

# MPI for Python and the new MPI ABI



جامعة الملك عبد الله  
للعلوم والتقنية  
King Abdullah University of  
Science and Technology

Lisandro Dalcin  
[dalcinl@gmail.com](mailto:dalcinl@gmail.com)  
EuroMPI/Australia 2024



# What is MPI for Python (**mpi4py**)?

- Unofficial Python bindings for MPI
- API based on MPI-2 C++ bindings (didn't age well)
- Almost all MPI features are supported
  - Best suited for MPI-4 implementations
  - Also works with MPI-1, MPI-2, MPI-3 implementations

# Features - MPI-1

- Point to point communication
  - blocking (send/recv)
  - nonblocking (isend/irecv + test/wait)
- Collective operations
  - Synchronization (barrier)
  - Communication (broadcast, scatter/gather)
  - Global reductions (reduce, scan)
- Process groups, communication domains, virtual topologies

# Features - MPI-2

- Parallel I/O (read/write)
- Dynamic process management (spawn, connect/accept)
- One-sided operations, a.k.a. RMA (put/get/accumulate)

# Features - MPI-3

- Matching probes (thread-safety)
- Non blocking collectives (ibARRIER, ibcast, igather)
- Neighbor collectives (neighbor allgather/alltoall)

# Features - MPI-4

- Persistent collectives
- Partitioned communication
- Sessions
- Large-count APIs, finally! 🚀

# Current MPI API (excluding MPI\_T)

	constants & routines
MPI-1	242
MPI-2	357
MPI-3	108
MPI-4	91 162
	960

# Hello World!

```
from mpi4py import MPI

rank = MPI.COMM_WORLD.Get_rank()
size = MPI.COMM_WORLD.Get_size()
name = MPI.Get_processor_name()

print(
    "Hello, World! I am process",
    f"{rank} of {size} on {name}"
)
```

```
$ mpiexec -n 5 python helloworld.py
```



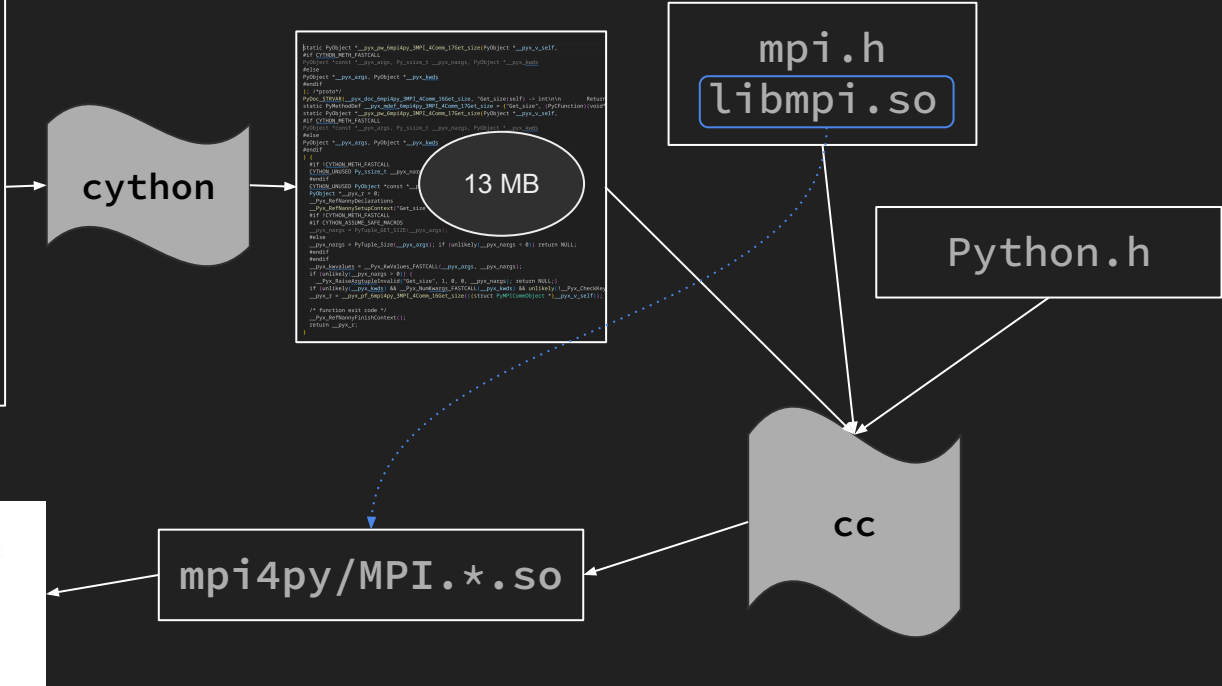



# From source to binary - pip install mpi4py

```
cdef extern from "<mpi.h>":  
    ctypedef ... MPI_Comm  
    int MPI_Comm_size(MPI_Comm, int*)  
    int MPI_Comm_rank(MPI_Comm, int*)  
  
cdef class Comm:  
  
    cdef MPI_Comm ob_mpi  
  
    def Get_size(self) -> int:  
        cdef int size  
        CHKERR( MPI_Comm_size(self.ob_mpi, &size) )  
        return size  
  
    def Get_rank(self) -> int:  
        cdef int rank  
        CHKERR( MPI_Comm_rank(self.ob_mpi, &rank) )  
        return rank
```



```
static PyObject *_pyc_mpiobj_mpi_comm_tor(PyObject *_pyc_self,  
    int argc, char **argv)  
{  
    PyObject *py_mpiobj_mpi_comm_tor = PyDict_GetItem(PyDict_Keys(self),  
        "mpiobj_mpi_comm_tor");  
    if (!py_mpiobj_mpi_comm_tor) Py_RETURN_NONE;  
    return PyDict_GetItem(PyDict_Keys(self), "mpiobj_mpi_comm_tor");  
}  
  
PyCDEF PyObject *_pyc_mpiobj_mpi_comm_tor(PyObject *_pyc_self,  
    int argc, char **argv)  
{  
    PyObject *py_mpiobj_mpi_comm_tor = PyDict_GetItem(PyDict_Keys(self),  
        "mpiobj_mpi_comm_tor");  
    if (!py_mpiobj_mpi_comm_tor) Py_RETURN_NONE;  
    return PyDict_GetItem(PyDict_Keys(self), "mpiobj_mpi_comm_tor");  
}  
  
PyCDEF PyObject *_pyc_mpiobj_mpi_comm_tor(PyObject *_pyc_self,  
    int argc, char **argv)  
{  
    PyObject *py_mpiobj_mpi_comm_tor = PyDict_GetItem(PyDict_Keys(self),  
        "mpiobj_mpi_comm_tor");  
    if (!py_mpiobj_mpi_comm_tor) Py_RETURN_NONE;  
    return PyDict_GetItem(PyDict_Keys(self), "mpiobj_mpi_comm_tor");  
}  
  
PyCDEF PyObject *_pyc_mpiobj_mpi_comm_tor(PyObject *_pyc_self,  
    int argc, char **argv)  
{  
    PyObject *py_mpiobj_mpi_comm_tor = PyDict_GetItem(PyDict_Keys(self),  
        "mpiobj_mpi_comm_tor");  
    if (!py_mpiobj_mpi_comm_tor) Py_RETURN_NONE;  
    return PyDict_GetItem(PyDict_Keys(self), "mpiobj_mpi_comm_tor");  
}  
  
PyCDEF PyObject *_pyc_mpiobj_mpi_comm_tor(PyObject *_pyc_self,  
    int argc, char **argv)  
{  
    PyObject *py_mpiobj_mpi_comm_tor = PyDict_GetItem(PyDict_Keys(self),  
        "mpiobj_mpi_comm_tor");  
    if (!py_mpiobj_mpi_comm_tor) Py_RETURN_NONE;  
    return PyDict_GetItem(PyDict_Keys(self), "mpiobj_mpi_comm_tor");  
}  
  
PyCDEF PyObject *_pyc_mpiobj_mpi_comm_tor(PyObject *_pyc_self,  
    int argc, char **argv)  
{  
    PyObject *py_mpiobj_mpi_comm_tor = PyDict_GetItem(PyDict_Keys(self),  
        "mpiobj_mpi_comm_tor");  
    if (!py_mpiobj_mpi_comm_tor) Py_RETURN_NONE;  
    return PyDict_GetItem(PyDict_Keys(self), "mpiobj_mpi_comm_tor");  
}  
  
PyCDEF PyObject *_pyc_mpiobj_mpi_comm_tor(PyObject *_pyc_self,  
    int argc, char **argv)  
{  
    PyObject *py_mpiobj_mpi_comm_tor = PyDict_GetItem(PyDict_Keys(self),  
        "mpiobj_mpi_comm_tor");  
    if (!py_mpiobj_mpi_comm_tor) Py_RETURN_NONE;  
    return PyDict_GetItem(PyDict_Keys(self), "mpiobj_mpi_comm_tor");  
}
```

13 MB



# Third-party pre-built mpi4py binaries

- macOS: Homebrew (Open MPI) and MacPorts (MPICH)
- Linux distributions: MPICH or Open MPI or both (via modules/alternatives)
- conda-forge <https://github.com/conda-forge/mpi4py-feedstock>

mpi4py-feedstock/recipe/conda\_build\_config.yaml

```
mpi:  
- mpich      # [not win]  
- openmpi    # [not win]  
- impi       # [not osx and x86_64]  
- msmapi     # [win]
```

65 builds

# Building binary Python wheels for distribution

<https://github.com/mpi4py/mpi4py-publish>

Python 3.6 - 3.13  
PyPy 3.7 - 3.10  
MPI 3.1 / 4.0 / 4.1

Linux	x86_64 aarch64 ppc64le	MPICH Open MPI (Intel MPI)
macOS	arm64 x86_64	MPICH Open MPI
Windows	AMD64	Intel MPI MSMPI



```
mpi4py-4.0.0-cp312-cp312-manylinux_2_17_x86_64.manylinux2014_x86_64.whl
```

```
mpi4py/__init__.py
```

```
...
```

```
mpi4py/MPI.mpi31-mpich.cpython-312-x86_64-linux-gnu.so
```

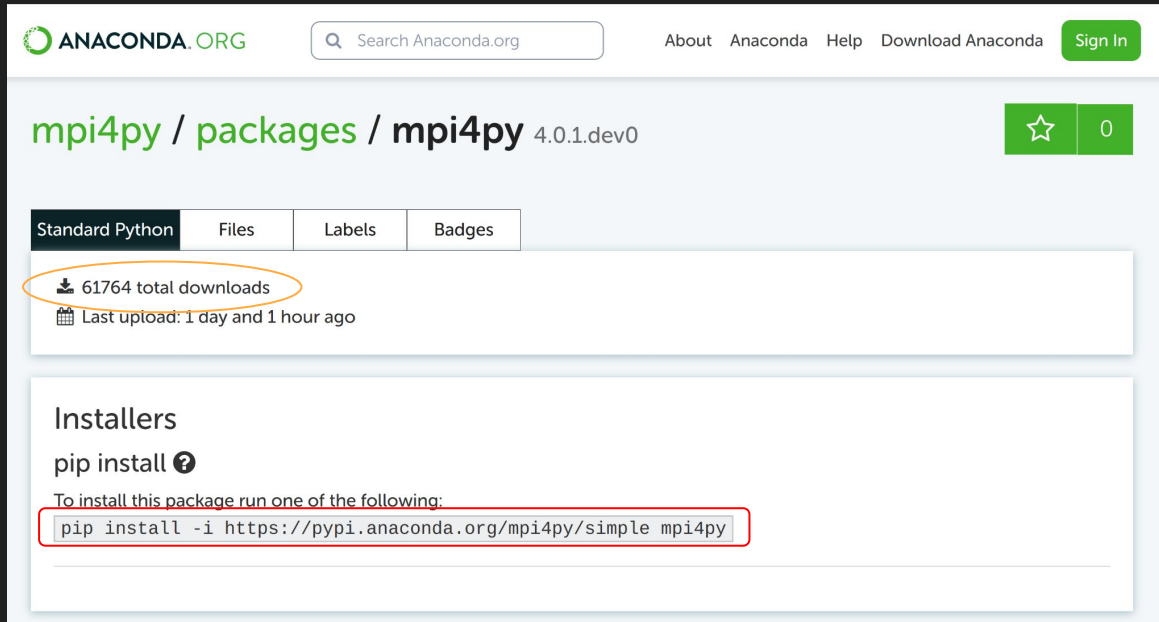
```
mpi4py/MPI.mpi40-mpich.cpython-312-x86_64-linux-gnu.so
```

```
mpi4py/MPI.mpi41-mpich.cpython-312-x86_64-linux-gnu.so
```

```
mpi4py/MPI.mpi31-openmpi.cpython-312-x86_64-linux-gnu.so
```

# Distributing binary Python wheels

<https://anaconda.org/mpi4py/mpi4py>



ANACONDA.ORG Search Anaconda.org About Anaconda Help Download Anaconda Sign In

mpi4py / packages / mpi4py 4.0.1.dev0 ☆ 0

Standard Python Files Labels Badges

61764 total downloads  
Last upload: 1 day and 1 hour ago



























### Installers

pip install ?

To install this package run one of the following:

```
pip install -i https://pypi.anaconda.org/mpi4py/simple mpi4py
```

# Is this useful? Does anyone cares?

<input type="checkbox"/>	Type	Size	Name	Uploaded	Downloads	Labels
<input type="checkbox"/>	Standard Python	1.5 MB	   mpi4py-3.1.6-cp38-cp38-manylinux_2_17_x86_64.manylinux2014_x86_64.whl	 5 months and 1 day ago	20584	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	1.5 MB	   mpi4py-3.1.6-cp311-cp311-manylinux_2_17_x86_64.manylinux2014_x86_64.whl	 5 months and 1 day ago	18548	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	1.5 MB	   mpi4py-3.1.5-cp38-cp38-manylinux_2_17_x86_64.manylinux2014_x86_64.whl	 6 months and 23 days ago	11976	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	3.3 MB	   mpi4py-4.0.0-cp38-cp38-manylinux_2_17_x86_64.manylinux2014_x86_64.whl	 1 month and 18 days ago	4179	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	3.3 MB	   mpi4py-4.0.0-cp312-cp312-manylinux_2_17_x86_64.manylinux2014_x86_64.whl	 1 month and 18 days ago	3609	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	3.3 MB	   mpi4py-4.0.0-cp311-cp311-manylinux_2_17_x86_64.manylinux2014_x86_64.whl	 1 month and 18 days ago	666	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	1.5 MB	   mpi4py-3.1.5-cp312-cp312-manylinux_2_17_x86_64.manylinux2014_x86_64.whl	 6 months and 23 days ago	462	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	1.3 MB	   mpi4py-3.1.5-cp312-cp312-macosx_11_0_arm64.whl	 6 months and 23 days ago	154	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	1.3 MB	   mpi4py-3.1.5-cp312-cp312-macosx_10_9_x86_64.whl	 6 months and 23 days ago	94	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	1.5 MB	   mpi4py-3.1.5-cp310-cp310-manylinux_2_17_x86_64.manylinux2014_x86_64.whl	 6 months and 23 days ago	94	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	1.3 MB	   mpi4py-3.1.5-cp310-cp310-macosx_11_0_arm64.whl	 6 months and 23 days ago	54	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	1.4 MB	   mpi4py-3.1.5-cp310-cp310-manylinux_2_17_aarch64.manylinux2014_aarch64.whl	 6 months and 23 days ago	40	<a href="#">main</a>
<input type="checkbox"/>	Standard Python	2.8 MB	   mpi4py-4.0.0-cp38-cp38-macosx_11_0_arm64.whl	 1 month and 18 days ago	37	<a href="#">main</a>

# Preliminary MPI ABI support

<https://github.com/mpi4py/mpi4py-testing/actions/workflows/abi.yml>

- Build step
  - Use MPI stubs ([https://github.com/mpiwg-abi/header\\_and\\_stub\\_library](https://github.com/mpiwg-abi/header_and_stub_library))
  - Generate Python wheel the usual way
- Test step
  - Build MPICH with `./configure --enable-mpi-abi`
  - Install Python wheel from build step
  - Run full mpi4py test suite

# A Future With (MPI ABI) Hope

- Use MPI API/ABI with weak symbols (multiple MPI std versions)
- Use `Py_LIMITED_API` (mpi4py would require Python  $\geq 3.11$ )

Just one Python wheel per OS & Arch!

```
mpi4py-4.0.0-cp311-cp311-manylinux_2_17_x86_64.manylinux2014_x86_64.whl
```

```
mpi4py-4.0.0-cp312-cp312-manylinux_2_17_x86_64.manylinux2014_x86_64.whl
```

```
mpi4py-4.0.0-cp313-cp313-manylinux_2_17_x86_64.manylinux2014_x86_64.whl
```

```
mpi4py/__init__.py
```

```
...
```

```
mpi4py/MPI.mpi31-mpich.cpython-313-x86_64-linux-gnu.so
```

```
mpi4py/MPI.mpi40-mpich.cpython-313-x86_64-linux-gnu.so
```

```
mpi4py/MPI.mpi41-mpich.cpython-313-x86_64-linux-gnu.so
```

```
mpi4py/MPI.mpi31-openmpi.cpython-313-x86_64-linux-gnu.so
```



```
mpi4py-4.0.1-cp311-abi3-manylinux_2_17_x86_64.whl
```

```
mpi4py/__init__.py
```

```
...
```

```
mpi4py/MPI.abi3.so
```



# Time Machine

If I could go back in time carrying the MPI ABI in my pocket...

- Get rid of dealing with C compilers and pre-built binaries
- Dynamically load the MPI library at runtime
- Call MPI routines via `ctypes` (built-in) or `cffi` (much nicer)
- Pure Python code working with any MPI ABI implementation
- Slightly slower in hotspots (usual Python overhead)



# Thanks!

<https://github.com/mmpi4py/mmpi4py>

<https://anaconda.org/mmpi4py/mmpi4py>

<https://github.com/mmpiwg-abi/abi-issues>

[https://github.com/mmpiwg-abi/header\\_and\\_stub\\_library](https://github.com/mmpiwg-abi/header_and_stub_library)