all of the connections.

neurokernel.mpi.Worker

class `neurokernel.mpi.Worker` (*ctrl_tag=1*, *args, **kwargs)

MPI worker class.

This class repeatedly executes a work method.

Parameters `ctrl_tag` (*int*) – MPI tag to identify control messages transmitted to worker nodes.

`__init__` (*ctrl_tag=1*, *args, **kwargs)

Methods

<code>__init__([ctrl_tag])</code>	
<code>do_work()</code>	Work method.
<code>post_run()</code>	Code to run after main loop.
<code>pre_run()</code>	Code to run before main loop.
<code>recv_parent([tag])</code>	Receive data from parent process.
<code>recv_peer([source, tag])</code>	
<code>run()</code>	Main body of worker process.
<code>send_parent(data[, tag])</code>	Send data to parent process.
<code>send_peer(data, dest[, tag])</code>	Send data to peer process.

neurokernel.mpi.WorkerManager

class `neurokernel.mpi.WorkerManager` (*ctrl_tag=1*)

Self-launching MPI worker manager.

This class may be used to construct an MPI application consisting of

- a manager process that spawns MPI processes that execute the `run()` methods of several subclasses of the `Worker` class;
- worker processes that perform some processing task; and

The application should NOT be started via `mpiexec`.

Parameters `ctrl_tag` (*int*) – MPI tag to identify control messages transmitted to worker nodes.

May not be equal to `mpi4py.MPI.ANY_TAG`

Notes

This class requires MPI-2 dynamic processing management.

See also:

Worker

`__init__(ctrl_tag=1)`

Methods

<code>__init__([ctrl_tag])</code>	
<code>add(target, *args, **kwargs)</code>	Add a worker to an MPI application.
<code>process_worker_msg(msg)</code>	Process the specified deserialized message from a worker.
<code>quit()</code>	Tell the workers to quit.
<code>recv([tag])</code>	Receive data from child process.
<code>send(data, dest[, tag])</code>	Send data to child process.
<code>spawn()</code>	Spawn MPI processes for and execute each of the managed targets.
<code>start([steps])</code>	Tell the workers to start processing data.
<code>stop()</code>	Tell the workers to stop processing data.
<code>wait()</code>	Wait for execution to complete.

1.3.3 Support Classes and Functions

Path-Like Port Identifier Handling

<i>Selector</i>	Validated and expanded port selector.
<i>SelectorMethods</i>	Class for manipulating and using path-like selectors.
<i>SelectorParser</i>	This class implements a parser for path-like selectors that can be associated with elements in a sequential data structure such as a Pandas DataFrame; in the latter case, each level of the selector corresponds to a level of a Pandas MultiIndex.

neurokernel.plsel.Selector

class `neurokernel.plsel.Selector(s)`

Validated and expanded port selector.

Parameters `s` (*Selector*, *str*, or *unicode*) – Existing Selector class instance or string representation. The selector may not be ambiguous. If an existing Selector instance is specified, the new instance is a copy of the existing instance.

str

str – String representation of selector.

expanded

tuple of tuples – Expanded selector.

max_levels

int – Maximum number of levels in selector.

`__init__(s)`

Methods

<code>__init__(s)</code>	
<code>add(*sels)</code>	Combine the identifiers in multiple selectors into a single selector.
<code>add_str(*s)</code>	Combine the identifiers in multiple selector strings into a single selector.
<code>concat(*sels)</code>	Concatenate the identifiers in multiple selectors element-wise.
<code>prod(*sels)</code>	Compute the product of identifiers in multiple selectors.
<code>union(*sels)</code>	Compute the union of the identifiers in multiple selectors.

neurokernel.plsel.SelectorMethods

class neurokernel.plsel.SelectorMethods

Class for manipulating and using path-like selectors.

Contains class methods for expanding selectors, selecting rows from a Pandas DataFrame using a selector, etc.

The class can also be used to create new MultiIndex instances from selectors that can be fully expanded into an explicit set of identifiers (and therefore contain no ambiguous symbols such as '*' or '[':]').

`__init__()`

`x.__init__(...)` initializes x; see `help(type(x))` for signature

Methods

<code>are_consecutive(int_list)</code>	Check whether a list of integers is consecutive.
<code>are_disjoint(*selectors)</code>	Check whether several selectors are disjoint.
<code>collapse(selector)</code>	Collapse a selector into a single string.
<code>count_ports(selector)</code>	Count number of distinct port identifiers in unambiguous selector.
<code>expand(selector[, pad_len])</code>	Expand an unambiguous selector into a list of identifiers.
<code>get_index(df, selector[, start, stop, names])</code>	Return index corresponding to rows selected by specified selector.
<code>get_tuples(df, selector[, start, stop])</code>	Return tuples containing index labels selected by specified selector.
<code>index_to_selector(idx)</code>	Convert an index into an expanded port selector.
<code>is_ambiguous(selector)</code>	Check whether a selector cannot be expanded into an explicit list of identifiers.
<code>is_expandable(selector)</code>	Check whether a selector can be expanded into multiple identifiers.

Continued on next page

Table 1.16 – continued from previous page

<code>is_identifier(s)</code>	Check whether a selector or token sequence can identify a single port.
<code>is_in(s, t)</code>	Check whether all of the identifiers in one selector are comprised by another.
<code>is_selector(s)</code>	Check whether a string or sequence is a valid selector.
<code>is_selector_empty(selector)</code>	Check whether a string or sequence is an empty selector.
<code>is_selector_seq(s)</code>	Check whether a sequence is a valid selector.
<code>is_selector_str(s)</code>	Check whether a string is a valid selector.
<code>make_index(selector[, names])</code>	Create an index from the specified selector.
<code>make_index_two_concat(sel_0, sel_1[, names])</code>	Create an index from two selectors concatenated elementwise.
<code>make_index_two_prod(sel_0, sel_1[, names])</code>	Create an index from the product of two selectors.
<code>max_levels(selector)</code>	Return maximum number of token levels in selector.
<code>p_error(p)</code>	
<code>p_level(p)</code>	level : ASTERISK
<code>p_selector_comma_selector(p)</code>	selector : selector COMMA selector
<code>p_selector_dotplus_selector(p)</code>	selector : selector DOTPLUS selector
<code>p_selector_level(p)</code>	selector : level
<code>p_selector_paren_selector(p)</code>	selector : LPAREN selector RPAREN
<code>p_selector_plus_selector(p)</code>	selector : selector PLUS selector
<code>p_selector_selector_level(p)</code>	selector : selector level
<code>p_selector_selector_plus_level(p)</code>	selector : selector PLUS level
<code>pad_parsed(selector[, pad_len, inplace])</code>	Pad token lists in a parsed selector to some maximum length.
<code>pad_selector(selector[, pad_len])</code>	Expand and pad a selector with blank tokens.
<code>pad_tuple_list(tuple_list, pad_len)</code>	Pad a list of tuples with blank strings.
<code>parse(selector[, pad_len])</code>	Parse a selector string into tokens.
<code>select(df, selector[, start, stop])</code>	Select rows from DataFrame using a path-like selector.
<code>t_ASTERISK(t)</code>	/*
<code>t_COMMA(t)</code>	,
<code>t_DOTPLUS(t)</code>	.+
<code>t_INTEGER(t)</code>	?d+
<code>t_INTEGER_SET(t)</code>	?[(?:d+,?)+]
<code>t_INTERVAL(t)</code>	?[d*:d*]
<code>t_LPAREN(t)</code>	(
<code>t_PLUS(t)</code>	+
<code>t_RPAREN(t)</code>)
<code>t_STRING(t)</code>	/[^*/[()[:.,d][^+*/[()[:.,,]*
<code>t_STRING_SET(t)</code>	?[(?:[^+*/[()[:.,d][^+*/[()[:.,,]*,?)+]
<code>t_error(t)</code>	
<code>to_identifier(s)</code>	Convert an expanded selector/token sequence into a single port identifier string.
<code>tokenize(selector)</code>	Tokenize a selector string.
<code>tokens_to_str(tokens)</code>	Convert expanded/parsed token sequence into a single selector string.

neurokernel.plsel.SelectorParser

class neurokernel.plsel.SelectorParser

This class implements a parser for path-like selectors that can be associated with elements in a sequential data structure such as a Pandas DataFrame; in the latter case, each level of the selector corresponds to a level of a

Pandas MultiIndex. An index level may either be denoted by a string label (e.g., 'foo') or a numerical index (e.g., 0, 1, 2); a selector level may additionally be a list of strings (e.g., ['foo,bar']) or integers (e.g., '[0,2,4]') or continuous intervals (e.g., '[0:5]'). The '*' symbol matches any value in a level, while a range with an open upper bound (e.g., '[5:]') will match all integers greater than or equal to the lower bound.

Examples of valid selectors include

Selector	Comments
/foo/bar	
/foo+/bar	equivalent to /foo/bar
/foo/[qux,bar]	
/foo/bar[0]	
/foo/bar/[0]	equivalent to /foo/bar[0]
/foo/bar/0	equivalent to /foo/bar[0]
/foo/bar[0,1]	
/foo/bar[0:5]	
/foo*/baz	
/foo*/baz[5]	
/foo/bar,/baz/qux	
(/foo,/bar)+/baz	equivalent to /foo/baz,/bar/baz
/[foo,bar].+/[0:2]	equivalent to /foo[0],/bar[1]

Notes

An empty string is deemed to be a valid selector.

Since there is no need to maintain multiple instances of the lexer/parser used to process path-like selectors, they are associated with the class rather than class instances; likewise, all of the class' methods are classmethods.

Numerical indices in selectors are assumed to be zero-based. Intervals do not include the end element (i.e., like numpy, not like Pandas).

```
__init__()
    x.__init__(...) initializes x; see help(type(x)) for signature
```

Methods

<code>p_error(p)</code>	
<code>p_level(p)</code>	level : ASTERISK
<code>p_selector_comma_selector(p)</code>	selector : selector COMMA selector
<code>p_selector_dotplus_selector(p)</code>	selector : selector DOTPLUS selector
<code>p_selector_level(p)</code>	selector : level
<code>p_selector_paren_selector(p)</code>	selector : LPAREN selector RPAREN
<code>p_selector_plus_selector(p)</code>	selector : selector PLUS selector
<code>p_selector_selector_level(p)</code>	selector : selector level
<code>p_selector_selector_plus_level(p)</code>	selector : selector PLUS level
<code>pad_parsed(selector[, pad_len, inplace])</code>	Pad token lists in a parsed selector to some maximum length.
<code>parse(selector[, pad_len])</code>	Parse a selector string into tokens.
<code>t_ASTERISK(t)</code>	/*
<code>t_COMMA(t)</code>	,

Continued on next page

Table 1.17 – continued from previous page

<code>t_DOTPLUS(t)</code>	<code>.+</code>
<code>t_INTEGER(t)</code>	<code>/?d+</code>
<code>t_INTEGER_SET(t)</code>	<code>/?[(?:d+,?)+]</code>
<code>t_INTERVAL(t)</code>	<code>/?[d*:d*]</code>
<code>t_LPAREN(t)</code>	<code>(</code>
<code>t_PLUS(t)</code>	<code>+</code>
<code>t_RPAREN(t)</code>	<code>)</code>
<code>t_STRING(t)</code>	<code>/[^*/[()?:.,d][^+*/[()?:.,]*</code>
<code>t_STRING_SET(t)</code>	<code>/?[(?:[+*/[()?:.,d][^+*/[()?:.,]*,?)+]</code>
<code>t_error(t)</code>	
<code>tokenize(selector)</code>	Tokenize a selector string.

GPU Port Mappers

<code>GPUPortMapper</code>	Maps a PyCUDA GPUArray to/from path-like port identifiers.
----------------------------	--

neurokernel.pm_gpu.GPUPortMapper

class `neurokernel.pm_gpu.GPUPortMapper` (*selector*, *data=None*, *portmap=None*, *make_copy=True*)

Maps a PyCUDA GPUArray to/from path-like port identifiers.

`__init__` (*selector*, *data=None*, *portmap=None*, *make_copy=True*)

Methods

<code>__init__(selector[, data, portmap, make_copy])</code>	
<code>copy()</code>	Return copy of this port mapper.
<code>equals(other)</code>	Check whether this mapper is equivalent to another mapper.
<code>from_index(idx, data[, portmap])</code>	
<code>from_pm(pm)</code>	Create a new port mapper instance given an existing instance.
<code>get(selector)</code>	
<code>get_by_inds(inds)</code>	Retrieve mapped data specified by integer index.
<code>get_inds_nonzero()</code>	
<code>get_map(selector)</code>	Retrieve integer indices associated with selector.
<code>get_ports(f)</code>	Select ports using a data selection function.
<code>get_ports_as_inds(f)</code>	Select integer indices corresponding to ports in map.
<code>get_ports_nonzero()</code>	
<code>inds_to_ports(inds)</code>	Convert list of integer indices to port identifiers.
<code>ports_to_inds(selector)</code>	Convert port selector to list of integer indices.
<code>set(selector, data)</code>	
<code>set_by_inds(inds, data)</code>	Set mapped data by integer indices.
<code>set_by_inds_array(inds, data)</code>	Set mapped data with array by integer indices.
<code>set_by_inds_scalar(inds, data)</code>	Set mapped data with scalar by integer indices.
<code>set_map(selector, portmap)</code>	Set mapped integer index associated with selector.

Python Port Mappers

<i>BasePortMapper</i>	Maps integer sequence to/from path-like port identifiers.
<i>PortMapper</i>	Maps a numpy array to/from path-like port identifiers.

neurokernel.pm.BasePortMapper

class neurokernel.pm.**BasePortMapper** (*selector, portmap=None*)
 Maps integer sequence to/from path-like port identifiers.

Examples

```
>>> pm = BasePortMapper('/[a,b][0:2]')
>>> print pm.ports_to_inds('/b[0:2]')
array([2, 3])
>>> print pm.inds_to_ports([0, 1])
[('a', 0), ('a', 1)]
```

Parameters

- **selector** (*str, unicode, or sequence*) – Selector string (e.g., `'/foo[0:2]'`) or sequence of token sequences (e.g., `[['foo', (0, 2)]]`) to map to *data*.
- **portmap** (*sequence of int*) – Integer indices to map to port identifiers. If no map is specified, it is assumed to be an array of consecutive integers from 0 through one less than the number of ports.

index

pandas.MultiIndex – Index of port identifiers.

portmap

pandas.Series – Map of port identifiers to integer indices.

Notes

The selectors may not contain any `*` or `[:]` characters. A single port identifier may be mapped to multiple integer indices, but not vice-versa.

`__init__` (*selector, portmap=None*)

Methods

<code>__init__(selector[, portmap])</code>	
<code>copy()</code>	Return copy of this port mapper.
<code>equals(pm)</code>	Check whether this mapper is equivalent to another mapper.
<code>from_index(idx[, portmap])</code>	Create port mapper from a Pandas index and a sequence of integer indices.

Continued on next page

Table 1.21 – continued from previous page

<code>from_pm(pm)</code>	Create a new port mapper instance given an existing instance.
<code>get_map(selector)</code>	Retrieve integer indices associated with selector.
<code>inds_to_ports(inds)</code>	Convert list of integer indices to port identifiers.
<code>ports_to_inds(selector)</code>	Convert port selector to list of integer indices.
<code>set_map(selector, portmap)</code>	Set mapped integer index associated with selector.

neurokernel.pm.PortMapper

class `neurokernel.pm.PortMapper` (*selector*, *data=None*, *portmap=None*, *make_copy=True*)
 Maps a numpy array to/from path-like port identifiers.

Examples

```
>>> data = np.array([1, 0, 3, 2, 5, 2])
>>> pm = PortMapper('/d[0:5]', data)
>>> print pm['/d[1]']
array([0])
>>> print pm['/d[2:4]']
array([3, 2])
```

Parameters

- **selector** (*str*, *unicode*, or *sequence*) – Selector string (e.g., `'/foo[0:2]'`) or sequence of token sequences (e.g., `[['foo', (0, 2)]]`) to map to *data*.
- **data** (*numpy.ndarray*) – 1D data array to map to ports. If no data array is specified, port identifiers will still be mapped to their sequential indices but `__getitem__()` and `__setitem__()` will raise exceptions if invoked.
- **portmap** (*sequence of int*) – Integer indices to map to port identifiers. If no map is specified, it is assumed to be an array of consecutive integers from 0 through one less than the number of ports.
- **make_copy** (*bool*) – If True, map a copy of the specified data array to the specified port identifiers.

data

numpy.ndarray – Data that has been mapped to ports.

dtype

numpy.dtype – Type of mapped data.

index

pandas.MultiIndex – Index of port identifiers.

portmap

pandas.Series – Map of port identifiers to integer indices into *data*.

Notes

The selectors may not contain any `'*'` or `'[:]'` characters.

`__init__` (*selector*, *data=None*, *portmap=None*, *make_copy=True*)

Methods

<code>__init__(selector[, data, portmap, make_copy])</code>	
<code>copy()</code>	Return copy of this port mapper.
<code>equals(other)</code>	Check whether this mapper is equivalent to another mapper.
<code>from_index(idx, data[, portmap])</code>	
<code>from_pm(pm)</code>	Create a new port mapper instance given an existing instance.
<code>get(selector)</code>	Retrieve mapped data specified by given selector.
<code>get_by_inds(inds)</code>	Retrieve mapped data specified by integer index.
<code>get_inds_nonzero()</code>	Select indices of ports with nonzero data.
<code>get_map(selector)</code>	Retrieve integer indices associated with selector.
<code>get_ports(f)</code>	Select ports using a data selection function.
<code>get_ports_as_inds(f)</code>	Select integer indices corresponding to ports in map.
<code>get_ports_nonzero()</code>	Select ports with nonzero data.
<code>inds_to_ports(inds)</code>	Convert list of integer indices to port identifiers.
<code>ports_to_inds(selector)</code>	Convert port selector to list of integer indices.
<code>set(selector, data)</code>	Set mapped data specified by given selector.
<code>set_by_inds(inds, data)</code>	Set mapped data by integer indices.
<code>set_map(selector, portmap)</code>	Set mapped integer index associated with selector.

XML Tools

<code>graph_to_nml_module</code>	Convert a module expressed as NetworkX graphs into Neurokernel NeuroML.
<code>graph_to_nml_pattern</code>	Convert a pattern expressed as a NetworkX graph into Neurokernel NeuroML.
<code>load</code>	Load a Neurokernel NeuroML document.
<code>nml_pattern_to_graph</code>	Convert a pattern expressed in Neurokernel NeuroML into a NetworkX graph.
<code>nml_module_to_graph</code>	Convert a module expressed in Neurokernel NeuroML into NetworkX graphs.
<code>write</code>	Write a Neurokernel NeuroML document to an XML file.

neurokernel.neuroml.utils.graph_to_nml_module

`neurokernel.neuroml.utils.graph_to_nml_module(g, i, id)`

Convert a module expressed as NetworkX graphs into Neurokernel NeuroML.

Parameters

- **g** (`networkx.DiGraph`) – Directed graph containing a module’s neurons and synapses
- **i** (`networkx.Graph`) – Undirected graph containing a module’s interface ports.
- **id** (`str`) – Module identifier.

Returns `module` – Module instance.

Return type `neurokernel.neuroml.Module`

neurokernel.neuroml.utils.graph_to_nml_pattern

neurokernel.neuroml.utils.**graph_to_nml_pattern**(*g*, *id*)

Convert a pattern expressed as a NetworkX graph into Neurokernel NeuroML.

Parameters

- **g** (*networkx.DiGraph*) – Directed graph containing the pattern’s ports (nodes) and connections (edges).
- **id** (*str*) – Pattern identifier.

Returns **pattern** – pattern instance.

Return type neurokernel.neuroml.Pattern

neurokernel.neuroml.utils.load

neurokernel.neuroml.utils.**load**(*file*)

Load a Neurokernel NeuroML document.

Parameters **file** (*str* or *file*) – Input file name or handle.

Returns **nk_doc** – Neurokernel NeuroML document root.

Return type neurokernel.neuroml.NeurokernelDoc

neurokernel.neuroml.utils.nml_pattern_to_graph

neurokernel.neuroml.utils.**nml_pattern_to_graph**(*pattern*)

Convert a pattern expressed in Neurokernel NeuroML into a NetworkX graph.

Parameters **pattern** (*neurokernel.neuroml.Pattern*) – Pattern instance.

Returns **g** – Directed graph containing the pattern’s ports (nodes) and connections (edges).

Return type networkx.DiGraph

neurokernel.neuroml.utils.nml_module_to_graph

neurokernel.neuroml.utils.**nml_module_to_graph**(*module*)

Convert a module expressed in Neurokernel NeuroML into NetworkX graphs.

Parameters **module** (*neurokernel.neuroml.Module*) – Module instance.

Returns

- **g** (*networkx.DiGraph*) – Directed graph containing a module’s neurons and synapses
- **i** (*networkx.Graph*) – Undirected graph containing a module’s interface ports.

neurokernel.neuroml.utils.write

neurokernel.neuroml.utils.**write**(*nk_doc*, *file*, *root_name='nk'*)

Write a Neurokernel NeuroML document to an XML file.

Parameters

- **nk_doc** (*neurokernel.neuroml.NeurokernelDoc*) – Neurokernel NeuroML document root object.
- **file** (*str or file*) – Output file name or handle.
- **root_name** (*str*) – Document root name.

Context Managers

<i>ExceptionOnSignal</i>	Raise a specific exception when the specified signal is detected.
<i>IgnoreKeyboardInterrupt</i>	Ignore keyboard interrupts.
<i>IgnoreSignal</i>	Ignore the specified signal.
<i>OnKeyboardInterrupt</i>	Respond to keyboard interrupt with specified handler.
<i>TryExceptionOnSignal</i>	Check for exception raised in response to specific signal.

neurokernel.ctx_managers.ExceptionOnSignal

`neurokernel.ctx_managers.ExceptionOnSignal(*args, **kws)`
 Raise a specific exception when the specified signal is detected.

neurokernel.ctx_managers.IgnoreKeyboardInterrupt

`neurokernel.ctx_managers.IgnoreKeyboardInterrupt(*args, **kws)`
 Ignore keyboard interrupts.

neurokernel.ctx_managers.IgnoreSignal

`neurokernel.ctx_managers.IgnoreSignal(*args, **kws)`
 Ignore the specified signal.

neurokernel.ctx_managers.OnKeyboardInterrupt

`neurokernel.ctx_managers.OnKeyboardInterrupt(*args, **kws)`
 Respond to keyboard interrupt with specified handler.

neurokernel.ctx_managers.TryExceptionOnSignal

`neurokernel.ctx_managers.TryExceptionOnSignal(*args, **kws)`
 Check for exception raised in response to specific signal.

GPU Tools

<i>bufint</i>	Return buffer interface to GPU or numpy array.
<i>set_by_inds</i>	Set values in a GPUArray by index.

Continued on next page

Table 1.25 – continued from previous page

<code>set_by_inds_from_inds</code>	Set values in a GPUArray by index from indexed values in another GPUArray.
<code>set_realloc</code>	Transfer data into a GPUArray instance.

neurokernel.tools.gpu.bufint

`neurokernel.tools.gpu.bufint(a)`

Return buffer interface to GPU or numpy array.

Parameters *a* (`pycuda.gpuarray.GPUArray` or `numpy.ndarray`) – GPU or numpy array.

Returns *b* – Buffer interface to array. Returns None if *a* has a length of 0.

Return type `buffer`

neurokernel.tools.gpu.set_by_inds

`neurokernel.tools.gpu.set_by_inds(dest_gpu, ind, src_gpu, ind_which='dest')`

Set values in a GPUArray by index.

Parameters

- **dest_gpu** (`pycuda.gpuarray.GPUArray`) – GPUArray instance to modify.
- **ind** (`pycuda.gpuarray.GPUArray` or `numpy.ndarray`) – 1D array of element indices to set. Must have an integer dtype.
- **src_gpu** (`pycuda.gpuarray.GPUArray`) – GPUArray instance from which to set values.
- **ind_which** (`str`) – If set to 'dest', set the elements in *dest_gpu* with indices *ind* to the successive values in *src_gpu*; the lengths of *ind* and *src_gpu* must be equal. If set to 'src', set the successive values in *dest_gpu* to the values in *src_gpu* with indices *ind*; the lengths of *ind* and *dest_gpu* must be equal.

Examples

```
>>> import pycuda.gpuarray as gpuarray
>>> import pycuda.autoint
>>> import numpy as np
>>> from nk.tools.gpu import set_by_inds
>>> dest_gpu = gpuarray.to_gpu(np.arange(5, dtype=np.float32))
>>> ind = gpuarray.to_gpu(np.array([0, 2, 4]))
>>> src_gpu = gpuarray.to_gpu(np.array([1, 1, 1], dtype=np.float32))
>>> set_by_inds(dest_gpu, ind, src_gpu, 'dest')
>>> np.allclose(dest_gpu.get(), np.array([1, 1, 1, 3, 1], dtype=np.float32))
True
>>> dest_gpu = gpuarray.to_gpu(np.zeros(3, dtype=np.float32))
>>> ind = gpuarray.to_gpu(np.array([0, 2, 4]))
>>> src_gpu = gpuarray.to_gpu(np.arange(5, dtype=np.float32))
>>> set_by_inds(dest_gpu, ind, src_gpu, 'src')
>>> np.allclose(dest_gpu.get(), np.array([0, 2, 4], dtype=np.float32))
True
```

Notes

Only supports 1D index arrays.

May not be efficient for certain index patterns because of lack of inability to coalesce memory operations.

neurokernel.tools.gpu.set_by_inds_from_inds

`neurokernel.tools.gpu.set_by_inds_from_inds(dest_gpu, ind_dest, src_gpu, ind_src)`

Set values in a GPUArray by index from indexed values in another GPUArray.

Parameters

- **dest_gpu** (`pycuda.gpuarray.GPUArray`) – GPUArray instance to modify.
- **ind_dest** (`pycuda.gpuarray.GPUArray` or `numpy.ndarray`) – 1D array of element indices in `dest_gpu` to set. Must have an integer dtype.
- **src_gpu** (`pycuda.gpuarray.GPUArray`) – GPUArray instance from which to set values.
- **ind_src** (`pycuda.gpuarray.GPUArray` or `numpy.ndarray`) – 1D array of element indices in `src_gpu` to copy. Must have an integer dtype and be the same length as `ind_dest`.

Examples

```
>>> import pycuda.gpuarray as gpuarray
>>> import pycuda.autoinit
>>> import numpy as np
>>> from nk.tools.gpu import set_by_inds_from_inds
>>> dest_gpu = gpuarray.to_gpu(np.zeros(5, dtype=np.float32))
>>> ind_dest = gpuarray.to_gpu(np.array([0, 2, 4]))
>>> src_gpu = gpuarray.to_gpu(np.arange(5, 10, dtype=np.float32))
>>> ind_src = gpuarray.to_gpu(np.array([2, 3, 4]))
>>> gpu.set_by_inds_from_inds(dest_gpu, ind_dest, src_gpu, ind_src)
>>> assert np.allclose(dest_gpu.get(), np.array([7, 0, 8, 0, 9], dtype=np.
↳float32))
True
```

neurokernel.tools.gpu.set_realloc

`neurokernel.tools.gpu.set_realloc(x_gpu, data)`

Transfer data into a GPUArray instance.

Copies the contents of a numpy array into a GPUArray instance. If the array has a different type or dimensions than the instance, the GPU memory used by the instance is reallocated and the instance updated appropriately.

Parameters

- **x_gpu** (`pycuda.gpuarray.GPUArray`) – GPUArray instance to modify.
- **data** (`numpy.ndarray`) – Array of data to transfer to the GPU.

Examples

```
>>> import pycuda.gpuarray as gpuarray
>>> import pycuda.autoinit
>>> import numpy as np
>>> import misc
>>> x = np.asarray(np.random.rand(5), np.float32)
>>> x_gpu = gpuarray.to_gpu(x)
>>> x = np.asarray(np.random.rand(10, 1), np.float64)
>>> set_realloc(x_gpu, x)
>>> np.allclose(x, x_gpu.get())
True
```

MPI Tools

MPIOutput

neurokernel.tools.mpi.MPIOutput

neurokernel.tools.mpi.MPIOutput = <Mock spec='str' id='140616444231504'>

ZeroMQ Tools

get_random_port

Return available random ZeroMQ port.

is_poll_in

Check for incoming data on a socket using a poller.

ZMQOutput

neurokernel.tools.zmq.get_random_port

neurokernel.tools.zmq.get_random_port (min_port=49152, max_port=65536, max_tries=100)
Return available random ZeroMQ port.

neurokernel.tools.zmq.is_poll_in

neurokernel.tools.zmq.is_poll_in (sock, poller, timeout=100)
Check for incoming data on a socket using a poller.

neurokernel.tools.zmq.ZMQOutput

neurokernel.tools.zmq.ZMQOutput = <Mock spec='str' id='140616444230224'>

Visualization Tools

imshow

Display the specified image file using matplotlib.

show_pydot

Display a networkx graph using pydot.

Continued on next page

Table 1.28 – continued from previous page

<code>show_pygraphviz</code>	Display a networkx graph using pygraphviz.
------------------------------	--

neurokernel.tools.plot.imdisp

`neurokernel.tools.plot.imdisp` (*f*)
 Display the specified image file using matplotlib.

neurokernel.tools.plot.show_pydot

`neurokernel.tools.plot.show_pydot` (*g*)
 Display a networkx graph using pydot.

neurokernel.tools.plot.show_pygraphviz

`neurokernel.tools.plot.show_pygraphviz` (*g*, *prog*='dot', *graph_attr*={}, *node_attr*={},
edge_attr={})

Display a networkx graph using pygraphviz.

Parameters

- **prog** (*str*) – Executable for generating the image.
- **graph_attr** (*dict*) – Global graph display attributes.
- **node_attr** (*dict*) – Global node display attributes.
- **edge_attr** (*dict*) – Global edge display attributes.

Logging Tools

<code>log_exception</code>	Log the specified exception data using twiggy.
<code>set_excepthook</code>	Set the exception hook to use the specified logger.
<code>setup_logger</code>	Setup a twiggy logger.

neurokernel.tools.logging.log_exception

`neurokernel.tools.logging.log_exception` (*type*, *value*, *tb*, *logger*=<Mock name='mock.log'
id='140616444276816'>, *multiline*=False)

Log the specified exception data using twiggy.

Parameters

- **tb** (*value*,) – Parameters expected by `traceback.print_exception`.
- **logger** (*twiggy.logger.Logger*) – Logger to use. `twiggy.log` is assumed by default.
- **multiline** (*bool*) – If True, print exception using multiple log lines.

neurokernel.tools.logging.set_excepthook

`neurokernel.tools.logging.set_excepthook` (*logger*, *multiline*=False)
 Set the exception hook to use the specified logger.

Parameters

- **logger** (*twiggy.logger.Logger*) – Configured logger.
- **multiline** (*bool*) – If True, log exception messages on multiple lines.

neurokernel.tools.logging.setup_logger

```
neurokernel.tools.logging.setup_logger (name="", level=<Mock
name='mock.levels.DEBUG'
id='140616444277136'>, fmt=<Mock
name='mock.formats.line_format'
id='140616444277392'>, fmt_name=<built-in
method format of str object>, screen=None,
file_name=None, mpi_comm=None,
zmq_addr=None, log_exceptions=True, multi-
line=False)
```

Setup a twiggy logger.

Parameters

- **name** (*str*) – Logger name.
- **level** (*twiggy.levels.LogLevel*) – Logging level.
- **fmt** (*twiggy.formats.LineFormat*) – Logging formatter class instance.
- **fmt_name** (*function*) – Function with one parameter that formats the message name.
- **screen** (*bool*) – Create output stream handler to the screen if True.
- **file_name** (*str*) – Create output handler to specified file.
- **mpi_comm** (*mpi4py.MPI.Intracomm*) – If not None, use MPI I/O with the specified communicator for output file handler. Ignored if the *file_name* parameter is not specified.
- **zmq_addr** (*str*) – ZeroMQ socket address.
- **log_exceptions** (*bool*) – If True, exception messages are written to the logger.
- **multiline** (*bool*) – If True, log exception messages on multiple lines.

Returns

- **logger** (*twiggy.logger.Logger*) – Configured logger.
- *Bug*
- —
- *To use the ZeroMQ output class with multiprocessing, it must be added*
- *as an emitter within each process.*

Other

LoggerMixin(name[, log_on])	Mixin that provides a per-instance logger that can be turned off.
-----------------------------	---

<code>catch_exception</code>	Catch and report exceptions when executing a function.
<code>rand_bin_matrix</code>	Generate a rectangular binary matrix with randomly distributed nonzero entries.

neurokernel.tools.misc.catch_exception

neurokernel.tools.misc.**catch_exception**(*func*, *disp*, **args*, ***kwargs*)

Catch and report exceptions when executing a function.

If an exception occurs while executing the specified function, the exception's message and the line number where it occurred (in the innermost traceback frame) are displayed.

Examples

```
>>> import sys
>>> def f(x): x/0
>>> catch_exception(f, sys.stdout.write, 1)
f: integer division or modulo by zero (...:1)
```

Parameters

- **func** (*function*) – Function to execute.
- **disp** (*function*) – Function to use to display exception message.
- **args** (*list*) – Function arguments.
- **kwargs** (*dict*) – Named function arguments.

neurokernel.tools.misc.rand_bin_matrix

neurokernel.tools.misc.**rand_bin_matrix**(*sh*, *N*, *dtype*=<Mock name='mock.double' id='140616444278608'>)

Generate a rectangular binary matrix with randomly distributed nonzero entries.

Examples

```
>>> m = rand_bin_matrix((2, 3), 3)
>>> set(m.flatten()) == set([0, 1])
True
```

Parameters

- **sh** (*tuple*) – Shape of generated matrix.
- **N** (*int*) – Number of entries to set to 1.
- **dtype** (*dtype*) – Generated matrix data type.

1.4 Authors & Acknowledgements

The Neurokernel Project was begun in July 2011 by [Aurel A. Lazar](#) at Columbia University's Department of Electrical Engineering after extensive discussions held during a [research seminar on Massively Parallel Neural Computation](#). The Neurokernel Development Team currently comprises the following [Bionet Group](#) researchers:

- [Lev Givon](#)
- [Konstantinos Psychas](#)
- [Nikul H. Ukani](#)
- [Chung-Heng Yeh](#)
- [Yiyin Zhou](#)

Past contributors who have participated in the project include

- [Daniel S. Chevitarese](#)

The Neurokernel logo is based upon the logo of the [FlyJunkies web site](#), used with permission for non-profit use by the site maintainers [Gavin Davis](#) and [Fraser Perry](#).

1.5 Licenses

1.5.1 Neurokernel

Copyright (c) 2012-2015, Neurokernel Development Team All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the names of the copyright holders nor the names of any contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

1.5.2 libNeuroML

Copyright (c) 2012, libNeuroML authors and contributors All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the names of the copyright holders nor the names of the contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

CHAPTER 2

Index

- genindex

Symbols

__init__() (neurokernel.core.Manager method), 12
 __init__() (neurokernel.core.Module method), 7
 __init__() (neurokernel.core_gpu.Manager method), 13
 __init__() (neurokernel.core_gpu.Module method), 8
 __init__() (neurokernel.mpi.Worker method), 15
 __init__() (neurokernel.mpi.WorkerManager method), 16
 __init__() (neurokernel.pattern.Interface method), 9
 __init__() (neurokernel.pattern.Pattern method), 11
 __init__() (neurokernel.plsel.Selector method), 17
 __init__() (neurokernel.plsel.SelectorMethods method), 17
 __init__() (neurokernel.plsel.SelectorParser method), 19
 __init__() (neurokernel.pm.BasePortMapper method), 21
 __init__() (neurokernel.pm.PortMapper method), 22
 __init__() (neurokernel.pm_gpu.GPUPortMapper method), 20
 __init__() (neurokernel.routing_table.RoutingTable method), 14

B

BasePortMapper (class in neurokernel.pm), 21
 bufint() (in module neurokernel.tools.gpu), 26

C

catch_exception() (in module neurokernel.tools.misc), 31
 connections (neurokernel.routing_table.RoutingTable attribute), 14
 copy() (neurokernel.routing_table.RoutingTable method), 14
 ctrl_tag (neurokernel.core.Manager attribute), 12
 ctrl_tag (neurokernel.core_gpu.Manager attribute), 13

D

data (neurokernel.core.Module attribute), 7
 data (neurokernel.core_gpu.Module attribute), 8
 data (neurokernel.pattern.Interface attribute), 9
 data (neurokernel.pattern.Pattern attribute), 11
 data (neurokernel.pm.PortMapper attribute), 22

dest_ids() (neurokernel.routing_table.RoutingTable method), 14

dtype (neurokernel.pm.PortMapper attribute), 22

E

ExceptionOnSignal() (in module neurokernel.ctx_managers), 25
 expanded (neurokernel.plsel.Selector attribute), 16

G

get_random_port() (in module neurokernel.tools.zmq), 28
 GPUPortMapper (class in neurokernel.pm_gpu), 20
 graph_to_nml_module() (in module neurokernel.neuroml.utils), 23
 graph_to_nml_pattern() (in module neurokernel.neuroml.utils), 24

H

has_node() (neurokernel.routing_table.RoutingTable method), 14

I

ids (neurokernel.routing_table.RoutingTable attribute), 14
 ids() (neurokernel.routing_table.RoutingTable method), 14
 IgnoreKeyboardInterrupt() (in module neurokernel.ctx_managers), 25
 IgnoreSignal() (in module neurokernel.ctx_managers), 25
 imdisp() (in module neurokernel.tools.plot), 29
 index (neurokernel.pattern.Interface attribute), 9
 index (neurokernel.pattern.Pattern attribute), 11
 index (neurokernel.pm.BasePortMapper attribute), 21
 index (neurokernel.pm.PortMapper attribute), 22
 Interface (class in neurokernel.pattern), 9
 interface (neurokernel.core.Module attribute), 7
 interface (neurokernel.core_gpu.Module attribute), 8
 interface (neurokernel.pattern.Pattern attribute), 11
 is_poll_in() (in module neurokernel.tools.zmq), 28

L

load() (in module neurokernel.neuroml.utils), 24
 log_exception() (in module neurokernel.tools.logging),
 29

M

Manager (class in neurokernel.core), 12
 Manager (class in neurokernel.core_gpu), 13
 max_levels (neurokernel.plsel.Selector attribute), 17
 Module (class in neurokernel.core), 6
 Module (class in neurokernel.core_gpu), 7
 modules (neurokernel.core.Manager attribute), 12
 modules (neurokernel.core_gpu.Manager attribute), 13
 MPIOutput (in module neurokernel.tools.mpi), 28

N

nml_module_to_graph() (in module neurokernel.neuroml.utils), 24
 nml_pattern_to_graph() (in module neurokernel.neuroml.utils), 24

O

OnKeyboardInterrupt() (in module neurokernel.ctx_managers), 25

P

Pattern (class in neurokernel.pattern), 10
 pm (neurokernel.core.Module attribute), 7
 pm (neurokernel.core_gpu.Module attribute), 8
 portmap (neurokernel.pm.BasePortMapper attribute), 21
 portmap (neurokernel.pm.PortMapper attribute), 22
 PortMapper (class in neurokernel.pm), 22

R

rand_bin_matrix() (in module neurokernel.tools.misc), 31
 rank_to_id (neurokernel.core.Manager attribute), 12
 rank_to_id (neurokernel.core_gpu.Manager attribute), 13
 routing_table (neurokernel.core.Manager attribute), 12
 routing_table (neurokernel.core_gpu.Manager attribute),
 13
 RoutingTable (class in neurokernel.routing_table), 14

S

Selector (class in neurokernel.plsel), 16
 SelectorMethods (class in neurokernel.plsel), 17
 SelectorParser (class in neurokernel.plsel), 18
 set_by_inds() (in module neurokernel.tools.gpu), 26
 set_by_inds_from_inds() (in module neurokernel.tools.gpu), 27
 set_excepthook() (in module neurokernel.tools.logging),
 29
 set_realloc() (in module neurokernel.tools.gpu), 27
 setup_logger() (in module neurokernel.tools.logging), 30

show_pydot() (in module neurokernel.tools.plot), 29
 show_pygraphviz() (in module neurokernel.tools.plot), 29
 src_ids() (neurokernel.routing_table.RoutingTable
 method), 14
 str (neurokernel.plsel.Selector attribute), 16
 subtable() (neurokernel.routing_table.RoutingTable
 method), 14

T

to_df() (neurokernel.routing_table.RoutingTable
 method), 14
 TryExceptionOnSignal() (in module neurokernel.ctx_managers), 25

W

Worker (class in neurokernel.mpi), 15
 WorkerManager (class in neurokernel.mpi), 15
 write() (in module neurokernel.neuroml.utils), 24

Z

ZMQOutput (in module neurokernel.tools.zmq), 28