
PETSc for Python

Release 3.9.1

Lisandro Dalcin

May 02, 2018

Contents

1 Overview	2
1.1 Components	2
2 Installation	2
2.1 Using <code>pip</code> or <code>easy_install</code>	2
2.2 Using <code>distutils</code>	3
3 Tutorial	4
4 Citations	4
References	5
Index	6

Abstract

This document describes `petsc4py`, a Python port to the PETSc libraries.

PETSc (the Portable, Extensible Toolkit for Scientific Computation) is a suite of data structures and routines for the scalable (parallel) solution of scientific applications modeled by partial differential equations. It employs the MPI standard for all message-passing communication.

This package provides an important subset of PETSc functionalities and uses NumPy to efficiently manage input and output of array data.

A *good friend* of `petsc4py` is:

- `mpi4py`: Python bindings for MPI, the *Message Passing Interface*.

Other projects depends on `petsc4py`:

- `slepc4py`: Python bindings for SLEPc, the *Scalable Library for Eigenvalue Problem Computations*.

1 Overview

PETSc is a suite of data structures and routines for the scalable (parallel) solution of scientific applications modeled by partial differential equations. It employs the MPI standard for all message-passing communication.

PETSc is intended for use in large-scale application projects [*petsc-efficient*], and several ongoing computational science projects are built around the PETSc libraries. With strict attention to component interoperability, PETSc facilitates the integration of independently developed application modules, which often most naturally employ different coding styles and data structures.

PETSc is easy to use for beginners [*petsc-user-ref*]. Moreover, its careful design allows advanced users to have detailed control over the solution process. PETSc includes an expanding suite of parallel linear and nonlinear equation solvers that are easily used in application codes written in C, C++, and Fortran. PETSc provides many of the mechanisms needed within parallel application codes, such as simple parallel matrix and vector assembly routines that allow the overlap of communication and computation.

1.1 Components

PETSc is designed with an object-oriented style. Almost all user-visible types are abstract interfaces with implementations that may be chosen at runtime. Those objects are managed through handles to opaque data structures which are created, accessed and destroyed by calling appropriate library routines.

PETSc consists of a variety of components. Each component manipulates a particular family of objects and the operations one would like to perform on these objects. These components provide the functionality required for many parallel solutions of PDEs.

- Vec** Provides the vector operations required for setting up and solving large-scale linear and nonlinear problems. Includes easy-to-use parallel scatter and gather operations, as well as special-purpose code for handling ghost points for regular data structures.
- Mat** A large suite of data structures and code for the manipulation of parallel sparse matrices. Includes four different parallel matrix data structures, each appropriate for a different class of problems.
- PC** A collection of sequential and parallel preconditioners, including (sequential) ILU(k), LU, and (both sequential and parallel) block Jacobi, overlapping additive Schwarz methods and (through Block-Solve95) ILU(0) and ICC(0).
- KSP** Parallel implementations of many popular Krylov subspace iterative methods, including GMRES, CG, CGS, Bi-CG-Stab, two variants of TFQMR, CR, and LSQR. All are coded so that they are immediately usable with any preconditioners and any matrix data structures, including matrix-free methods.
- SNES** Data-structure-neutral implementations of Newton-like methods for nonlinear systems. Includes both line search and trust region techniques with a single interface. Employs by default the above data structures and linear solvers. Users can set custom monitoring routines, convergence criteria, etc.
- TS** Code for the time evolution of solutions of PDEs. In addition, provides pseudo-transient continuation techniques for computing steady-state solutions.

2 Installation

2.1 Using pip or easy_install

You can use **pip** to install `petsc4py` and its dependencies (`mpi4py` is optional but highly recommended):

```
$ pip install [--user] numpy mpi4py
$ pip install [--user] petsc petsc4py
```

Alternatively, you can use **easy_install** (deprecated):

```
$ easy_install petsc4py
```

If you already have a working PETSc installation, set environment variables `PETSC_DIR` and `PETSC_ARCH` to appropriate values and next use **pip**:

```
$ export PETSC_DIR=/path/to/petsc
$ export PETSC_ARCH=arch-linux2-c-opt
$ pip install petsc4py
```

2.2 Using distutils

Requirements

You need to have the following software properly installed in order to build *PETSc for Python*:

- Any MPI implementation¹ (e.g., **MPICH** or **Open MPI**), built with shared libraries.
- A matching version of **PETSc** built with shared libraries.
- **NumPy** package.

Downloading

The *PETSc for Python* package is available for download at the project website generously hosted by Bitbucket. You can use **curl** or **wget** to get a release tarball.

- Using **curl**:

```
$ curl -O https://bitbucket.org/petsc/petsc4py/petsc4py-X.Y.Z.tar.gz
```

- Using **wget**:

```
$ wget https://bitbucket.org/petsc/petsc4py/petsc4py-X.Y.Z.tar.gz
```

Building

After unpacking the release tarball:

```
$ tar -zxf petsc4py-X.Y.Z.tar.gz
$ cd petsc4py-X.Y.Z
```

the distribution is ready for building.

Note: **Mac OS X** users employing a Python distribution built with **universal binaries** may need to set the environment variables `MACOSX_DEPLOYMENT_TARGET`, `SDKROOT`, and `ARCHFLAGS` to appropriate values. As an example, assume your Mac is running **Snow Leopard** on a **64-bit Intel** processor and you want to override the hard-wired cross-development SDK in Python configuration, your environment should be modified like this:

¹ Unless you have appropriately configured and built PETSc without MPI (configure option `--with-mpi=0`).

```
$ export MACOSX_DEPLOYMENT_TARGET=10.6
$ export SDKROOT=/
$ export ARCHFLAGS='-arch x86_64'
```

Some environment configuration is needed to inform the location of PETSc. You can set (using **setenv**, **export** or what applies to you shell or system) the environment variables `PETSC_DIR`, and `PETSC_ARCH` indicating where you have built/installed PETSc:

```
$ export PETSC_DIR=/usr/local/petsc
$ export PETSC_ARCH=arch-linux2-c-opt
```

Alternatively, you can edit the file `setup.cfg` and provide the required information below the `[config]` section:

```
[config]
petsc_dir = /usr/local/petsc
petsc_arch = arch-linux2-c-opt
...
```

Finally, you can build the distribution by typing:

```
$ python setup.py build
```

Installing

After building, the distribution is ready for installation.

If you have root privileges (either by log-in as the root user or by using **sudo**) and you want to install *PETSc for Python* in your system for all users, just do:

```
$ python setup.py install
```

The previous steps will install the `petsc4py` package at standard location `prefix/lib/pythonX.Y/site-packages`.

If you do not have root privileges or you want to install *PETSc for Python* for your private use, just do:

```
$ python setup.py install --user
```

3 Tutorial

XXX To be written ... Any contribution welcome!

4 Citations

If PETSc for Python been significant to a project that leads to an academic publication, please acknowledge that fact by citing the project.

- L. Dalcin, P. Kler, R. Paz, and A. Cosimo, *Parallel Distributed Computing using Python*, Advances in Water Resources, 34(9):1124-1139, 2011. <http://dx.doi.org/10.1016/j.advwatres.2011.04.013>

- S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. Gropp, D. Karpeyev, D. Kaushik, M. Knepley, D. May, L. Curfman McInnes, R. Mills, T. Munson, K. Rupp, P. Sanan, B. Smith, S. Zampini, H. Zhang, and H. Zhang, *PETSc Users Manual*, ANL-95/11 - Revision 3.9, 2018. <http://www.mcs.anl.gov/petsc/petsc-current/docs/manual.pdf>

References

- [petsc-user-ref] S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. Gropp, D. Karpeyev, D. Kaushik, M. Knepley, D. May, L. Curfman McInnes, R. Mills, T. Munson, K. Rupp, P. Sanan, B. Smith, S. Zampini, H. Zhang, and H. Zhang, *PETSc Users Manual*, ANL-95/11 - Revision 3.9, 2018. <http://www.mcs.anl.gov/petsc/petsc-current/docs/manual.pdf>
- [petsc-efficient] Satish Balay, Victor Eijkhout, William D. Gropp, Lois Curfman McInnes and Barry F. Smith. Efficient Management of Parallelism in Object Oriented Numerical Software Libraries. *Modern Software Tools in Scientific Computing*. E. Arge, A. M. Bruaset and H. P. Langtangen, editors. 163–202. Birkhauser Press. 1997.

Index

A

ARCHFLAGS, 3

E

environment variable

ARCHFLAGS, 3

MACOSX_DEPLOYMENT_TARGET, 3

PETSC_ARCH, 3, 4

PETSC_DIR, 3, 4

SDKROOT, 3

M

MACOSX_DEPLOYMENT_TARGET, 3

P

PETSC_ARCH, 3, 4

PETSC_DIR, 3, 4

S

SDKROOT, 3