

---

# hypr Documentatation

*Release 3.1.0*

**Lawrence Livermore National Laboratory**

**January 23, 2026**



# CONTENTS

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Overview of Features . . . . .	3
1.2	Getting More Information . . . . .	4
1.3	How to get started . . . . .	4
1.3.1	Installing hypre . . . . .	4
1.3.2	Choosing a conceptual interface . . . . .	5
1.3.3	Writing your code . . . . .	6
<b>2</b>	<b>Structured-Grid System Interface (Struct)</b>	<b>7</b>
2.1	Setting Up the Struct Grid . . . . .	7
2.2	Setting Up the Struct Stencil . . . . .	8
2.3	Setting Up the Struct Matrix . . . . .	9
2.4	Setting Up the Struct Right-Hand-Side Vector . . . . .	10
2.5	Symmetric Matrices . . . . .	11
<b>3</b>	<b>Semi-Structured-Grid System Interface (SStruct)</b>	<b>13</b>
3.1	Block-Structured Grids with Stencils . . . . .	13
3.2	Block-Structured Grids with Finite Elements . . . . .	16
3.3	Structured Adaptive Mesh Refinement . . . . .	18
<b>4</b>	<b>Linear-Algebraic System Interface (IJ)</b>	<b>21</b>
4.1	IJ Matrix Interface . . . . .	21
4.2	IJ Vector Interface . . . . .	22
4.3	A Scalable Interface . . . . .	23
<b>5</b>	<b>Solvers and Preconditioners</b>	<b>25</b>
5.1	SMG . . . . .	27
5.2	PFMG . . . . .	27
5.3	SysPFMG . . . . .	27
5.4	SplitSolve . . . . .	27
5.5	SSAMG . . . . .	28
5.5.1	Configuration options . . . . .	28
5.5.2	Algorithm notes . . . . .	29
5.5.3	Minimal C example . . . . .	29
5.6	Hybrid . . . . .	30
5.7	BoomerAMG . . . . .	31
5.7.1	Parameter Options . . . . .	31
5.7.2	Coarsening Options . . . . .	31
5.7.3	Interpolation Options . . . . .	31
5.7.4	Non-Galerkin Options . . . . .	32

5.7.5	Smoother Options	32
5.7.6	AMG for systems of PDEs	33
5.7.7	Special AMG Cycles	33
5.7.8	GPU-supported Options	33
5.7.9	Memory locations and execution policies	33
5.7.10	Miscellaneous	35
5.8	AMS	35
5.8.1	Overview	35
5.8.2	Sample Usage	36
5.8.3	High-order Discretizations	39
5.8.4	Non-conforming AMR Grids	39
5.9	ADS	40
5.9.1	Overview	40
5.9.2	Sample Usage	41
5.9.3	High-order Discretizations	43
5.10	Multigrid Reduction (MGR)	43
5.11	FSAI	44
5.11.1	Parameter Settings	44
5.11.2	FSAI as Smoother to BoomerAMG	45
5.11.3	Implementation Notes	45
5.12	ParaSails	46
5.12.1	Parameter Settings	46
5.12.2	Preconditioning Nearly Symmetric Matrices	47
5.13	ILU	47
5.13.1	Overview	47
5.13.2	User-level functions	48
5.13.3	ILU as Smoother for BoomerAMG	50
5.13.4	GPU support	50
5.14	Euclid	51
5.14.1	Overview	52
5.14.2	Setting Options: Examples	53
5.14.3	Options Summary	53
5.15	PILUT: Parallel Incomplete Factorization	54
5.15.1	Parameters:	55
5.16	LOBPCG Eigensolver	55
<b>6</b>	<b>Mixed Precision</b>	<b>57</b>
6.1	Calling functions with fixed precision	57
6.2	Calling functions with multiple precisions	57
<b>7</b>	<b>General Information</b>	<b>59</b>
7.1	Getting the Source Code	59
7.2	Building the Library	59
7.2.1	1. Using autotools (Configure & Make)	60
7.2.2	2. Using CMake (Windows, macOS, Linux, etc.)	60
7.2.3	3. Using Spack (Recommended for HPC environments)	61
7.3	Build System Options	62
7.4	GPU Build Options	64
7.4.1	Building Umpire	66
7.5	Make Targets	67
7.6	Using the Library	68
7.7	Testing the Library	69
7.8	Linking to the Library	70
7.8.1	Using CMake	70

7.8.2	Using Autotools . . . . .	70
7.8.3	Shared Library Considerations . . . . .	70
7.9	Error Flags . . . . .	71
7.10	Bug Reporting and General Support . . . . .	72
7.11	Calling HYPRE from Other Languages . . . . .	73
<b>8</b>	<b>API</b>	<b>75</b>
8.1	Struct System and Object Interface . . . . .	75
8.2	SStruct System and Object Interface . . . . .	82
8.3	IJ System Interface . . . . .	99
8.4	ParCSR Object Interface . . . . .	107
8.5	Struct Solvers . . . . .	109
8.6	SStruct Solvers . . . . .	121
8.7	ParCSR Solvers . . . . .	131
8.8	Krylov Solvers . . . . .	201
8.9	Eigensolvers . . . . .	211
8.10	Utilities . . . . .	212
	<b>Bibliography</b>	<b>221</b>
	<b>Index</b>	<b>225</b>



Copyright (c) 1998 Lawrence Livermore National Security, LLC and other HYPRE Project Developers. See the top-level COPYRIGHT file for details.

SPDX-License-Identifier: (Apache-2.0 OR MIT)



## INTRODUCTION

This manual describes hypre, a software library of high performance preconditioners and solvers for the solution of large, sparse linear systems of equations on massively parallel computers [FaJY2004]. The hypre library was created with the primary goal of providing users with advanced parallel preconditioners. The library features parallel multigrid solvers for both structured and unstructured grid problems. For ease of use, these solvers are accessed from the application code via hypre’s conceptual linear system interfaces [FaJY2005] (abbreviated to *conceptual interfaces* throughout much of this manual), which allow a variety of natural problem descriptions.

This introductory chapter provides an overview of the various features in hypre, discusses further sources of information on hypre, and offers suggestions on how to get started.

### 1.1 Overview of Features

**Scalable preconditioners provide efficient solution on today’s and tomorrow’s systems:** hypre contains several families of preconditioner algorithms focused on the scalable solution of *very large* sparse linear systems. (Note that small linear systems, systems that are solvable on a sequential computer, and dense systems are all better addressed by other libraries that are designed specifically for them.) hypre includes “grey box” algorithms that use more than just the matrix to solve certain classes of problems more efficiently than general-purpose libraries. This includes algorithms such as structured multigrid.

**Suite of common iterative methods provides options for a spectrum of problems:** hypre provides several of the most commonly used Krylov-based iterative methods to be used in conjunction with its scalable preconditioners. This includes methods for nonsymmetric systems such as GMRES and methods for symmetric matrices such as Conjugate Gradient.

**Intuitive grid-centric interfaces obviate need for complicated data structures and provide access to advanced solvers:** hypre has made a major step forward in usability from earlier generations of sparse linear solver libraries in that users do not have to learn complicated sparse matrix data structures. Instead, hypre does the work of building these data structures for the user through a variety of conceptual interfaces, each appropriate to different classes of users. These include stencil-based structured/semi-structured interfaces most appropriate for finite-difference applications; a finite-element based unstructured interface; and a linear-algebra based interface. Each conceptual interface provides access to several solvers without the need to write new interface code.

**User options accommodate beginners through experts:** hypre allows a spectrum of expertise to be applied by users. The beginning user can get up and running with a minimal amount of effort. More expert users can take further control of the solution process through various parameters.

**Configuration options to suit your computing system:** hypre allows a simple and flexible installation on a wide variety of computing systems. Users can tailor the installation to match their computing system. Options include debug and optimized modes, the ability to change required libraries such as MPI and BLAS, a sequential mode, and modes enabling threads for certain solvers. On most systems, however, hypre can be built by simply typing `configure` followed by `make`, or by using CMake [CMakeWeb].

**Interfaces in multiple languages provide greater flexibility for applications:** hypr is written in C and provides an interface for Fortran users.

## 1.2 Getting More Information

This user's manual consists of chapters describing each conceptual interface, a chapter detailing the various linear solver options available, detailed installation information, and the API reference. In addition to this manual, a number of other information sources for hypr are available.

- **Reference Manual:** This is equivalent to Chapter *API* in this user manual, but it can also be built as a separate document. The reference manual comprehensively lists all of the interface and solver functions available in hypr. It is ideal for determining the various options available for a particular solver or for viewing the functions provided to describe a problem for a particular interface.
- **Example Problems:** A suite of example problems is provided with the hypr installation. These examples reside in the `examples` subdirectory and demonstrate various features of the hypr library. Associated documentation may be accessed by viewing the `README.html` file in that same directory.
- **hyprdrive:** Users who want runtime-configurable solver setups may benefit from `hyprdrive`, a high-level interface library built on top of hypr. It provides both a standalone driver for matrix/vector files and a lightweight C API that accepts YAML input from file or in-memory string. This is useful for prototyping, benchmarking, and fine-tuning hypr solver and preconditioner configurations without recompiling application code. See [hyprdrive documentation](#) for more details.
- **Papers, Presentations, etc.:** Articles and presentations related to the hypr software library and the solvers available in the library are available from the hypr web page at <http://www.llnl.gov/CASC/hypr/>.
- **Mailing List:** The mailing list `hypr-announce` can be subscribed to through the hypr web page at <http://www.llnl.gov/CASC/hypr/>. The development team uses this list to announce new releases of hypr. It cannot be posted to by users.

## 1.3 How to get started

### 1.3.1 Installing hypr

As previously noted, on most systems hypr can be built by typing `./configure` followed by `make` in the top-level source directory. Alternatively, `CMake` can be used and is the recommended method for building hypr on Windows. For more detailed instructions, refer to the `INSTALL` file provided in the hypr distribution or see the *General Information* section of this manual.

**Note**

- To run in parallel, hyre requires an installation of MPI.
- Configuration of hyre with threads requires an implementation of OpenMP. Currently, only a subset of hyre is threaded.
- To run on GPUs, hyre must be configured with support for the appropriate GPU toolkit. Available options are CUDA for NVIDIA GPUs, HIP for AMD GPUs, and SYCL for Intel GPUs.
- The hyre library currently does not directly support complex-valued systems.

### 1.3.2 Choosing a conceptual interface

An important decision to make before writing any code is to choose an appropriate conceptual interface. These conceptual interfaces are intended to represent the way that applications developers naturally think of their linear problem and to provide natural interfaces for them to pass the data that defines their linear system into hyre. Essentially, these conceptual interfaces can be considered convenient utilities for helping a user build a matrix data structure for hyre solvers and preconditioners. The top row of [Figure 1.1](#) illustrates a number of conceptual interfaces. Generally, the conceptual interfaces are denoted by different types of computational grids, but other application features might also be used, such as geometrical information. For example, applications that use structured grids (such as in the left-most interface in [Figure 1.1](#)) typically view their linear problems in terms of stencils and grids. On the other hand, applications that use unstructured grids and finite elements typically view their linear problems in terms of elements and element stiffness matrices. Finally, the right-most interface is the standard linear-algebraic (matrix rows/columns) way of viewing the linear problem.

The hyre library currently supports three conceptual interfaces, and typically the appropriate choice for a given problem is fairly obvious, e.g. a structured-grid interface is clearly inappropriate for an unstructured-grid application.

- **Structured-Grid System Interface (Struct):** This interface is appropriate for applications whose grids consist of unions of logically rectangular grids with a fixed stencil pattern of nonzeros at each grid point. This interface supports only a single unknown per grid point. See Chapter [Structured-Grid System Interface \(Struct\)](#) for details.
- **Semi-Structured-Grid System Interface (SStruct):** This interface is appropriate for applications whose grids are mostly structured, but with some unstructured features. Examples include block-structured grids, composite grids in structured adaptive mesh refinement (AMR) applications, and overset grids. This interface supports multiple unknowns per cell. See Chapter [Semi-Structured-Grid System Interface \(SStruct\)](#) for details.
- **Linear-Algebraic System Interface (IJ):** This is the traditional linear-algebraic interface. It can be used as a last resort by users for whom the other grid-based interfaces are not appropriate. It requires more work on the user's part, though still less than building parallel sparse data structures. General solvers and preconditioners are available through this interface, but not specialized solvers which need more information. Our experience is that users with legacy codes, in which they already have code for building matrices in particular formats, find the IJ interface relatively easy to use. See Chapter [Linear-Algebraic System Interface \(IJ\)](#) for details.

Figure 1.1: : Graphic illustrating the notion of conceptual linear system interfaces.

Generally, a user should choose the most specific interface that matches their application, because this will allow them to use specialized and more efficient solvers and preconditioners without losing access to more general solvers. For example, the second row of [Figure 1.1](#) is a set of linear solver algorithms. Each linear solver group requires different information from the user through the conceptual interfaces. So, the geometric multigrid algorithm (GMG) listed in the left-most box, for example, can only be used with the left-most conceptual interface. On the other hand, the ILU algorithm in the right-most box may be used with any conceptual interface. Matrix requirements for each solver and preconditioner are provided in Chapter [Solvers and Preconditioners](#) and in Chapter [API](#). Your desired solver strategy may influence your choice of conceptual interface. A typical user will select a single Krylov method and a single preconditioner to solve their system.

The third row of [Figure 1.1](#) is a list of data layouts or matrix/vector storage schemes. The relationship between linear solver and storage scheme is similar to that of the conceptual interface and linear solver. Note that some of the interfaces in `hypr` currently only support one matrix/vector storage scheme choice. The conceptual interface, the desired solvers and preconditioners, and the matrix storage class must all be compatible.

### 1.3.3 Writing your code

As discussed in the previous section, the following decisions should be made before writing any code:

- Choose a conceptual interface.
- Choose your desired solver strategy.
- Look up matrix requirements for each solver and preconditioner.
- Choose a matrix storage class that is compatible with your solvers and preconditioners and your conceptual interface.

Once the previous decisions have been made, it is time to code your application to call `hypr`. At this point, reviewing the previously mentioned example codes provided with the `hypr` library may prove very helpful. The example codes demonstrate the following general structure of the application calls to `hypr`:

- **Build any necessary auxiliary structures for your chosen conceptual interface.** This includes, e.g., the grid and stencil structures if you are using the structured-grid interface.
- **Build the matrix, solution vector, and right-hand-side vector through your chosen conceptual interface.** Each conceptual interface provides a series of calls for entering information about your problem into `hypr`.
- **Build solvers and preconditioners and set solver parameters (optional).** Some parameters like convergence tolerance are the same across solvers, while others are solver specific.
- **Call the solve function for the solver.**
- **Retrieve desired information from solver.** Depending on your application, there may be different things you may want to do with the solution vector. Also, performance information such as number of iterations is typically available, though it may differ from solver to solver.

The subsequent chapters of this User's Manual provide the details needed to more fully understand the function of each conceptual interface and each solver. Remember that a comprehensive list of all available functions is provided in Chapter [API](#), and the provided example codes may prove helpful as templates for your specific application.

## STRUCTURED-GRID SYSTEM INTERFACE (STRUCT)

In order to get access to the most efficient and scalable solvers for scalar structured-grid applications, users should use the `Struct` interface described in this chapter. This interface will also provide access (this is not yet supported) to solvers in `hypr` that were designed for unstructured-grid applications and sparse linear systems in general. These additional solvers are usually provided via the linear-algebraic interface (IJ) described in Chapter *Linear-Algebraic System Interface (IJ)*.

Figure 2.1 gives an example of the type of grid currently supported by the `Struct` interface. The interface uses a finite-difference or finite-volume style, and currently supports only scalar PDEs (i.e., one unknown per gridpoint).

Figure 2.1: : 2D structured grid example that is distributed across two processors.

There are four basic steps involved in setting up the linear system to be solved:

1. set up the grid,
2. set up the stencil,
3. set up the matrix,
4. set up the right-hand-side vector.

To describe each of these steps in more detail, consider solving the 2D Laplacian problem

$$\begin{cases} \nabla^2 u = f, & \text{in the domain,} \\ u = 0, & \text{on the boundary.} \end{cases} \quad (2.1)$$

Assume (2.1) is discretized using standard 5-pt finite-volumes on the uniform grid pictured in Figure 2.1, and assume that the problem data is distributed across two processes as depicted.

### 2.1 Setting Up the Struct Grid

The grid is described via a global *index space*, i.e., via integer singles in 1D, tuples in 2D, or triples in 3D (see Figure 2.2).

Figure 2.2: : Boxes in index space. A box is a collection of abstract cell-centered indices, described by its minimum and maximum indices. Here, two boxes are illustrated.

The integers may have any value, negative or positive. The global indexes allow `hypr` to discern how data is related spatially, and how it is distributed across the parallel machine. The basic component of the grid is a *box*: a collection of abstract cell-centered indices in index space, described by its “lower” and “upper” corner indices. The scalar grid data is always associated with cell centers, unlike the more general `SSTRUCT` interface which allows data to be associated with box indices in several different ways.

Each process describes that portion of the grid that it “owns”, one box at a time. For example, the global grid in [Figure 2.1](#) can be described in terms of three boxes, two owned by process 0, and one owned by process 1. The following is the code (with visual annotations) for setting up the grid on process 0 (the code for process 1 is similar).

1:	2:	3:	4:
----	----	----	----

```

HYPRE_StructGrid grid;
int ndim          = 2;
int ilower[][2]  = {{-3,1}, {0,1}};
int iupper[][2]  = {{-1,2}, {2,4}};

/* Create the grid object */
1: HYPRE_StructGridCreate(MPI_COMM_WORLD, ndim, &grid);

/* Set grid extents for the first box */
2: HYPRE_StructGridSetExtents(grid, ilower[0], iupper[0]);

/* Set grid extents for the second box */
3: HYPRE_StructGridSetExtents(grid, ilower[1], iupper[1]);

/* Assemble the grid */
4: HYPRE_StructGridAssemble(grid);

```

The images along the top illustrate the result of the numbered lines of code. The `Create()` routine creates an empty 2D grid object that lives on the `MPI_COMM_WORLD` communicator. The `SetExtents()` routine adds a new box to the grid. The `Assemble()` routine is a collective call (i.e., must be called on all processes from a common synchronization point), and finalizes the grid assembly, making the grid “ready to use”.

## 2.2 Setting Up the Struct Stencil

The geometry of the discretization stencil is described by an array of indexes, each representing a relative offset from any given gridpoint on the grid. For example, the geometry of the 5-pt stencil for the example problem being considered can be represented by the list of index offsets shown in [Figure 2.3](#).

Figure 2.3: : Representation of the 5-point discretization stencil for the example problem.

Figure 2.4: : Alternate representation of the stencil configuration shown in [Figure 2.3](#).

Here, the  $(0, 0)$  entry represents the “center” coefficient, and is the 0th stencil entry. The  $(0, -1)$  entry represents the “south” coefficient, and is the 3rd stencil entry. And so on.

On process 0 or 1, the following code (with visual annotations) will set up the stencil in [Figure 2.3](#). The stencil must be the same on all processes.

1:	2:	3:
4:	5:	6:

```

HYPRE_StructStencil stencil;
int ndim          = 2;

```

(continues on next page)

(continued from previous page)

```

int size      = 5;
int entry;
int offsets[][2] = {{0,0}, {-1,0}, {1,0}, {0,-1}, {0,1}};

/* Create the stencil object */
1: HYPRE_StructStencilCreate(ndim, size, &stencil);

/* Set stencil entries */
for (entry = 0; entry < size; entry++)
{
2-6:   HYPRE_StructStencilSetElement(stencil, entry, offsets[entry]);
}

/* Thats it! There is no assemble routine */

```

The Create() routine creates an empty 2D, 5-pt stencil object. The SetElement() routine defines the geometry of the stencil and assigns the stencil numbers for each of the stencil entries. None of the calls are collective calls.

## 2.3 Setting Up the Struct Matrix

The matrix is set up in terms of the grid and stencil objects described in Sections *Setting Up the Struct Grid* and *Setting Up the Struct Stencil*. The coefficients associated with each stencil entry will typically vary from gridpoint to gridpoint, but in the example problem being considered, they are as follows over the entire grid (except at boundaries; see below):

$$\begin{bmatrix} & -1 & \\ -1 & 4 & -1 \\ & -1 & \end{bmatrix}. \quad (2.2)$$

On process 0, the following code sets up matrix values associated with the center (entry 0) and south (entry 3) stencil entries as given by (2.2) and Figure 2.3 (boundaries are ignored here temporarily).

```

HYPRE_StructMatrix A;
double values[36];
int stencil_indices[2] = {0,3};
int i;

HYPRE_StructMatrixCreate(MPI_COMM_WORLD, grid, stencil, &A);
HYPRE_StructMatrixInitialize(A);

for (i = 0; i < 36; i += 2)
{
  values[i] = 4.0;
  values[i+1] = -1.0;
}

HYPRE_StructMatrixSetBoxValues(A, ilower[0], iupper[0], 2,
                               stencil_indices, values);
HYPRE_StructMatrixSetBoxValues(A, ilower[1], iupper[1], 2,
                               stencil_indices, values);

/* set boundary conditions */
...

```

(continues on next page)

```
HYPRE_StructMatrixAssemble(A);
```

The `Create()` routine creates an empty matrix object. The `Initialize()` routine indicates that the matrix coefficients (or values) are ready to be set. This routine may or may not involve the allocation of memory for the coefficient data, depending on the implementation. The optional `Set` routines mentioned later in this chapter and in Chapter *API*, should be called before this step. The `SetBoxValues()` routine sets the matrix coefficients for some set of stencil entries over the gridpoints in some box. Note that the box need not correspond to any of the boxes used to create the grid, but values should be set for all gridpoints that this process “owns”. The `Assemble()` routine is a collective call, and finalizes the matrix assembly, making the matrix “ready to use”.

Matrix coefficients that reach outside of the boundary should be set to zero. For efficiency reasons, hypr does not do this automatically. The most natural time to insure this is when the boundary conditions are being set, and this is most naturally done after the coefficients on the grid’s interior have been set. For example, during the implementation of the Dirichlet boundary condition on the lower boundary of the grid in [Figure 2.1](#), the south coefficient must be set to zero. To do this on process 0, the following code could be used:

```
int  ilower[2] = {-3, 1};
int  iupper[2] = { 2, 1};

/* create matrix and set interior coefficients */
...

/* implement boundary conditions */
...

for (i = 0; i < 12; i++)
{
    values[i] = 0.0;
}

i = 3;
HYPRE_StructMatrixSetBoxValues(A, ilower, iupper, 1, &i, values);

/* complete implementation of boundary conditions */
...
```

## 2.4 Setting Up the Struct Right-Hand-Side Vector

The right-hand-side vector is set up similarly to the matrix set up described in Section *Setting Up the Struct Matrix* above. The main difference is that there is no stencil (note that a stencil currently does appear in the interface, but this will eventually be removed).

On process 0, the following code sets up the right-hand-side vector values.

```
HYPRE_StructVector  b;
double             values[18];
int                i;

HYPRE_StructVectorCreate(MPI_COMM_WORLD, grid, &b);
HYPRE_StructVectorInitialize(b);
```

(continues on next page)

(continued from previous page)

```

for (i = 0; i < 18; i++)
{
    values[i] = 0.0;
}

HYPRE_StructVectorSetBoxValues(b, ilower[0], iupper[0], values);
HYPRE_StructVectorSetBoxValues(b, ilower[1], iupper[1], values);

HYPRE_StructVectorAssemble(b);

```

The `Create()` routine creates an empty vector object. The `Initialize()` routine indicates that the vector coefficients (or values) are ready to be set. This routine follows the same rules as its corresponding `Matrix` routine. The `SetBoxValues()` routine sets the vector coefficients over the gridpoints in some box, and again, follows the same rules as its corresponding `Matrix` routine. The `Assemble()` routine is a collective call, and finalizes the vector assembly, making the vector “ready to use”.

## 2.5 Symmetric Matrices

Some solvers and matrix storage schemes provide capabilities for significantly reducing memory usage when the coefficient matrix is symmetric. In this situation, each off-diagonal coefficient appears twice in the matrix, but only one copy needs to be stored. The `Struct` interface provides support for matrix and solver implementations that use symmetric storage via the `SetSymmetric()` routine.

To describe this in more detail, consider again the 5-pt finite-volume discretization of (2.1) on the grid pictured in [Figure 2.1](#). Because the discretization is symmetric, only half of the off-diagonal coefficients need to be stored. To turn symmetric storage on, the following line of code needs to be inserted somewhere between the `Create()` and `Initialize()` calls.

```
HYPRE_StructMatrixSetSymmetric(A, 1);
```

The coefficients for the entire stencil can be passed in as before. Note that symmetric storage may or may not actually be used, depending on the underlying storage scheme. Currently in `hyre`, the `Struct` interface always uses symmetric storage.

To most efficiently utilize the `Struct` interface for symmetric matrices, notice that only half of the off-diagonal coefficients need to be set. To do this for the example being considered, we simply need to redefine the 5-pt stencil of [Section \*Setting Up the Struct Stencil\*](#) to an “appropriate” 3-pt stencil, then set matrix coefficients (as in [Section \*Setting Up the Struct Matrix\*](#)) for these three stencil elements *only*. For example, we could use the following stencil

$$\begin{bmatrix} (0, 1) \\ (0, 0) & (1, 0) \end{bmatrix}. \quad (2.3)$$

This 3-pt stencil provides enough information to recover the full 5-pt stencil geometry and associated matrix coefficients.



## SEMI-STRUCTURED-GRID SYSTEM INTERFACE (SSTRUCT)

The `SStruct` interface is appropriate for applications with grids that are mostly—but not entirely—structured, e.g. block-structured grids (see [Figure 3.2](#)), composite grids in structured adaptive mesh refinement (AMR) applications (see [Figure 3.11](#)), and overset grids. In addition, it supports more general PDEs than the `Struct` interface by allowing multiple variables (system PDEs) and multiple variable types (e.g. cell-centered, face-centered, etc.). The interface provides access to data structures and linear solvers in `hypr` that are designed for semi-structured grid problems, but also to the most general data structures and solvers.

The `SStruct` grid is composed out of a number of structured grid *parts*, where the physical inter-relationship between the parts is arbitrary. Each part is constructed out of two basic components: boxes (see [Figure 2.2](#)) and *variables*. Variables represent the actual unknown quantities in the grid, and are associated with the box indices in a variety of ways, depending on their types. In `hypr`, variables may be cell-centered, node-centered, face-centered, or edge-centered. Face-centered variables are split into x-face, y-face, and z-face, and edge-centered variables are split into x-edge, y-edge, and z-edge. See [Figure 3.1](#) for an illustration in 2D.

Figure 3.1: : Grid variables in `hypr` are referenced by the abstract cell-centered index to the left and down in 2D (analogously in 3D). In the figure, index  $(i, j)$  is used to reference the variables in black. The variables in grey—although contained in the pictured cell—are not referenced by the  $(i, j)$  index.

The `SStruct` interface uses a *graph* to allow nearly arbitrary relationships between part data. The graph is constructed from stencils or finite element stiffness matrices plus some additional data-coupling information set by the `GraphAddEntries()` routine. Two other methods for relating part data are the `GridSetNeighborPart()` and `GridSetSharedPart()` routines, which are particularly well suited for block-structured grid problems. The latter is useful for finite element codes.

There are five basic steps involved in setting up the linear system to be solved:

1. set up the grid,
2. set up the stencils (if needed),
3. set up the graph,
4. set up the matrix,
5. set up the right-hand-side vector.

### 3.1 Block-Structured Grids with Stencils

In this section, we describe how to use the `SStruct` interface to define block-structured grid problems. We do this primarily by example, paying particular attention to the construction of stencils and the use of the `GridSetNeighborPart()` interface routine.

Consider the solution of the diffusion equation

$$-\nabla \cdot (D\nabla u) + \sigma u = f \tag{3.1}$$

on the block-structured grid in Figure 3.2, where  $D$  is a scalar diffusion coefficient, and  $\sigma \geq 0$ . The discretization [MoRS1998] introduces three different types of variables: cell-centered,  $x$ -face, and  $y$ -face. The three discretization stencils that couple these variables are given Figure 3.3. The information in this figure is essentially all that is needed to describe the nonzero structure of the linear system we wish to solve.

Figure 3.2: : Example of a block-structured grid with five logically-rectangular blocks and three variables types: cell-centered,  $x$ -face, and  $y$ -face. Discretization stencils for the cell-centered (left),  $x$ -face (middle), and  $y$ -face (right) variables are also pictured.

Figure 3.3: : Example of stencil connections involving different variable types.

Figure 3.4: : One possible labeling of the grid in Figure 3.2.

The grid in Figure 3.2 is defined in terms of five separate logically-rectangular parts as shown in Figure 3.4, and each part is given a unique label between 0 and 4. Each part consists of a single box with lower index (1, 1) and upper index (4, 4) (see Section *Setting Up the Struct Grid*), and the grid data is distributed on five processes such that data associated with part  $p$  lives on process  $p$ . Note that in general, parts may be composed out of arbitrary unions of boxes, and indices may consist of non-positive integers (see Figure 2.2). Also note that the SStruct interface expects a domain-based data distribution by boxes, but the actual distribution is determined by the user and simply described (in parallel) through the interface.

1:	2:	3:
4:	5:	6:

```
HYPRE_SStructGrid grid;
int ndim = 2, nparts = 5, nvars = 3, part = 3;
int extents[][2] = {{1,1}, {4,4}};
int vartypes[] = {HYPRE_SSTRUCT_VARIABLE_CELL,
                  HYPRE_SSTRUCT_VARIABLE_XFACE,
                  HYPRE_SSTRUCT_VARIABLE_YFACE};

int nb2_n_part      = 2,          nb4_n_part      = 4;
int nb2_exts[][2]  = {{1,0}, {4,0}}, nb4_exts[][2] = {{0,1}, {0,4}};
int nb2_n_exts[][2] = {{1,1}, {1,4}}, nb4_n_exts[][2] = {{4,1}, {4,4}};
int nb2_map[2]     = {1,0},      nb4_map[2]     = {0,1};
int nb2_dir[2]     = {1,-1},     nb4_dir[2]     = {1,1};

1: HYPRE_SStructGridCreate(MPI_COMM_WORLD, ndim, nparts, &grid);

/* Set grid extents and grid variables for part 3 */
2: HYPRE_SStructGridSetExtents(grid, part, extents[0], extents[1]);
3: HYPRE_SStructGridSetVariables(grid, part, nvars, vartypes);

/* Set spatial relationship between parts 3 and 2, then parts 3 and 4 */
4: HYPRE_SStructGridSetNeighborPart(grid, part, nb2_exts[0], nb2_exts[1],
    nb2_n_part, nb2_n_exts[0], nb2_n_exts[1], nb2_map, nb2_dir);
5: HYPRE_SStructGridSetNeighborPart(grid, part, nb4_exts[0], nb4_exts[1],
```

(continues on next page)

Figure 3.5: : Grid Setup Process. The “icons” illustrate the result of the numbered lines of code.

(continued from previous page)

```

nb4_n_part, nb4_n_exts[0], nb4_n_exts[1], nb4_map, nb4_dir);
6: HYPRE_SStructGridAssemble(grid);

```

Code on process 3 for setting up the grid in [Figure 3.2](#).

As with the `Struct` interface, each process describes that portion of the grid that it “owns”, one box at a time. [Figure 3.5](#) shows the code for setting up the grid on process 3 (the code for the other processes is similar). The “icons” at the top of the figure illustrate the result of the numbered lines of code. Process 3 needs to describe the data pictured in the bottom-right of the figure. That is, it needs to describe part 3 plus some additional neighbor information that ties part 3 together with the rest of the grid. The `Create()` routine creates an empty 2D grid object with five parts that lives on the `MPI_COMM_WORLD` communicator. The `SetExtents()` routine adds a new box to the grid. The `SetVariables()` routine associates three variables of type cell-centered,  $x$ -face, and  $y$ -face with part 3.

At this stage, the description of the data on part 3 is complete. However, the spatial relationship between this data and the data on neighboring parts is not yet defined. To do this, we need to relate the index space for part 3 with the index spaces of parts 2 and 4. More specifically, we need to tell the interface that the two grey boxes neighboring part 3 in the bottom-right of [Figure 3.5](#) also correspond to boxes on parts 2 and 4. This is done through the two calls to the `SetNeighborPart()` routine. We discuss only the first call, which describes the grey box on the right of the figure. Note that this grey box lives outside of the box extents for the grid on part 3, but it can still be described using the index-space for part 3 (recall [Figure 2.2](#)). That is, the grey box has extents  $(1, 0)$  and  $(4, 0)$  on part 3’s index-space, which is outside of part 3’s grid. The arguments for the `SetNeighborPart()` call are simply the lower and upper indices on part 3 and the corresponding indices on part 2. The final two arguments to the routine indicate that the positive  $x$ -direction on part 3 (i.e., the  $i$  component of the tuple  $(i, j)$ ) corresponds to the positive  $y$ -direction on part 2 and that the positive  $y$ -direction on part 3 corresponds to the positive  $x$ -direction on part 2.

The `Assemble()` routine is a collective call (i.e., must be called on all processes from a common synchronization point), and finalizes the grid assembly, making the grid “ready to use”.

With the neighbor information, it is now possible to determine where off-part stencil entries couple. Take, for example, any shared part boundary such as the boundary between parts 2 and 3. Along these boundaries, some stencil entries reach outside of the part. If no neighbor information is given, these entries are effectively zeroed out, i.e., they don’t participate in the discretization. However, with the additional neighbor information, when a stencil entry reaches into a neighbor box it is then coupled to the part described by that neighbor box information.

Another important consequence of the use of the `SetNeighborPart()` routine is that it can declare variables on different parts as being the same. For example, the face variables on the boundary of parts 2 and 3 are recognized as being shared by both parts (prior to the `SetNeighborPart()` call, there were two distinct sets of variables). Note also that these variables are of different types on the two parts; on part 2 they are  $x$ -face variables, but on part 3 they are  $y$ -face variables.

For brevity, we consider only the description of the  $y$ -face stencil in [Figure 3.2](#), i.e. the third stencil in the figure. To do this, the stencil entries are assigned unique labels between 0 and 8 and their “offsets” are described relative to the “center” of the stencil. This process is illustrated in [Figure 3.6](#). Nine calls are made to the routine `HYPRE_SStructStencilSetEntry()`. As an example, the call that describes stencil entry 5 in the figure is given the entry number 5, the offset  $(-1, 0)$ , and the identifier for the  $x$ -face variable (the variable to which this entry couples). Recall from [Figure 3.1](#) the convention used for referencing variables of different types. The geometry description uses the same convention, but with indices numbered relative to the referencing index  $(0, 0)$  for the stencil’s center. [Figure 3.8](#) shows the code for setting up the graph.

Figure 3.6: : Assignment of labels and geometries to the  $y$ -face stencil in [Figure 3.2](#).

Figure 3.7: : Alternate representation of the stencil configuration shown in Figure 3.7.

1:	2:	3:
----	----	----

Figure 3.8: : Graph Setup Process. The icons illustrate the result of the numbered lines of code.

```

HYPRE_SStructGraph graph;
HYPRE_SStructStencil c_stencil, x_stencil, y_stencil;
int c_var = 0, x_var = 1, y_var = 2;
int part;

1: HYPRE_SStructGraphCreate(MPI_COMM_WORLD, grid, &graph);

/* Set the cell-centered, x-face, and y-face stencils for each part */
for (part = 0; part < 5; part++)
{
2:   HYPRE_SStructGraphSetStencil(graph, part, c_var, c_stencil);
   HYPRE_SStructGraphSetStencil(graph, part, x_var, x_stencil);
   HYPRE_SStructGraphSetStencil(graph, part, y_var, y_stencil);
}

3: HYPRE_SStructGraphAssemble(graph);

```

Code on process 3 for setting up the graph for Figure 3.2.

With the above, we now have a complete description of the nonzero structure for the matrix. The matrix coefficients are then easily set in a manner similar to what is described in Section *Setting Up the Struct Matrix* using routines `MatrixSetValues()` and `MatrixSetBoxValues()` in the `SStruct` interface. As before, there are also `AddTo` variants of these routines. Likewise, setting up the right-hand-side is similar to what is described in Section *Setting Up the Struct Right-Hand-Side Vector*. See the hypr reference manual for details.

An alternative approach for describing the above problem through the interface is to use the `GraphAddEntries()` routine instead of the `GridSetNeighborPart()` routine. In this approach, the five parts would be explicitly “sewn” together by adding non-stencil couplings to the matrix graph. The main downside to this approach for block-structured grid problems is that variables along block boundaries are no longer considered to be the same variables on the corresponding parts that share these boundaries. For example, any face variable along the boundary between parts 2 and 3 in Figure 3.2 would represent two different variables that live on different parts. To “sew” the parts together correctly, we would need to explicitly select one of these variables as the representative that participates in the discretization, and make the other variable a dummy variable that is decoupled from the discretization by zeroing out appropriate entries in the matrix. All of these complications are avoided by using the `GridSetNeighborPart()` for this example.

## 3.2 Block-Structured Grids with Finite Elements

In this section, we describe how to use the `SStruct` interface to define block-structured grid problems with finite elements. We again do this by example, paying particular attention to the use of the FEM interface routines and the `GridSetSharedPart()` routine. See example code `ex14.c` for a complete implementation.

Consider a nodal finite element (FEM) discretization of the Laplace equation on the star-shaped grid in Figure 3.9. The local FEM stiffness matrix in the figure describes the coupling between the grid variables. Although we could still describe this problem using stencils as in Section *Block-Structured Grids with Stencils*, an FEM-based approach (available in hypr version 2.6.0b and later) is a more natural alternative.

Figure 3.9: : Example of a star-shaped grid with six logically-rectangular blocks and one nodal variable. Each block has an angle at the origin given by  $\gamma = \pi/3$ . The finite element stiffness matrix (right) is given in terms of the pictured variable ordering (left).

The grid in Figure 3.9 is defined in terms of six separate logically-rectangular parts, and each part is given a unique label between 0 and 5. Each part consists of a single box with lower index (1, 1) and upper index (9, 9), and the grid data is distributed on six processes such that data associated with part  $p$  lives on process  $p$ .

1:	2:	3:
4:	5:	6:

Figure 3.10: : FEM Grid Setup Process. The “icons” illustrate the result of the numbered lines of code.

```
HYPRE_SStructGrid grid;
int ndim = 2, nparts = 6, nvars = 1, part = 0;
int ilower[2] = {1,1}, iupper[2] = {9,9};
int vartypes[] = {HYPRE_SSTRUCT_VARIABLE_NODE};
int ordering[12] = {0,-1,-1, 0,+1,-1, 0,+1,+1, 0,-1,+1};

int s_part = 2;
int ilo[2] = {1,1}, iup[2] = {1,9}, offset[2] = {-1,0};
int s_ilo[2] = {1,1}, s_iup[2] = {9,1}, s_offset[2] = {0,-1};
int map[2] = {1,0};
int dir[2] = {-1,1};

1: HYPRE_SStructGridCreate(MPI_COMM_WORLD, ndim, nparts, &grid);

/* Set grid extents, grid variables, and FEM ordering for part 0 */
2: HYPRE_SStructGridSetExtents(grid, part, ilower, iupper);
3: HYPRE_SStructGridSetVariables(grid, part, nvars, vartypes);
4: HYPRE_SStructGridSetFEMOrdering(grid, part, ordering);

/* Set shared variables for parts 0 and 1 (0 and 2/3/4/5 not shown) */
5: HYPRE_SStructGridSetSharedPart(grid, part, ilo, iup, offset,
    s_part, s_ilo, s_iup, s_offset, map, dir);

6: HYPRE_SStructGridAssemble(grid);
```

Code on process 0 for setting up the grid in Figure 3.9.

As in Section *Block-Structured Grids with Stencils*, each process describes that portion of the grid that it “owns”, one box at a time. Figure 3.10 shows the code for setting up the grid on process 0 (the code for the other processes is similar). Process 0 needs to describe the data pictured in the bottom-right of the figure. That is, it needs to describe part 0 plus some additional information about shared data with other parts on the grid. The `SetFEMOrdering()` routine sets the ordering of the unknowns in an element (an element is always a grid cell in hypr). This determines the ordering of the data passed into the routines `MatrixAddFEMValues()` and `VectorAddFEMValues()` discussed later.

At this point, the layout of the data on part 0 is complete, but there is no relationship to the rest of the grid. To couple the parts, we need to tell hypr that some of the boundary variables on part 0 are shared with other parts, i.e., they are the same as some of the variables on other parts. This is done through five calls to the `SetSharedPart()` routine.

Only the first call is shown in the figure; the other four calls are similar. The arguments to this routine are the same as `SetNeighborPart()` with the addition of two new offset arguments, named `offset` and `s_offset` in the figure. Each offset represents a pointer from the cell center to one of the following: all variables in the cell (no nonzeros in offset); all variables on a face (only 1 nonzero); all variables on an edge (2 nonzeros); all variables at a point (3 nonzeros). The two offsets must be consistent with each other.

The graph is set up similarly to [Figure 3.8](#), except that the stencil calls are replaced by calls to `GraphSetFEM()`. The nonzero pattern of the stiffness matrix can also be set by calling the optional routine `GraphSetFEMSparsity()`.

Matrix and vector values are set one element at a time. For the example in this section, calls on part 0 would have the following form:

```
int part = 0;
int index[2] = {i, j};
double m_values[16] = {...};
double v_values[4] = {...};

HYPRE_SStructMatrixAddFEMValues(A, part, index, m_values);
HYPRE_SStructVectorAddFEMValues(v, part, index, v_values);
```

Here, `m_values` contains local stiffness matrix values and `v_values` contains local variable values. The global matrix and vector are assembled internally by hypr, using the shared variables to couple the parts.

### 3.3 Structured Adaptive Mesh Refinement

We now briefly discuss how to use the `SStruct` interface in a structured AMR application. Consider Poisson's equation on the simple cell-centered example grid illustrated in [Figure 3.11](#). For structured AMR applications, each refinement level should be defined as a unique part. There are two parts in this example: part 0 is the global coarse grid and part 1 is the single refinement patch. Note that the coarse unknowns underneath the refinement patch (gray dots in [Figure 3.11](#)) are not real physical unknowns; the solution in this region is given by the values on the refinement patch. In setting up the composite grid matrix [[McCo1989](#)] for hypr the equations for these "dummy" unknowns should be uncoupled from the other unknowns (this can easily be done by setting all off-diagonal couplings to zero in this region).

Figure 3.11: : Structured AMR grid example. Shaded regions correspond to process 0, unshaded to process 1. The grey dots are dummy variables.

In the example, parts are distributed across the same two processes with process 0 having the "left" half of both parts. The composite grid is then set up part-by-part by making calls to `GridSetExtents()` just as was done in [Section Block-Structured Grids with Stencils](#) and [Figure 3.5](#) (no `SetNeighborPart` calls are made in this example). Note that in the interface there is no required rule relating the indexing on the refinement patch to that on the global coarse grid; they are separate parts and thus each has its own index space. In this example, we have chosen the indexing such that refinement cell  $(2i, 2j)$  lies in the lower left quadrant of coarse cell  $(i, j)$ . Then the stencil is set up. In this example we are using a finite volume approach resulting in the standard 5-point stencil in [Section Setting Up the Struct Grid](#) in both parts.

The grid and stencil are used to define all intra-part coupling in the graph, the non-zero pattern of the composite grid matrix. The inter-part coupling at the coarse-fine interface is described by `GraphAddEntries()` calls. This coupling in the composite grid matrix is typically the composition of an interpolation rule and a discretization formula. In this example, we use a simple piecewise constant interpolation, i.e. the solution value in a coarse cell is equal to the solution value at the cell center. Then the flux across a portion of the coarse-fine interface is approximated by a difference of the solution values on each side. As an example, consider approximating the flux across the left interface of cell  $(6, 6)$  in [Figure 3.12](#). Let  $h$  be the coarse grid mesh size, and consider a local coordinate system with the origin at the center

of cell (6, 6). We approximate the flux as follows

$$\begin{aligned} \int_{-h/4}^{h/4} u_x(-h/4, s) ds &\approx \frac{h}{2} u_x(-h/4, 0) \approx \frac{h}{2} \frac{u(0, 0) - u(-3h/4, 0)}{3h/4} \\ &\approx \frac{2}{3} (u_{6,6} - u_{2,3}). \end{aligned}$$

The first approximation uses the midpoint rule for the edge integral, the second uses a finite difference formula for the derivative, and the third the piecewise constant interpolation to the solution in the coarse cell. This means that the equation for the variable at cell (6, 6) involves not only the stencil couplings to (6, 7) and (7, 6) on part 1 but also non-stencil couplings to (2, 3) and (3, 2) on part 0. These non-stencil couplings are described by `GraphAddEntries()` calls. The syntax for this call is simply the part and index for both the variable whose equation is being defined and the variable to which it couples. After these calls, the non-zero pattern of the matrix (and the graph) is complete. Note that the “west” and “south” stencil couplings simply “drop off” the part, and are effectively zeroed out (currently, this is only supported for the `HYPRE_PARCSR` object type, and these values must be manually zeroed out for other object types; see `MatrixSetObjectType()` in the reference manual).

Figure 3.12: : Coupling for equation at corner of refinement patch. Black lines (solid and broken) are stencil couplings. Gray line are non-stencil couplings.

The remaining step is to define the actual numerical values for the composite grid matrix. This can be done by either `MatrixSetValues()` calls to set entries in a single equation, or by `MatrixSetBoxValues()` calls to set entries for a box of equations in a single call. The syntax for the `MatrixSetValues()` call is a part and index for the variable whose equation is being set and an array of entry numbers identifying which entries in that equation are being set. The entry numbers may correspond to stencil entries or non-stencil entries.



## LINEAR-ALGEBRAIC SYSTEM INTERFACE (IJ)

The IJ interface described in this chapter is the lowest common denominator for specifying linear systems in hypre. This interface provides access to general sparse-matrix solvers in hypre, not to the specialized solvers that require more problem information.

### 4.1 IJ Matrix Interface

As with the other interfaces in hypre, the IJ interface expects to get data in distributed form because this is the only scalable approach for assembling matrices on thousands of processes. Matrices are assumed to be distributed by blocks of rows as follows:

$$\begin{bmatrix} A_0 \\ A_1 \\ \vdots \\ A_{P-1} \end{bmatrix}$$

In the above example, the matrix is distributed across the  $P$  processes,  $0, 1, \dots, P-1$  by blocks of rows. Each submatrix  $A_p$  is “owned” by a single process and its first and last row numbers are given by the global indices `ilower` and `iupper` in the `Create()` call below.

The following example code illustrates the basic usage of the IJ interface for building matrices:

```
MPI_Comm      comm;
HYPRE_IJMatrix ij_matrix;
HYPRE_ParCSRMatrix parcsr_matrix;
int           ilower, iupper;
int           jlower, jupper;
int           nrows;
int           *ncols;
int           *rows;
int           *cols;
double        *values;

HYPRE_IJMatrixCreate(comm, ilower, iupper, jlower, jupper, &ij_matrix);
HYPRE_IJMatrixSetObjectType(ij_matrix, HYPRE_PARCSR);
HYPRE_IJMatrixInitialize(ij_matrix);

/* set matrix coefficients */
HYPRE_IJMatrixSetValues(ij_matrix, nrows, ncols, rows, cols, values);
...
/* add-to matrix coefficients, if desired */
HYPRE_IJMatrixAddToValues(ij_matrix, nrows, ncols, rows, cols, values);
```

(continues on next page)

(continued from previous page)

```

...
HYPRE_IJMatrixAssemble(ij_matrix);
HYPRE_IJMatrixGetObject(ij_matrix, (void **) &parcsr_matrix);

```

The `Create()` routine creates an empty matrix object that lives on the `comm` communicator. This is a collective call (i.e., must be called on all processes from a common synchronization point), with each process passing its own row extents, `ilower` and `iupper`. The row partitioning must be contiguous, i.e., `iupper` for process `i` must equal `ilower`–1 for process `i+1`. Note that this allows matrices to have 0- or 1-based indexing. The parameters `jlower` and `jupper` define a column partitioning, and should match `ilower` and `iupper` when solving square linear systems. See Chapter [API](#) for more information.

The `SetObjectType()` routine sets the underlying matrix object type to `HYPRE_PARCSR` (this is the only object type currently supported). The `Initialize()` routine indicates that the matrix coefficients (or values) are ready to be set. This routine may or may not involve the allocation of memory for the coefficient data, depending on the implementation. The optional `SetRowSizes()` and `SetDiagOffdSizes()` routines mentioned later in this chapter and in Chapter [API](#), should be called before this step.

The `SetValues()` routine sets matrix values for some number of rows (`nrows`) and some number of columns in each row (`ncols`). The actual row and column numbers of the matrix values to be set are given by `rows` and `cols`. The coefficients can be modified with the `AddToValues()` routine. If `AddToValues()` is used to add to a value that previously didn't exist, it will set this value. Note that while `AddToValues()` will add to values on other processors, `SetValues()` does not set values on other processors. Instead if a user calls `SetValues()` on processor `i` to set a matrix coefficient belonging to processor `j`, processor `i` will erase all previous occurrences of this matrix coefficient, so they will not contribute to this coefficient on processor `j`. The actual coefficient has to be set on processor `j`.

The `Assemble()` routine is a collective call, and finalizes the matrix assembly, making the matrix “ready to use”. The `GetObject()` routine retrieves the built matrix object so that it can be passed on to hypr solvers that use the ParCSR internal storage format. Note that this is not an expensive routine; the matrix already exists in ParCSR storage format, and the routine simply returns a “handle” or pointer to it. Although we currently only support one underlying data storage format, in the future several different formats may be supported.

One can preset the row sizes of the matrix in order to reduce the execution time for the matrix specification. One can specify the total number of coefficients for each row, the number of coefficients in the row that couple the diagonal unknown to (Diag) unknowns in the same processor domain, and the number of coefficients in the row that couple the diagonal unknown to (Offd) unknowns in other processor domains:

```

HYPRE_IJMatrixSetRowSizes(ij_matrix, sizes);
HYPRE_IJMatrixSetDiagOffdSizes(matrix, diag_sizes, offdiag_sizes);

```

Once the matrix has been assembled, the sparsity pattern cannot be altered without completely destroying the matrix object and starting from scratch. However, one can modify the matrix values of an already assembled matrix. To do this, first call the `Initialize()` routine to re-initialize the matrix, then set or add-to values as before, and call the `Assemble()` routine to re-assemble before using the matrix. Re-initialization and re-assembly are very cheap, essentially a no-op in the current implementation of the code.

## 4.2 IJ Vector Interface

The following example code illustrates the basic usage of the IJ interface for building vectors:

```

MPI_Comm      comm;
HYPRE_IJVector ij_vector;
HYPRE_ParVector par_vector;
int           jlower, jupper;

```

(continues on next page)

(continued from previous page)

```

int          nvalues;
int          *indices;
double       *values;

HYPRE_IJVectorCreate(comm, jlower, jupper, &ij_vector);
HYPRE_IJVectorSetObjectType(ij_vector, HYPRE_PARCSR);
HYPRE_IJVectorInitialize(ij_vector);

/* set vector values */
HYPRE_IJVectorSetValues(ij_vector, nvalues, indices, values);
...

HYPRE_IJVectorAssemble(ij_vector);
HYPRE_IJVectorGetObject(ij_vector, (void **) &par_vector);

```

The `Create()` routine creates an empty vector object that lives on the `comm` communicator. This is a collective call, with each process passing its own index extents, `jlower` and `jupper`. The names of these extent parameters begin with a `j` because we typically think of matrix-vector multiplies as the fundamental operation involving both matrices and vectors. For matrix-vector multiplies, the vector partitioning should match the column partitioning of the matrix (which also uses the `j` notation). For linear system solves, these extents will typically match the row partitioning of the matrix as well.

The `SetObjectType()` routine sets the underlying vector storage type to `HYPRE_PARCSR` (this is the only storage type currently supported). The `Initialize()` routine indicates that the vector coefficients (or values) are ready to be set. This routine may or may not involve the allocation of memory for the coefficient data, depending on the implementation.

The `SetValues()` routine sets the vector values for some number (`nvalues`) of indices. The values can be modified with the `AddToValues()` routine. Note that while `AddToValues()` will add to values on other processors, `SetValues()` does not set values on other processors. Instead if a user calls `SetValues()` on processor  $i$  to set a value belonging to processor  $j$ , processor  $i$  will erase all previous occurrences of this matrix coefficient, so they will not contribute to this value on processor  $j$ . The actual value has to be set on processor  $j$ .

The `Assemble()` routine is a trivial collective call, and finalizes the vector assembly, making the vector “ready to use”. The `GetObject()` routine retrieves the built vector object so that it can be passed on to hypr solvers that use the `ParVector` internal storage format.

Vector values can be modified in much the same way as with matrices by first re-initializing the vector with the `Initialize()` routine.

### 4.3 A Scalable Interface

As explained in the previous sections, problem data is passed to the hypr library in its distributed form. However, as is typically the case for a parallel software library, some information regarding the global distribution of the data will be needed for hypr to perform its function. In particular, a solver algorithm requires that a processor obtain “nearby” data from other processors in order to complete the solve. While a processor may easily determine what data it needs from other processors, it may not know which processor owns the data it needs. Therefore, processors must determine their communication partners, or neighbors.

The straightforward approach to determining neighbors involves constructing a global partition of the data. This approach, however, requires  $O(P)$  storage and computations and is not scalable for machines with tens of thousands of processors. The *assumed partition* algorithm was developed to address this problem [BaFY2006]. It is the approach used in hypr.



## SOLVERS AND PRECONDITIONERS

There are several solvers available in hypr via different conceptual interfaces:

Solvers	System Interfaces		
	Struct	SStruct	IJ
Jacobi	X	X	
SMG	X	X	
PFMG	X	X	
Split		X	
SysPFMG		X	
SSAMG		X	
BoomerAMG		X	X
AMS		X	X
ADS		X	X
MGR			X
FSAI			X
ParaSails		X	X
ILU			X
Euclid		X	X
PILUT		X	X
PCG	X	X	X
GMRES	X	X	X
FlexGMRES	X	X	X
LGMRES	X	X	X
BiCGSTAB	X	X	X
Hybrid	X	X	X
LOBPCG	X	X	X

The procedure for setup and use of solvers and preconditioners is largely the same. We will refer to them both as solvers in the sequel except when noted. In normal usage, the preconditioner is chosen and constructed before the solver, and then handed to the solver as part of the solver's setup. In the following, we assume the most common usage pattern in which a single linear system is set up and then solved with a single righthand side. We comment later on considerations for other usage patterns.

**Setup:**

1. **Pass to the solver the information defining the problem.** In the typical user cycle, the user has passed this information into a matrix through one of the conceptual interfaces prior to setting up the solver. In this situation, the problem definition information is then passed to the solver by passing the constructed matrix into the solver. As described before, the matrix and solver must be compatible, in that the matrix must provide the services needed by the solver. Krylov solvers, for example, need only a matrix-vector multiplication. Most preconditioners, on the other hand, have additional requirements such as access to the matrix coefficients.

2. **Create the solver/preconditioner** via the `Create()` routine.
3. **Choose parameters for the preconditioner and/or solver.** Parameters are chosen through the `Set()` calls provided by the solver. Throughout hypr, we have made our best effort to give all parameters reasonable defaults if not chosen. However, for some preconditioners/solvers the best choices for parameters depend on the problem to be solved. We give recommendations in the individual sections on how to choose these parameters. Note that in hypr, convergence criteria can be chosen after the preconditioner/solver has been setup. For a complete set of all available parameters see Chapter [API](#).
4. **Pass the preconditioner to the solver.** For solvers that are not preconditioned, this step is omitted. The preconditioner is passed through the `SetPrecond()` call.
5. **Set up the solver.** This is just the `Setup()` routine. At this point the matrix and right hand side is passed into the solver or preconditioner. Note that the actual right hand side is not used until the actual solve is performed.

At this point, the solver/preconditioner is fully constructed and ready for use.

#### Use:

1. **Set convergence criteria.** Convergence can be controlled by the number of iterations, as well as various tolerances such as relative residual, preconditioned residual, etc. Like all parameters, reasonable defaults are used. Users are free to change these, though care must be taken. For example, if an iterative method is used as a preconditioner for a Krylov method, a constant number of iterations is usually required.
2. **Solve the system.** This is just the `Solve()` routine.

#### Finalize:

1. **Free the solver or preconditioner.** This is done using the `Destroy()` routine.

#### Synopsis

In general, a solver (let's call it SOLVER) is set up and run using the following routines, where  $A$  is the matrix,  $b$  the right hand side and  $x$  the solution vector of the linear system to be solved:

```
/* Create Solver */
int HYPRE_SOLVERCreate(MPI_COMM_WORLD, &solver);

/* Set certain parameters if desired */
HYPRE_SOLVERSetTol(solver, 1.e-8);
...

/* Set up Solver */
HYPRE_SOLVERSetup(solver, A, b, x);

/* Solve the system */
HYPRE_SOLVERsolve(solver, A, b, x);

/* Destroy the solver */
HYPRE_SOLVERDestroy(solver);
```

In the following sections, we will give brief descriptions of the available hypr solvers with some suggestions on how to choose the parameters as well as references for users who are interested in a more detailed description and analysis of the solvers. A complete list of all routines that are available can be found in Chapter [API](#).

## 5.1 SMG

SMG is a parallel semicoarsening multigrid solver for the linear systems arising from finite difference, finite volume, or finite element discretizations of the diffusion equation,

$$\nabla \cdot (D\nabla u) + \sigma u = f$$

on logically rectangular grids. The code solves both 2D and 3D problems with discretization stencils of up to 9-point in 2D and up to 27-point in 3D. See [Scha1998], [BrFJ2000], [FaJo2000] for details on the algorithm and its parallel implementation/performance.

SMG is a particularly robust method. The algorithm semicoarsens in the z-direction and uses plane smoothing. The xy plane-solves are effected by one V-cycle of the 2D SMG algorithm, which semicoarsens in the y-direction and uses line smoothing.

## 5.2 PFMG

PFMG is a parallel semicoarsening multigrid solver similar to SMG. See [AsFa1996], [FaJo2000] for details on the algorithm and its parallel implementation/performance.

The main difference between the two methods is in the smoother: PFMG uses simple pointwise smoothing. As a result, PFMG is not as robust as SMG, but is much more efficient per V-cycle.

## 5.3 SysPFMG

SysPFMG is a parallel semicoarsening multigrid solver for systems of elliptic PDEs. It is a generalization of PFMG, with the interpolation defined only within the same variable. The relaxation is of nodal type- all variables at a given point location are simultaneously solved for in the relaxation.

Although SysPFMG is implemented through the SStruct interface, it can be used only for problems with one grid part. Ideally, the solver should handle any of the seven variable types (cell-, node-, xface-, yface-, zface-, xedge-, yedge-, and zedge-based). However, it has been completed only for cell-based variables.

## 5.4 SplitSolve

SplitSolve is a parallel block Gauss-Seidel solver for semi-structured problems with multiple parts. For problems with only one variable, it can be viewed as a domain-decomposition solver with no grid overlapping.

Consider a multiple part problem given by the linear system  $Ax = b$ . Matrix  $A$  can be decomposed into a structured intra-variable block diagonal component  $M$  and a component  $N$  consisting of the inter-variable blocks and any unstructured connections between the parts. SplitSolve performs the iteration

$$x_{k+1} = \tilde{M}^{-1}(b + Nx_k),$$

where  $\tilde{M}^{-1}$  is a decoupled block-diagonal V(1,1) cycle, a separate cycle for each part and variable type. There are two V-cycle options, SMG and PFMG.

## 5.5 SSAMG

SSAMG is a semi-structured algebraic multigrid solver designed for problems defined through hypr's SStruct interface [MaFaYa23]. It targets applications with grids composed of multiple logically rectangular parts (e.g., block-structured, overset, and structured adaptive mesh refinement grids) and supports coupling between parts. The system matrix is naturally viewed as

$$A = S + U,$$

where  $S$  represents structured, stencil-based intra-part couplings and  $U$  captures unstructured inter-part couplings.

SSAMG builds a multilevel hierarchy algebraically, combining structured, directional coarsening within each part with algebraic treatment of couplings across parts. Compared to fully unstructured AMG, SSAMG exploits regularity for performance while allowing more general problem structure than strictly structured methods like PFMG.

### Note

SSAMG currently targets single-type variables on each part. Support for multiple variables per part is limited.

### 5.5.1 Configuration options

Below is a summary of configuration routines. See the reference manual for full details and defaults.

- `HYPRE_SStructSSAMGSetTol`: Convergence tolerance on the relative residual.
- `HYPRE_SStructSSAMGSetMaxIter`: Maximum number of SSAMG iterations.
- `HYPRE_SStructSSAMGSetMaxLevels`: Limit the multigrid hierarchy depth.
- `HYPRE_SStructSSAMGSetRelChange`: Require relative change of iterates to be small in addition to the residual criterion.
- `HYPRE_SStructSSAMGSetZeroGuess` / `SetNonZeroGuess`: Specify initial guess semantics.
- `HYPRE_SStructSSAMGSetInterpType`: Interpolation choice within parts:
  - `-1`: Structured interpolation only (default)

- 0: Structured plus classical modified unstructured interpolation
- HYPRE\_SStructSSAMGSetRelaxType: Smoother selection:
  - 0: Jacobi
  - 1: Weighted Jacobi (default)
  - 2: L1-Jacobi
  - 10: Red/Black Gauss–Seidel (symmetric RB/BR)
- HYPRE\_SStructSSAMGSetRelaxWeight: Jacobi weight (used by Jacobi variants).
- HYPRE\_SStructSSAMGSetNumPreRelax / SetNumPostRelax: Number of sweeps before/after coarse correction on each level.
- HYPRE\_SStructSSAMGSetNumCoarseRelax: Number of sweeps on the coarsest level (when not delegating to BoomerAMG).
- HYPRE\_SStructSSAMGSetSkipRelax: Enable skip-relax on levels whose coarsening directions repeat those from previous cycles (useful for isotropic problems).
- HYPRE\_SStructSSAMGSetCoarseSolverType: Coarse solver selection:
  - 0: Weighted Jacobi (default)
  - 1: BoomerAMG (hybrid hierarchy on coarse levels)
- HYPRE\_SStructSSAMGSetDxyz: Provide a per-part grid-spacing metric used to guide coarsening direction choices.
- HYPRE\_SStructSSAMGSetMaxCoarseSize: Limit the maximum coarse problem size (set to zero to disable).
- HYPRE\_SStructSSAMGSetLogging: Enable internal logging.
- HYPRE\_SStructSSAMGSetPrintLevel / SetPrintFreq: Control verbosity and output frequency.
- HYPRE\_SStructSSAMGGetNumIterations / GetFinalRelativeResidualNorm: Retrieve iteration count and final relative residual norm.

## 5.5.2 Algorithm notes

- **Per-part semi-coarsening:** Each part selects a dominant coupling direction (heuristically related to an effective grid spacing metric) and coarsens by a factor of two in that direction. This adapts to anisotropy that varies across parts.
- **Structured interpolation:** Prolongation is constructed within each part from the structured component of the operator and uses the transpose as restriction. This limits stencil growth and captures heterogeneity.
- **Hybrid coarse solve:** Delegating coarse levels to BoomerAMG can reduce the overall number of levels in the multigrid hierarchy while retaining efficiency on fine levels.

## 5.5.3 Minimal C example

The snippet below shows how to create and configure an SSAMG solver and solves a problem with an assembled SStruct matrix  $A$  and vectors  $b$ ,  $x$ . Only solver creation and configuration are shown. For better robustness, we recommend using SSAMG as a preconditioner for a Krylov solver.

```
HYPRE_SStructSolver ssamg;
HYPRE_SStructSSAMGCreate(MPI_COMM_WORLD, &ssamg);

HYPRE_SStructSSAMGSetTol(ssamg, 1e-8);
```

(continues on next page)

(continued from previous page)

```

HYPRE_SStructSSAMGSetMaxIter(ssamg, 50);
HYPRE_SStructSSAMGSetInterpType(ssamg, -1);      /* Structured-only interpolation */
HYPRE_SStructSSAMGSetRelaxType(ssamg, 1);        /* Weighted Jacobi */
HYPRE_SStructSSAMGSetNumPreRelax(ssamg, 1);
HYPRE_SStructSSAMGSetNumPostRelax(ssamg, 1);
HYPRE_SStructSSAMGSetSkipRelax(ssamg, 1);        /* Skip relaxation on certain levels */
HYPRE_SStructSSAMGSetCoarseSolverType(ssamg, 1); /* Switch to BoomerAMG on coarse level.
↪ */

HYPRE_SStructSSAMGSetup(ssamg, A, b, x);
HYPRE_SStructSSAMGSolve(ssamg, A, b, x);

HYPRE_Int iters; HYPRE_Real relres;
HYPRE_SStructSSAMGGetNumIterations(ssamg, &iters);
HYPRE_SStructSSAMGGetFinalRelativeResidualNorm(ssamg, &relres);

HYPRE_SStructSSAMGDestroy(ssamg);

```

## 5.6 Hybrid

The hybrid solver is designed to detect whether a multigrid preconditioner is needed when solving a linear system and possibly avoid the expensive setup of a preconditioner if a system can be solved efficiently with a diagonally scaled Krylov solver, e.g. a strongly diagonally dominant system. It first uses a diagonally scaled Krylov solver, which can be chosen by the user (the default is conjugate gradient, but one should use GMRES if the matrix of the linear system to be solved is nonsymmetric). It monitors how fast the Krylov solver converges. If there is not sufficient progress, the algorithm switches to a preconditioned Krylov solver.

If used through the `Struct` interface, the solver is called `StructHybrid` and can be used with the preconditioners `SMG` and `PFMG` (default). It is called `ParCSRHybrid`, if used through the `IJ` interface and is used here with `BoomerAMG`. The user can determine the average convergence speed by setting a convergence tolerance  $0 \leq \theta < 1$  via the routine `HYPRE_StructHybridSetConvergenceTol` or `HYPRE_ParCSRHybridSetConvergenceTol`. The default setting is 0.9.

The average convergence factor  $\rho_i = \left( \frac{\|r_i\|}{\|r_0\|} \right)^{1/i}$  is monitored within the chosen Krylov solver, where  $r_i = b - Ax_i$  is the  $i$ -th residual. Convergence is considered too slow when

$$\left( 1 - \frac{|\rho_i - \rho_{i-1}|}{\max(\rho_i, \rho_{i-1})} \right) \rho_i > \theta.$$

When this condition is fulfilled the hybrid solver switches from a diagonally scaled Krylov solver to a preconditioned solver.

## 5.7 BoomerAMG

BoomerAMG is a parallel implementation of the algebraic multigrid method [RuSt1987]. It can be used both as a solver or as a preconditioner. The user can choose between various different parallel coarsening techniques, interpolation and relaxation schemes. The default settings for CPUs, HMIS (coarsening 8) combined with a distance-two interpolation (6) truncated to 4 or 5 elements per row, should work fairly well for two- and three-dimensional diffusion problems. Additional reduction in complexity and increased scalability can often be achieved using one or two levels of aggressive coarsening.

### 5.7.1 Parameter Options

Various BoomerAMG functions and options are mentioned below. However, for a complete listing and description of all available functions, see the reference manual.

BoomerAMG's `Create` function differs from the synopsis in that it has only one parameter `HYPRE_BoomerAMGCreate(HYPRE_Solver *solver)`. It uses the communicator of the matrix `A`.

### 5.7.2 Coarsening Options

Coarsening can be set by the user using the function `HYPRE_BoomerAMGSetCoarsenType`. A detailed description of various coarsening techniques can be found in [HeYa2002], [Yang2005].

Various coarsening techniques are available:

- the Cleary-Luby-Jones-Plassman (CLJP) coarsening,
- parallel versions of the classical RS coarsening described in [HeYa2002].
- the Falgout coarsening which is a combination of CLJP and the classical RS coarsening algorithm,
- CGC and CGC-E coarsenings [GrMS2006a], [GrMS2006b],
- PMIS and HMIS coarsening algorithms which lead to coarsenings with lower complexities [DeYH2004] as well as
- aggressive coarsening, which can be applied to any of the coarsening techniques mentioned above and thus achieving much lower complexities and lower memory use [Stue1999].

To use aggressive coarsening users have to set the number of levels to which they want to apply aggressive coarsening (starting with the finest level) via `HYPRE_BoomerAMGSetAggNumLevels`. Since aggressive coarsening requires long range interpolation, multipass interpolation is always used on levels with aggressive coarsening, unless the user specifies another long-range interpolation suitable for aggressive coarsening via `HYPRE_BoomerAMGSetAggInterpType`.

Note that the default coarsening for CPUs is HMIS, for GPUs PMIS [DeYH2004].

### 5.7.3 Interpolation Options

Various interpolation techniques can be set using `HYPRE_BoomerAMGSetInterpType`:

- the “classical” interpolation (0) as defined in [RuSt1987],
- direct interpolation (3) [Stue1999],
- standard interpolation (8) [Stue1999],

- an extended “classical” interpolation, which is a long range interpolation and is recommended to be used with PMIS and HMIS coarsening for harder problems (6) [DFNY2008],
- distance-two interpolation based on matrix operations (17) [LiSY2021],
- multipass interpolation (4) [Stue1999],
- two-stage interpolation [Yang2010],
- Jacobi interpolation [Stue1999],
- the “classical” interpolation modified for hyperbolic PDEs (2).

Jacobi interpolation is only used to improve certain interpolation operators and can be used with `HYPRE_BoomerAMGSetPostInterpType`. Since some of the interpolation operators might generate large stencils, it is often possible and recommended to control complexity and truncate the interpolation operators using `HYPRE_BoomerAMGSetTruncFactor` and/or `HYPRE_BoomerAMGSetPMaxElmts`, or `HYPRE_BoomerAMGSetJacobiTruncTheshold` (for Jacobi interpolation only).

Note that the default interpolation is extended+i interpolation [DFNY2008] truncated to 4 elements per row, for CPUs, and a version of this interpolation based on matrix operations for GPUs [LiSY2021].

### 5.7.4 Non-Galerkin Options

In order to reduce communication, there is a non-Galerkin coarse grid sparsification option available [FaSc2014]. This option can be used by itself or with existing strategies to reduce communication such as aggressive coarsening and HMIS coarsening. To use, call `HYPRE_BoomerAMGSetNonGalerkTo1`, which gives BoomerAMG a list of level specific non-Galerkin drop tolerances. It is common to drop more aggressively on coarser levels. A common choice of drop-tolerances is [0.0, 0.01, 0.05] where the value of 0.0 will skip the non-Galerkin process on the first coarse level (level 1), use a drop-tolerance of 0.01 on the second coarse level (level 2) and then use 0.05 on all subsequent coarse levels. While still experimental, this capability has significantly improved performance on a variety of problems. See the `ij` driver for an example usage and the reference manual for more details.

### 5.7.5 Smoother Options

A good overview of parallel smoothers and their properties can be found in [BFKY2011]. Various of the described relaxation techniques are available:

- weighted Jacobi relaxation (0),
- a hybrid Gauss-Seidel / Jacobi relaxation scheme (3 4),
- a symmetric hybrid Gauss-Seidel / Jacobi relaxation scheme (6),
- 11-Gauss-Seidel or Jacobi (13 14 18 8),
- Chebyshev smoothers (16),
- two-stage Gauss-Seidel smoothers (11 12) [BKRHSMTY2021],
- hybrid block and Schwarz smoothers [Yang2004],
- Incomplete LU factorization, see *ILU as Smoother for BoomerAMG*.
- Factorized Sparse Approximate Inverse (FSAI), see *FSAI as Smoother to BoomerAMG*.

Point relaxation schemes can be set using `HYPRE_BoomerAMGSetRelaxType` or, if one wants to specifically set the up cycle, down cycle or the coarsest grid, with `HYPRE_BoomerAMGSetCycleRelaxType`. To use the more complicated smoothers, e.g. block, Schwarz, ILU smoothers, it is necessary to use `HYPRE_BoomerAMGSetSmoothType` and `HYPRE_BoomerAMGSetSmoothNumLevels`. There are further parameter choices for the individual smoothers, which are described in the reference manual. The default relaxation type is 11-Gauss-Seidel, using a forward solve on the down cycle and a backward solve on the up-cycle, to keep symmetry. Note that if BoomerAMG is used as a preconditioner

for conjugate gradient, it is necessary to use a symmetric smoother. Other symmetric options are weighted Jacobi or hybrid symmetric Gauss-Seidel.

## 5.7.6 AMG for systems of PDEs

If the user wants to solve systems of PDEs and can provide information on which variables belong to which function, BoomerAMG's systems AMG version can also be used. Functions that enable the user to access the systems AMG version are `HYPRE_BoomerAMGSetNumFunctions`, `HYPRE_BoomerAMGSetDofFunc` and `HYPRE_BoomerAMGSetNodal`.

There are basically two approaches to deal with matrices derived from systems of PDEs. The unknown-based approach (which is the default) treats variables corresponding to the same unknown or function separately, i.e., when coarsening or generating interpolation, connections between variables associated with different unknowns are ignored. This can work well for weakly coupled PDEs, but will be problematic for strongly coupled PDEs. For such problems, we recommend to use hypr's multigrid reduction (MGR) solver. The second approach, called the nodal approach, considers all unknowns at a physical grid point together such that coarsening, interpolation and relaxation occur in a point-wise fashion. It is possible and sometimes preferred to combine nodal coarsening with unknown-based interpolation. For this case, `HYPRE_BoomerAMGSetNodal` should be set  $> 1$ . For details see the reference manual.

If the user can provide the near null-space vectors, such as the rigid body modes for linear elasticity problems, an interpolation is available that will incorporate these vectors with `HYPRE_BoomerAMGSetInterpVectors` and `HYPRE_BoomerAMGSetInterpVecVariant`. This can lead to improved convergence and scalability [BaKY2010].

## 5.7.7 Special AMG Cycles

The default cycle is a V(1,1)-cycle, however it is possible to change the number of sweeps of the up- and down-cycle as well as the coarse grid. One can also choose a W-cycle, however for parallel processing this is not recommended, since it is not scalable.

BoomerAMG also provides an additive V(1,1)-cycle as well as a mult-additive V(1,1)-cycle and a simplified version [VaYa2014]. The additive variants can only be used with weighted Jacobi or 11-Jacobi smoothing.

## 5.7.8 GPU-supported Options

In general, most BoomerAMG options are GPU-enabled and do not require unified memory. If any selected option lacks GPU support, then unified memory is required for those parts to run on the CPU. When building hypr without `--enable-unified-memory`, make sure that all chosen BoomerAMG parameters are GPU-compatible. The currently supported GPU-enabled BoomerAMG options include:

- Coarsening: PMIS (8)
- Interpolation: direct (3), BAMG-direct (15), extended (14), extended+i (6) and extended+e (18)
- Aggressive coarsening
- Multipass interpolation with aggressive coarsening (4 or 8)
- Second-stage interpolation with aggressive coarsening: extended (5), extended+i (6) and extended+e (7)
- Smoother: Jacobi (7), 11-Jacobi (18), hybrid Gauss Seidel/SSOR (3 4 6), two-stage Gauss-Seidel (11,12) [BKRHSMTY2021], and Chebyshev (16)
- Relaxation order can be 0, lexicographic order, or C/F for (7) and (18)

## 5.7.9 Memory locations and execution policies

Hypr provides two user-level memory locations, `HYPRE_MEMORY_HOST` and `HYPRE_MEMORY_DEVICE`, where `HYPRE_MEMORY_HOST` is always the CPU memory while `HYPRE_MEMORY_DEVICE` can be mapped to different memory spaces based on the configure options of hypr. When built with `--with-cuda`, `--with-hip`, `--with-sycl`, or `--with-device-openmp`, `HYPRE_MEMORY_DEVICE` is the GPU device memory, and when built additionally with

--enable-unified-memory, it is the GPU unified memory (UM). For a non-GPU build, HYPRE\_MEMORY\_DEVICE is also mapped to the CPU memory. The default memory location of hypre's matrix and vector objects is HYPRE\_MEMORY\_DEVICE, which can be changed at runtime by HYPRE\_SetMemoryLocation(...).

The execution policies define the platform of running computations based on the memory locations of participating objects. The default policy is HYPRE\_EXEC\_HOST, i.e., executing on the host **if the objects are accessible from the host**. It can be adjusted by HYPRE\_SetExecutionPolicy(...). Clearly, this policy only affects objects in UM, since UM is accessible from **both CPUs and GPUs**.

A sample code of setting up IJ matrix  $A$  and solve  $Ax = b$  using AMG-preconditioned CG on GPUs is shown below.

```

cudaSetDevice(device_id); /* GPU binding */
...
HYPRE_Initialize(); /* must be the first HYPRE function call */
...
/* AMG in GPU memory (default) */
HYPRE_SetMemoryLocation(HYPRE_MEMORY_DEVICE);
/* setup AMG on GPUs */
HYPRE_SetExecutionPolicy(HYPRE_EXEC_DEVICE);
/* use hypre's SpGEMM instead of vendor implementation */
HYPRE_SetSpGemmUseVendor(FALSE);
/* use GPU RNG */
HYPRE_SetUseGpuRand(TRUE);
if (useHypreGpuMemPool)
{
    /* use hypre's GPU memory pool */
    HYPRE_SetGPUMemoryPoolSize(bin_growth, min_bin, max_bin, max_bytes);
}
else if (useUmpireGpuMemPool)
{
    /* or use Umpire GPU memory pool */
    HYPRE_SetUmpireUMPoolName("HYPRE_UM_POOL_TEST");
    HYPRE_SetUmpireDevicePoolName("HYPRE_DEVICE_POOL_TEST");
}
...
/* setup IJ matrix A */
HYPRE_IJMatrixCreate(comm, first_row, last_row, first_col, last_col, &ij_A);
HYPRE_IJMatrixSetObjectType(ij_A, HYPRE_PARCSR);
/* GPU pointers; efficient in large chunks */
HYPRE_IJMatrixAddToValues(ij_A, num_rows, num_cols, rows, cols, data);
HYPRE_IJMatrixAssemble(ij_A);
HYPRE_IJMatrixGetObject(ij_A, (void **) &parcsr_A);
...
/* setup AMG */
HYPRE_ParCSRPCGCreate(comm, &solver);
HYPRE_BoomerAMGCreate(&precon);
HYPRE_BoomerAMGSetRelaxType(precon, rlx_type); /* 3, 4, 6, 7, 18, 11, 12 */
HYPRE_BoomerAMGSetRelaxOrder(precon, FALSE); /* must be false */
HYPRE_BoomerAMGSetCoarsenType(precon, coarsen_type); /* 8 */
HYPRE_BoomerAMGSetInterpType(precon, interp_type); /* 3, 15, 6, 14, 18 */
HYPRE_BoomerAMGSetAggNumLevels(precon, agg_num_levels);
HYPRE_BoomerAMGSetAggInterpType(precon, agg_interp_type); /* 4, 5, 6, 7, 8 */
HYPRE_BoomerAMGSetKeepTranspose(precon, TRUE); /* keep transpose to avoid SpMTV */
HYPRE_BoomerAMGSetRAP2(precon, FALSE); /* RAP in two multiplications

```

(continues on next page)

(continued from previous page)

```

                                (default: FALSE) */
HYPRE_ParCSRPCGSetPrecond(solver, HYPRE_BoomerAMGSolve, HYPRE_BoomerAMGSetup,
                          precon);
HYPRE_PCGSetup(solver, parcsr_A, b, x);
...
/* solve */
HYPRE_PCGSolve(solver, parcsr_A, b, x);
...
HYPRE_Finalize(); /* must be the last HYPRE function call */

```

HYPRE\_Initialize() must be called and precede all the other HYPRE\_ functions, and HYPRE\_Finalize() must be called before exiting.

### 5.7.10 Miscellaneous

For best performance, it might be necessary to set certain parameters, which will affect both coarsening and interpolation. One important parameter is the strong threshold, which can be set using the function HYPRE\_BoomerAMGSetStrongThreshold. The default value is 0.25, which appears to be a good choice for diffusion problems. The choice of the strength threshold is problem dependent. For example, elasticity problems often require a larger strength threshold.

## 5.8 AMS

AMS (the Auxiliary-space Maxwell Solver) is a parallel unstructured Maxwell solver for edge finite element discretizations of the variational problem

$$\text{Find } \mathbf{u} \in \mathbf{V}_h : \quad (\alpha \nabla \times \mathbf{u}, \nabla \times \mathbf{v}) + (\beta \mathbf{u}, \mathbf{v}) = (\mathbf{f}, \mathbf{v}), \quad \text{for all } \mathbf{v} \in \mathbf{V}_h. \quad (5.1)$$

Here  $\mathbf{V}_h$  is the lowest order Nedelec (edge) finite element space, and  $\alpha > 0$  and  $\beta \geq 0$  are scalar, or SPD matrix coefficients. AMS was designed to be scalable on problems with variable coefficients, and allows for  $\beta$  to be zero in part or the whole domain. In either case the resulting problem is only semidefinite, and for solvability the right-hand side should be chosen to satisfy compatibility conditions.

AMS is based on the auxiliary space methods for definite Maxwell problems proposed in [HiXu2006]. For more details, see [KoVa2009].

### 5.8.1 Overview

Let  $\mathbf{A}$  and  $\mathbf{b}$  be the stiffness matrix and the load vector corresponding to (5.1). Then the resulting linear system of interest reads,

$$\mathbf{A} \mathbf{x} = \mathbf{b}. \quad (5.2)$$

The coefficients  $\alpha$  and  $\beta$  are naturally associated with the “stiffness” and “mass” terms of  $\mathbf{A}$ . Besides  $\mathbf{A}$  and  $\mathbf{b}$ , AMS requires the following additional user input:

1. The discrete gradient matrix  $G$  representing the edges of the mesh in terms of its vertices.  $G$  has as many rows as the number of edges in the mesh, with each row having two nonzero entries:  $+1$  and  $-1$  in the columns corresponding to the vertices composing the edge. The sign is determined based on the orientation of the edge. We require that  $G$  includes all (interior and boundary) edges and vertices.
2. The representations of the constant vector fields  $(1, 0, 0)$ ,  $(0, 1, 0)$ , and  $(0, 0, 1)$  in the  $\mathbf{V}_h$  basis, given as three vectors:  $G_x$ ,  $G_y$ , and  $G_z$ . Note that since no boundary conditions are imposed on  $G$ , the above vectors can be computed as  $G_x = Gx$ ,  $G_y = Gy$  and  $G_z = Gz$ , where  $x$ ,  $y$ , and  $z$  are vectors representing the coordinates of the vertices of the mesh.

In addition to the above quantities, AMS can utilize the following (optional) information:

- The Poisson matrices  $A_\alpha$  and  $A_\beta$ , corresponding to assembling of the forms  $(\alpha \nabla u, \nabla v) + (\beta u, v)$  and  $(\beta \nabla u, \nabla v)$  using standard linear finite elements on the same mesh.

Internally, AMS proceeds with the construction of the following additional objects:

- $A_G$  – a matrix associated with the mass term which is either  $G^T \mathbf{A} G$  or the Poisson matrix  $A_\beta$  (if given).
- $\mathbf{\Pi}$  – the matrix representation of the interpolation operator from vector linear to edge finite elements.
- $\mathbf{A}_\mathbf{\Pi}$  – a matrix associated with the stiffness term which is either  $\mathbf{\Pi}^T \mathbf{A} \mathbf{\Pi}$  or a block-diagonal matrix with diagonal blocks  $A_\alpha$  (if given).
- $B_G$  and  $\mathbf{B}_\mathbf{\Pi}$  – efficient (AMG) solvers for  $A_G$  and  $\mathbf{A}_\mathbf{\Pi}$ .

The solution procedure then is a three-level method using smoothing in the original edge space and subspace corrections based on  $B_G$  and  $\mathbf{B}_\mathbf{\Pi}$ . We can employ a number of options here utilizing various combinations of the smoother and solvers in additive or multiplicative fashion. If  $\beta$  is identically zero one can skip the subspace correction associated with  $G$ , in which case the solver is a two-level method.

## 5.8.2 Sample Usage

AMS can be used either as a solver or as a preconditioner. Below we list the sequence of hypr calls needed to create and use it as a solver. See example code `ex15.c` for a complete implementation. We start with the allocation of the `HYPRE_Solver` object:

```
HYPRE_Solver solver;
HYPRE_AMSCreate(&solver);
```

Next, we set a number of solver parameters. Some of them are optional, while others are necessary in order to perform the solver setup.

AMS offers the option to set the space dimension. By default we consider the dimension to be 3. The only other option is 2, and it can be set with the function given below. We note that a 3D solver will still work for a 2D problem, but it will be slower and will require more memory than necessary.

```
HYPRE_AMSSetDimension(solver, dim);
```

The user is required to provide the discrete gradient matrix  $G$ . AMS expects a matrix defined on the whole mesh with no boundary edges/nodes excluded. It is essential to **not** impose any boundary conditions on  $G$ . Regardless of which hypr conceptual interface was used to construct  $G$ , one can obtain a ParCSR version of it. This is the expected format in the following function.

```
HYPRE_AMSSetDiscreteGradient(solver, G);
```

In addition to  $G$ , we need one additional piece of information in order to construct the solver. The user has the option to choose either the coordinates of the vertices in the mesh or the representations of the constant vector fields in the edge element basis. In both cases three hypr parallel vectors should be provided. For 2D problems, the user can set the third vector to NULL. The corresponding function calls read:

```
HYPRE_AMSSetCoordinateVectors(solver, x, y, z);
```

or

```
HYPRE_AMSSetEdgeConstantVectors(solver, one_zero_zero, zero_one_zero, zero_zero_one);
```

The vectors `one_zero_zero`, `zero_one_zero` and `zero_zero_one` above correspond to the constant vector fields  $(1, 0, 0)$ ,  $(0, 1, 0)$  and  $(0, 0, 1)$ .

The remaining solver parameters are optional. For example, the user can choose a different cycle type by calling

```
HYPRE_AMSSetCycleType(solver, cycle_type); /* default value: 1 */
```

The available cycle types in AMS are:

- `cycle_type=1`: multiplicative solver (01210)
- `cycle_type=2`: additive solver (0 + 1 + 2)
- `cycle_type=3`: multiplicative solver (02120)
- `cycle_type=4`: additive solver (010 + 2)
- `cycle_type=5`: multiplicative solver (0102010)
- `cycle_type=6`: additive solver (1 + 020)
- `cycle_type=7`: multiplicative solver (0201020)
- `cycle_type=8`: additive solver (0(1 + 2)0)
- `cycle_type=11`: multiplicative solver (013454310)
- `cycle_type=12`: additive solver (0 + 1 + 3 + 4 + 5)
- `cycle_type=13`: multiplicative solver (034515430)
- `cycle_type=14`: additive solver (01(3 + 4 + 5)10)

Here we use the following convention for the three subspace correction methods: 0 refers to smoothing, 1 stands for BoomerAMG based on  $B_G$ , and 2 refers to a call to BoomerAMG for  $\mathbf{B}_\Pi$ . The values 3, 4 and 5 refer to the scalar subspaces corresponding to the  $x$ ,  $y$  and  $z$  components of  $\Pi$ .

The abbreviation  $xyyz$  for  $x, y, z \in \{0, 1, 2, 3, 4, 5\}$  refers to a multiplicative subspace correction based on solvers  $x$ ,  $y$ , and  $z$  (in that order). The abbreviation  $x + y + z$  stands for an additive subspace correction method based on  $x$ ,  $y$  and  $z$  solvers. The additive cycles are meant to be used only when AMS is called as a preconditioner. In our experience the choices `cycle_type=1, 5, 8, 11, 13` often produced fastest solution times, while `cycle_type=7` resulted in smallest number of iterations.

Additional solver parameters, such as the maximum number of iterations, the convergence tolerance and the output level, can be set with

```
HYPRE_AMSsetMaxIter(solver, maxit); /* default value: 20 */
HYPRE_AMSsetTol(solver, tol); /* default value: 1e-6 */
HYPRE_AMSsetPrintLevel(solver, print); /* default value: 1 */
```

More advanced parameters, affecting the smoothing and the internal AMG solvers, can be set with the following three functions:

```
HYPRE_AMSsetSmoothingOptions(solver, 2, 1, 1.0, 1.0);
HYPRE_AMSsetAlphaAMGOptions(solver, 10, 1, 3, 0.25, 0, 0);
HYPRE_AMSsetBetaAMGOptions(solver, 10, 1, 3, 0.25, 0, 0);
```

For (singular) problems where  $\beta = 0$  in the whole domain, different (in fact simpler) version of the AMS solver is offered. To allow for this simplification, use the following hypr call

```
HYPRE_AMSsetBetaPoissonMatrix(solver, NULL);
```

If  $\beta$  is zero only in parts of the domain, the problem is still singular, but the AMS solver will try to detect this and construct a non-singular preconditioner. Though this often works well in practice, AMS also provides a more ro-

bust version for solving such singular problems to very low convergence tolerances. This version takes advantage of additional information: the list of nodes which are interior to a zero-conductivity region provided by the function

```
HYPRE_AMSSetInteriorNodes(solver, HYPRE_ParVector interior_nodes);
```

A node is interior, if its entry in the `interior_nodes` array is 1.0. Based on this array, a restricted discrete gradient operator  $G_0$  is constructed, and AMS is then defined based on the matrix  $\mathbf{A} + \delta G_0^T G_0$  which is non-singular, and a small  $\delta > 0$  perturbation of  $\mathbf{A}$ . When iterating with this preconditioner, it is advantageous to project on the compatible subspace  $\text{Ker}(G_0^T)$ . This can be done periodically, or manually through the functions

```
HYPRE_AMSSetProjectionFrequency(solver, int projection_frequency);
HYPRE_AMSProjectOutGradients(solver, HYPRE_ParVector x);
```

Two additional matrices are constructed in the setup of the AMS method—one corresponding to the coefficient  $\alpha$  and another corresponding to  $\beta$ . This may lead to prohibitively high memory requirements, and the next two function calls may help to save some memory. For example, if the Poisson matrix with coefficient  $\beta$  (denoted by `Abeta`) is available then one can avoid one matrix construction by calling

```
HYPRE_AMSSetBetaPoissonMatrix(solver, Abeta);
```

Similarly, if the Poisson matrix with coefficient  $\alpha$  is available (denoted by `Aalpha`) the second matrix construction can also be avoided by calling

```
HYPRE_AMSSetAlphaPoissonMatrix(solver, Aalpha);
```

Note the following regarding these functions:

- Both of them change their input. More specifically, the diagonal entries of the input matrix corresponding to eliminated degrees of freedom (due to essential boundary conditions) are penalized.
- It is assumed that their essential boundary conditions of `A`, `Abeta` and `Aalpha` are on the same part of the boundary.
- `HYPRE_AMSSetAlphaPoissonMatrix` forces the AMS method to use a simpler, but weaker (in terms of convergence) method. With this option, the multiplicative AMS cycle is not guaranteed to converge with the default parameters. The reason for this is the fact the solver is not variationally obtained from the original matrix (it utilizes the auxiliary Poisson-like matrices `Abeta` and `Aalpha`). Therefore, it is recommended in this case to use AMS as preconditioner only.

After the above calls, the solver is ready to be constructed. The user has to provide the stiffness matrix `A` (in ParCSR format) and the hypr parallel vectors `b` and `x`. (The vectors are actually not used in the current AMS setup.) The setup call reads,

```
HYPRE_AMSSetup(solver, A, b, x);
```

It is important to note the order of the calling sequence. For example, do **not** call `HYPRE_AMSSetup` before calling `HYPRE_AMSSetDiscreteGradient` and one of the functions `HYPRE_AMSSetCoordinateVectors` or `HYPRE_AMSSetEdgeConstantVectors`.

Once the setup has completed, we can solve the linear system by calling

```
HYPRE_AMSolve(solver, A, b, x);
```

Finally, the solver can be destroyed with

```
HYPRE_AMSDestroy(&solver);
```

More details can be found in the files `ams.h` and `ams.c` located in the `parcsr_ls` directory.

### 5.8.3 High-order Discretizations

In addition to the interface for the lowest-order Nedelec elements described in the previous subsections, AMS also provides support for (arbitrary) high-order Nedelec element discretizations. Since the robustness of AMS depends on the performance of BoomerAMG on the associated (high-order) auxiliary subspace problems, we note that the convergence may not be optimal for large polynomial degrees  $k \geq 1$ .

In the high-order AMS interface, the user does not need to provide the coordinates of the vertices (or the representations of the constant vector fields in the edge basis), but instead should construct and pass the Nedelec interpolation matrix  $\mathbf{\Pi}$  which maps (high-order) vector nodal finite elements into the (high-order) Nedelec space. In other words,  $\mathbf{\Pi}$  is the (parallel) matrix representation of the interpolation mapping from  $P_k^3/Q_k^3$  into  $ND_k$ , see [HiXu2006], [KoVa2009]. We require this matrix as an input, since in the high-order case its entries very much depend on the particular choice of the basis functions in the edge and nodal spaces, as well as on the geometry of the mesh elements. The columns of  $\mathbf{\Pi}$  should use a node-based numbering, where the  $x/y/z$  components of the first node (vertex or high-order degree of freedom) should be listed first, followed by the  $x/y/z$  components of the second node and so on (see the documentation of `HYPRE_BoomerAMGSetDofFunc`).

Similarly to the Nedelec interpolation, the discrete gradient matrix  $G$  should correspond to the mapping  $\varphi \in P_k^3/Q_k^3 \mapsto \nabla\varphi \in ND_k$ , so even though its values are still independent of the mesh coordinates, they will not be  $\pm 1$ , but will be determined by the particular form of the high-order basis functions and degrees of freedom.

With these matrices, the high-order setup procedure is simply

```
HYPRE_AMSSetDimension(solver, dim);
HYPRE_AMSSetDiscreteGradient(solver, G);
HYPRE_AMSSetInterpolations(solver, Pi, NULL, NULL, NULL);
```

We remark that the above interface calls can also be used in the lowest-order case (or even other types of discretizations such as those based on the second family of Nedelec elements), but we recommend calling the previously described `HYPRE_AMSSetCoordinateVectors` instead, since this allows AMS to handle the construction and use of  $\mathbf{\Pi}$  internally.

Specifying the monolithic  $\mathbf{\Pi}$  limits the AMS cycle type options to those less than 10. Alternatively one can separately specify the  $x$ ,  $y$  and  $z$  components of  $\mathbf{\Pi}$ :

```
HYPRE_AMSSetInterpolations(solver, NULL, Pix, Piy, Piz);
```

which enables the use of AMS cycle types with index greater than 10. By definition,  $\mathbf{\Pi}^x\varphi = \mathbf{\Pi}(\varphi, 0, 0)$ , and similarly for  $\mathbf{\Pi}^y$  and  $\mathbf{\Pi}^z$ . Each of these matrices has the same sparsity pattern as  $G$ , but their entries depend on the coordinates of the mesh vertices.

Finally, both  $\mathbf{\Pi}$  and its components can be passed to the solver:

```
HYPRE_AMSSetInterpolations(solver, Pi, Pix, Piy, Piz);
```

which will duplicate some memory, but allows for experimentation with all available AMS cycle types.

### 5.8.4 Non-conforming AMR Grids

AMS could also be applied to problems with adaptive mesh refinement (AMR) posed on non-conforming quadrilateral/hexahedral meshes, see [GrKo2015] for more details.

On non-conforming grids (assuming also arbitrarily high-order elements), each finite element space has two versions: a conforming one, e.g.  $Q_k^c/ND_k^c$ , where the *hanging* degrees of freedom are constrained by the conforming (*real*) degrees of freedom, and a non-conforming one, e.g.  $Q_k^{nc}/ND_k^{nc}$  where the non-conforming degrees of freedom (hanging and real) are unconstrained. These spaces are related with the conforming prolongation and the pure restriction operators

$P$  and  $R$ , as well as the conforming and non-conforming version of the discrete gradient operator as follows:

$$\begin{array}{ccc}
 Q_k^c & \xrightarrow{G_c} & ND_k^c \\
 P_Q \uparrow & & \uparrow R_{ND} \\
 Q_k^{nc} & \xrightarrow{G_{nc}} & ND_k^{nc}
 \end{array}$$

Since the linear system is posed on  $ND_k^c$ , the user needs to provide the conforming discrete gradient matrix  $G_c$  to AMS, using `HYPRE_AMSSetDiscreteGradient`. This matrix is defined by the requirement that the above diagram commutes from  $Q_k^c$  to  $ND_k^{nc}$ , corresponding to the definition

$$G_c = R_{ND} G_{nc} P_Q,$$

i.e. the conforming gradient is computed by starting with a conforming nodal  $Q_k$  function, interpolating it in the hanging nodes, computing the gradient locally and representing it in the Nedelec space on each element (the non-conforming discrete gradient  $G_{nc}$  in the above formula), and disregarding the values in the hanging  $ND_k$  degrees of freedom.

Similar considerations imply that the conforming Nedelec interpolation matrix  $\Pi_c$  should be defined as

$$\Pi_c = R_{ND} \Pi_{nc} P_{Q^3},$$

with  $\Pi_{nc}$  computed element-wise as in the previous subsection. Note that in the low-order case,  $\Pi_c$  can be computed internally in AMS based only  $G_c$  and the conforming coordinates of the vertices  $x_c/y_c/z_c$ , see [GrKo2015].

## 5.9 ADS

The Auxiliary-space Divergence Solver (ADS) is a parallel unstructured solver similar to AMS, but targeting  $H(div)$  instead of  $H(curl)$  problems. Its usage and options are very similar to those of AMS, and in general the relationship between ADS and AMS is analogous to that between AMS and AMG.

Specifically ADS was designed for the scalable solution of linear systems arising in the finite element discretization of the variational problem

$$\text{Find } \mathbf{u} \in \mathbf{W}_h : \quad (\alpha \nabla \cdot \mathbf{u}, \nabla \cdot \mathbf{v}) + (\beta \mathbf{u}, \mathbf{v}) = (\mathbf{f}, \mathbf{v}), \quad \text{for all } \mathbf{v} \in \mathbf{W}_h, \quad (5.3)$$

where  $\mathbf{W}_h$  is the lowest order Raviart-Thomas (face) finite element space, and  $\alpha > 0$  and  $\beta > 0$  are scalar, or SPD matrix variable coefficients. It is based on the auxiliary space methods for  $H(div)$  problems proposed in [HiXu2006].

### 5.9.1 Overview

Let  $\mathbf{A}$  and  $\mathbf{b}$  be the stiffness matrix and the load vector corresponding to (5.3). Then the resulting linear system of interest reads,

$$\mathbf{A} \mathbf{x} = \mathbf{b}. \quad (5.4)$$

The coefficients  $\alpha$  and  $\beta$  are naturally associated with the ‘‘stiffness’’ and ‘‘mass’’ terms of  $\mathbf{A}$ . Besides  $\mathbf{A}$  and  $\mathbf{b}$ , ADS requires the following additional user input:

1. The discrete curl matrix  $C$  representing the faces of the mesh in terms of its edges.  $C$  has as many rows as the number of faces in the mesh, with each row having nonzero entries  $+1$  and  $-1$  in the columns corresponding to the edges composing the face. The sign is determined based on the orientation of the edges relative to the face. We require that  $C$  includes all (interior and boundary) faces and edges.
2. The discrete gradient matrix  $G$  representing the edges of the mesh in terms of its vertices.  $G$  has as many rows as the number of edges in the mesh, with each row having two nonzero entries:  $+1$  and  $-1$  in the columns corresponding to the vertices composing the edge. The sign is determined based on the orientation of the edge. We require that  $G$  includes all (interior and boundary) edges and vertices.

3. Vectors  $x$ ,  $y$ , and  $z$  representing the coordinates of the vertices of the mesh.

Internally, ADS proceeds with the construction of the following additional objects:

- $A_C$  – the curl-curl matrix  $C^T \mathbf{A} C$ .
- $\mathbf{\Pi}$  – the matrix representation of the interpolation operator from vector linear to face finite elements.
- $\mathbf{A}_{\mathbf{\Pi}}$  – the vector nodal matrix  $\mathbf{\Pi}^T \mathbf{A} \mathbf{\Pi}$ .
- $B_C$  and  $\mathbf{B}_{\mathbf{\Pi}}$  – efficient (AMS/AMG) solvers for  $A_C$  and  $\mathbf{A}_{\mathbf{\Pi}}$ .

The solution procedure then is a three-level method using smoothing in the original face space and subspace corrections based on  $B_C$  and  $\mathbf{B}_{\mathbf{\Pi}}$ . We can employ a number of options here utilizing various combinations of the smoother and solvers in additive or multiplicative fashion.

## 5.9.2 Sample Usage

ADS can be used either as a solver or as a preconditioner. Below we list the sequence of hypr calls needed to create and use it as a solver. We start with the allocation of the HYPRE\_Solver object:

```
HYPRE_Solver solver;
HYPRE_ADSCreate(&solver);
```

Next, we set a number of solver parameters. Some of them are optional, while others are necessary in order to perform the solver setup.

The user is required to provide the discrete curl and gradient matrices  $C$  and  $G$ . ADS expects a matrix defined on the whole mesh with no boundary faces, edges or nodes excluded. It is essential to **not** impose any boundary conditions on  $C$  or  $G$ . Regardless of which hypr conceptual interface was used to construct the matrices, one can always obtain a ParCSR version of them. This is the expected format in the following functions.

```
HYPRE_ADSSetDiscreteCurl(solver, C);
HYPRE_ADSSetDiscreteGradient(solver, G);
```

Next, ADS requires the coordinates of the vertices in the mesh as three hypr parallel vectors. The corresponding function call reads:

```
HYPRE_ADSSetCoordinateVectors(solver, x, y, z);
```

The remaining solver parameters are optional. For example, the user can choose a different cycle type by calling

```
HYPRE_ADSSetCycleType(solver, cycle_type); /* default value: 1 */
```

The available cycle types in ADS are:

- `cycle_type=1`: multiplicative solver (01210)
- `cycle_type=2`: additive solver (0 + 1 + 2)
- `cycle_type=3`: multiplicative solver (02120)
- `cycle_type=4`: additive solver (010 + 2)
- `cycle_type=5`: multiplicative solver (0102010)
- `cycle_type=6`: additive solver (1 + 020)
- `cycle_type=7`: multiplicative solver (0201020)
- `cycle_type=8`: additive solver (0(1 + 2)0)
- `cycle_type=11`: multiplicative solver (013454310)

- `cycle_type=12`: additive solver (0 + 1 + 3 + 4 + 5)
- `cycle_type=13`: multiplicative solver (034515430)
- `cycle_type=14`: additive solver (01(3 + 4 + 5)10)

Here we use the following convention for the three subspace correction methods: 0 refers to smoothing, 1 stands for AMS based on  $B_C$ , and 2 refers to a call to BoomerAMG for  $\mathbf{B}_\Pi$ . The values 3, 4 and 5 refer to the scalar subspaces corresponding to the  $x$ ,  $y$  and  $z$  components of  $\Pi$ .

The abbreviation  $xyyz$  for  $x, y, z \in \{0, 1, 2, 3, 4, 5\}$  refers to a multiplicative subspace correction based on solvers  $x$ ,  $y$ ,  $y$ , and  $z$  (in that order). The abbreviation  $x+y+z$  stands for an additive subspace correction method based on  $x$ ,  $y$  and  $z$  solvers. The additive cycles are meant to be used only when ADS is called as a preconditioner. In our experience the choices `cycle_type=1, 5, 8, 11, 13` often produced fastest solution times, while `cycle_type=7` resulted in smallest number of iterations.

Additional solver parameters, such as the maximum number of iterations, the convergence tolerance and the output level, can be set with

```
HYPRE_ADSSetMaxIter(solver, maxit);      /* default value: 20 */
HYPRE_ADSSetTol(solver, tol);           /* default value: 1e-6 */
HYPRE_ADSSetPrintLevel(solver, print);  /* default value: 1 */
```

More advanced parameters, affecting the smoothing and the internal AMS and AMG solvers, can be set with the following three functions:

```
HYPRE_ADSSetSmoothingOptions(solver, 2, 1, 1.0, 1.0);
HYPRE_ADSSetAMSOptions(solver, 11, 10, 1, 3, 0.25, 0, 0);
HYPRE_ADSSetAMGOptions(solver, 10, 1, 3, 0.25, 0, 0);
```

We note that the AMS cycle type, which is the second parameter of `HYPRE_ADSSetAMSOptions` should be greater than 10, unless the high-order interface of `HYPRE_ADSSetInterpolations` described in the next subsection is being used.

After the above calls, the solver is ready to be constructed. The user has to provide the stiffness matrix  $\mathbf{A}$  (in ParCSR format) and the hypr parallel vectors  $\mathbf{b}$  and  $\mathbf{x}$ . (The vectors are actually not used in the current ADS setup.) The setup call reads,

```
HYPRE_ADSSetup(solver, A, b, x);
```

It is important to note the order of the calling sequence. For example, do **not** call `HYPRE_ADSSetup` before calling each of the functions `HYPRE_ADSSetDiscreteCurl`, `HYPRE_ADSSetDiscreteGradient` and `HYPRE_ADSSetCoordinateVectors`.

Once the setup has completed, we can solve the linear system by calling

```
HYPRE_ADSSolve(solver, A, b, x);
```

Finally, the solver can be destroyed with

```
HYPRE_ADSDestroy(&solver);
```

More details can be found in the files `ads.h` and `ads.c` located in the `parcsr_ls` directory.

### 5.9.3 High-order Discretizations

Similarly to AMS, ADS also provides support for (arbitrary) high-order  $H(\text{div})$  discretizations. Since the robustness of ADS depends on the performance of AMS and BoomerAMG on the associated (high-order) auxiliary subspace problems, we note that the convergence may not be optimal for large polynomial degrees  $k \geq 1$ .

In the high-order ADS interface, the user does not need to provide the coordinates of the vertices, but instead should construct and pass the Raviart-Thomas and Nedelec interpolation matrices  $\mathbf{\Pi}_{RT}$  and  $\mathbf{\Pi}_{ND}$  which map (high-order) vector nodal finite elements into the (high-order) Raviart-Thomas and Nedelec space. In other words, these are the (parallel) matrix representation of the interpolation mappings from  $P_k^3/Q_k^3$  into  $RT_{k-1}$  and  $ND_k$ , see [HiXu2006], [KoVa2009]. We require these matrices as inputs, since in the high-order case their entries very much depend on the particular choice of the basis functions in the finite element spaces, as well as on the geometry of the mesh elements. The columns of the  $\mathbf{\Pi}$  matrices should use a node-based numbering, where the  $x/y/z$  components of the first node (vertex or high-order degree of freedom) should be listed first, followed by the  $x/y/z$  components of the second node and so on (see the documentation of `HYPRE_BoomerAMGSetDofFunc`). Furthermore, each interpolation matrix can be split into  $x$ ,  $y$  and  $z$  components by defining  $\mathbf{\Pi}^x \varphi = \mathbf{\Pi}(\varphi, 0, 0)$ , and similarly for  $\mathbf{\Pi}^y$  and  $\mathbf{\Pi}^z$ .

The discrete gradient and curl matrices  $G$  and  $C$  should correspond to the mappings  $\varphi \in P_k^3/Q_k^3 \mapsto \nabla \varphi \in ND_k$  and  $\mathbf{u} \in ND_k \mapsto \nabla \times \mathbf{u} \in RT_{k-1}$ , so even though their values are still independent of the mesh coordinates, they will not be  $\pm 1$ , but will be determined by the particular form of the high-order basis functions and degrees of freedom.

With these matrices, the high-order setup procedure is simply

```
HYPRE_ADSSetDiscreteCurl(solver, C);
HYPRE_ADSSetDiscreteGradient(solver, G);
HYPRE_ADSSetInterpolations(solver, RT_Pi, NULL, NULL, NULL,
                           ND_Pi, NULL, NULL, NULL);
```

We remark that the above interface calls can also be used in the lowest-order case (or even other types of discretizations), but we recommend calling the previously described `HYPRE_ADSSetCoordinateVectors` instead, since this allows ADS to handle the construction and use of the interpolations internally.

Specifying the monolithic  $\mathbf{\Pi}_{RT}$  limits the ADS cycle type options to those less than 10. Alternatively one can separately specify the  $x$ ,  $y$  and  $z$  components of  $\mathbf{\Pi}_{RT}$ .

```
HYPRE_ADSSetInterpolations(solver, NULL, RT_Pix, RT_Piy, RT_Piz,
                           ND_Pi, NULL, NULL, NULL);
```

which enables the use of ADS cycle types with index greater than 10. The same holds for  $\mathbf{\Pi}_{ND}$  and its components, e.g. to enable the subspace AMS cycle type greater than 10 we need to call

```
HYPRE_ADSSetInterpolations(solver, NULL, RT_Pix, RT_Piy, RT_Piz,
                           NULL, ND_Pix, ND_Piy, ND_Piz);
```

Finally, both  $\mathbf{\Pi}$  and their components can be passed to the solver:

```
HYPRE_ADSSetInterpolations(solver, RT_Pi, RT_Pix, RT_Piy, RT_Piz,
                           ND_Pi, ND_Pix, ND_Piy, ND_Piz);
```

which will duplicate some memory, but allows for experimentation with all available ADS and AMS cycle types.

## 5.10 Multigrid Reduction (MGR)

MGR is a parallel multigrid reduction solver and preconditioner designed to take advantage of use-provided information to solve systems of equations with multiple variable types. The algorithm is similar to two-stage preconditioner strategies and other reduction techniques like ARMS, but in a standard multigrid framework.

The MGR algorithm accepts information about the variables in block form from the user and uses it to define the appropriate C/F splitting for the multigrid scheme. The linear system solve proceeds with an F-relaxation solve on the F points, followed by a coarse grid correction. The coarse grid solve is handled by scalar AMG (BoomerAMG). MGR provides users with more control over the coarsening process, and can potentially be a starting point for designing multigrid-based physics-based preconditioners.

The following represents a minimal set of functions, and some optional functions, to call to use the MGR solver. For simplicity, we ignore the function parameters here, and refer the reader to the reference manual for more details on the parameters and their defaults.

- `HYPRE_MGRCreate`: Create the MGR solver object.
- `HYPRE_MGRSetCpointsByBlock`: Set up block data with information about coarse indexes for reduction. Here, the user specifies the number of reduction levels, as well as the coarse nodes for each level of the reduction. These coarse nodes are indexed by their index in the block of unknowns. This is used internally to tag the appropriate indexes of the linear system matrix as coarse nodes.
- (Optional) `HYPRE_MGRSetReservedCoarseNodes`: Prescribe a subset of nodes to be kept as coarse nodes until the coarsest level. These nodes are transferred onto the coarsest grid of the BoomerAMG coarse grid solver.
- (Optional) `HYPRE_MGRSetNonCpointsToFpoints`: Set points not prescribed as C points to be fixed as F points for intermediate levels. Setting this to 1 uses the user input to define the C/F splitting. Otherwise, a BoomerAMG coarsening routine is used to determine the C/F splitting for intermediate levels.
- (Optional) `HYPRE_MGRSetCoarseSolver`: This function sets the BoomerAMG solver to be used for the solve on the coarse grid. The user can define their own BoomerAMG solver with their preferred options and pass this to the MGR solver. Otherwise, an internal BoomerAMG solver is used as the coarse grid solver instead.
- `HYPRE_MGRSetup`: Setup and MGR solver object.
- `HYPRE_MGRSolve`: Solve the linear system.
- `HYPRE_MGRDestroy`: Destroy the MGR solver object

For more details about additional solver options and parameters, please refer to the reference manual. NOTE: The MGR solver is currently only supported by the IJ interface.

## 5.11 FSAI

FSAI is a parallel implementation of the Factorized Sparse Approximate Inverse preconditioner, initially proposed by [KoYe1993]. Given a symmetric positive definite matrix  $A$ , FSAI computes a triangular matrix  $G$  that approximates the inverse of the lower Cholesky factor ( $L$ ) of  $A$ . This computation is done by minimizing the Frobenius norm  $\|I - GL\|_F$  without explicit knowledge of  $L$ . The resulting preconditioner preserves the positive definiteness of  $A$  and is given by  $M^{-1} = G^T G$ .

One of the critical factors determining the quality of sparse approximate inverse preconditioners lies in choosing the sparsity pattern of  $G$ . While ParaSails employs *a priori* sparsity patterns, FSAI uses an iterative strategy that generates sparsity patterns on the fly, i.e., while computing their nonzero coefficient values concurrently. At every step of the iterative process, the sparsity pattern of a row of  $G$  is augmented with a fixed number of entries, leading to the most significant reduction of the conditioning number of  $GAG^T$ . Such a strategy is also called “adaptive FSAI” or “dynamic FSAI” and it can lead to more robust sparse approximate inverses than ParaSails. For more details on how it works, see [JaFe2015].

### 5.11.1 Parameter Settings

The accuracy and cost of FSAI are determined by three configurations parameters as shown in the table below

param	type	range	sug. values	default
max_steps	int	$\geq 0$	5, 10, 30	5
max_step_size	int	$\geq 0$	1, 3, 6	3
kap_tolerance	real	$\geq 0$	0.0, 1.E-2, 1.E-3	1.E-3

The first parameter, `max_steps`, controls the number of maximum steps used in the iterative algorithm. The second parameter, `max_step_size`, gives the maximum number of indices added to the sparsity pattern of  $G$  at each step. Lastly, the third parameter, `kap_tolerance`, is a floating-point value used to stop the inclusion of new indices to the sparsity pattern of  $G$  when the conditioning number of  $GAG^T$  stagnates. This can be disabled by setting `kap_tolerance = 0`. Naturally, the preconditioner quality increases for denser sparsity patterns of  $G$ , but so do its setup and solve costs. For a reasonable balance between accuracy and cost, we recommend that  $max\_steps * max\_step\_size \leq 30$ . The configuration parameters of FSAI can be set via the following calls:

```
HYPRE_FSAISetMaxSteps(HYPRE_Solver solver, HYPRE_Int max_steps);
HYPRE_FSAISetMaxStepSize(HYPRE_Solver solver, HYPRE_Int max_step_size);
HYPRE_FSAISetKapTolerance(HYPRE_Solver solver, HYPRE_Real kap_tolerance);
```

### 5.11.2 FSAI as Smoother to BoomerAMG

As discussed in [PaFa2019], the factorized sparse approximate inverse method can be an effective smoother to AMG for several reasons. Particularly, it leads to a symmetric operator, and thus allows AMG to be used as a preconditioner for the conjugate gradient solver. In `hypr`, FSAI can be used as a complex smoother to BoomerAMG by calling the functions:

```
HYPRE_BoomerAMGSetSmoothType(HYPRE_Solver solver, 4);
HYPRE_BoomerAMGSetSmoothNumLevels(HYPRE_Solver solver, HYPRE_Int num_levels);
```

where `num_levels` is the last multigrid level where FSAI is used. The configuration parameters of the FSAI smoother, as described above, can be set via the following calls:

```
HYPRE_BoomerAMGSetFSAIMaxSteps(HYPRE_Solver solver, HYPRE_Int max_steps);
HYPRE_BoomerAMGSetFSAIMaxStepSize(HYPRE_Solver solver, HYPRE_Int max_step_size);
HYPRE_BoomerAMGSetFSAIKapTolerance(HYPRE_Solver solver, HYPRE_Real kap_tolerance);
```

### 5.11.3 Implementation Notes

- When the matrix  $A$  is distributed across MPI tasks, FSAI considers only the block diagonal portions of  $A$  for computing  $G$ . The resulting preconditioner is effectively a block-Jacobi sparse approximate inverse in the MPI sense. Although this strategy reduces communication costs, it can degrade convergence performance when several tasks are used, especially when FSAI is employed as a preconditioner to a Krylov solver.
- The CPU version of FSAI supports threading via OpenMP. To enable it, users need to compile `hypr` with OpenMP support via the configure option `--with-openmp`. In this case, FSAI relies on an implementation of BLAS/LAPACK that is thread-safe. The one distributed internally with `hypr` fulfills this criterion, but care must be taken when linking `hypr` to external BLAS/LAPACK libraries. In HPC platforms, we recommend using vendor implementations of BLAS/LAPACK for better setup performance of FSAI, regardless of whether using OpenMP or not.
- The GPU version of FSAI is under development.

## 5.12 ParaSails

### Warning

ParaSails is not actively supported by the hypr development team. We recommend using *FSAI* for parallel sparse approximate inverse algorithms. This new implementation includes NVIDIA/AMD GPU support through the CUDA/HIP backends.

ParaSails is a parallel implementation of a sparse approximate inverse preconditioner, using *a priori* sparsity patterns and least-squares (Frobenius norm) minimization. Symmetric positive definite (SPD) problems are handled using a factored SPD sparse approximate inverse. General (nonsymmetric and/or indefinite) problems are handled with an unfactored sparse approximate inverse. It is also possible to precondition nonsymmetric but definite matrices with a factored, SPD preconditioner.

ParaSails uses *a priori* sparsity patterns that are patterns of powers of sparsified matrices. ParaSails also uses a post-filtering technique to reduce the cost of applying the preconditioner. In advanced usage not described here, the pattern of the preconditioner can also be reused to generate preconditioners for different matrices in a sequence of linear solves.

For more details about the ParaSails algorithm, see [Chow2000].

### 5.12.1 Parameter Settings

The accuracy and cost of ParaSails are parametrized by the real `thresh` and integer `nlevels` parameters,  $0 \leq \text{thresh} \leq 1$ ,  $0 \leq \text{nlevels}$ . Lower values of `thresh` and higher values of `nlevels` lead to more accurate, but more expensive preconditioners. More accurate preconditioners are also more expensive per iteration. The default values are `thresh = 0.1` and `nlevels = 1`. The parameters are set using `HYPRE_ParaSailsSetParams`.

Mathematically, given a symmetric matrix  $A$ , the pattern of the approximate inverse is the pattern of  $\tilde{A}^m$  where  $\tilde{A}$  is a matrix that has been sparsified from  $A$ . The sparsification is performed by dropping all entries in a symmetrically diagonally scaled  $A$  whose values are less than `thresh` in magnitude. The parameter `nlevel` is equivalent to  $m - 1$ . Filtering is a post-thresholding procedure. For more details about the algorithm, see [Chow2000].

The storage required for the ParaSails preconditioner depends on the parameters `thresh` and `nlevels`. The default parameters often produce a preconditioner that can be stored in less than the space required to store the original matrix. ParaSails does not need a large amount of intermediate storage in order to construct the preconditioner.

ParaSail's `Create` function differs from the synopsis in the following way:

```
int HYPRE_ParaSailsCreate(MPI_Comm comm, HYPRE_Solver *solver, int symmetry);
```

where `comm` is the MPI communicator.

The value of `symmetry` has the following meanings, to indicate the symmetry and definiteness of the problem, and to specify the type of preconditioner to construct:

value	meaning
0	nonsymmetric and/or indefinite problem, and nonsymmetric preconditioner
1	SPD problem, and SPD (factored) preconditioner
2	nonsymmetric, definite problem, and SPD (factored) preconditioner

For more information about the final case, see section *Preconditioning Nearly Symmetric Matrices*.

Parameters for setting up the preconditioner are specified using

```
int HYPRE_ParaSailsSetParams(HYPRE_Solver solver, double thresh,
                             int nlevel, double filter);
```

The parameters are used to specify the sparsity pattern and filtering value (see above), and are described with suggested values as follows:

param	type	range	sug. values	default	meaning
nlevel	integer	$\geq 0$	0, 1, 2	1	$m = 1 + \text{nlevel}$
thresh	real	$\geq 0$	0, 0.1, 0.01	0.1	thresh = thresh
		$< 0$	-0.75, -0.90		thresh auto-selected
filter	real	$\geq 0$	0, 0.05, 0.001	0.05	filter = filter
		$< 0$	-0.90		filter auto-selected

When `thresh`  $< 0$ , then a threshold is selected such that `thresh` represents the negative of the fraction of nonzero elements that are dropped. For example, if `thresh` =  $-0.9$  then  $\tilde{A}$  will contain approximately ten percent of the nonzeros in  $A$ .

When `filter`  $< 0$ , then a filter value is selected such that `filter` represents the negative of the fraction of nonzero elements that are dropped. For example, if `filter` =  $-0.9$  then approximately 90 percent of the entries in the computed approximate inverse are dropped.

## 5.12.2 Preconditioning Nearly Symmetric Matrices

A nonsymmetric, but definite and nearly symmetric matrix  $A$  may be preconditioned with a symmetric preconditioner  $M$ . Using a symmetric preconditioner has a few advantages, such as guaranteeing positive definiteness of the preconditioner, as well as being less expensive to construct.

The nonsymmetric matrix  $A$  must be definite, i.e.,  $(A + A^T)/2$  is SPD, and the *a priori* sparsity pattern to be used must be symmetric. The latter may be guaranteed by 1) constructing the sparsity pattern with a symmetric matrix, or 2) if the matrix is structurally symmetric (has symmetric pattern), then thresholding to construct the pattern is not used (i.e., zero value of the `thresh` parameter is used).

## 5.13 ILU

ILU is a suite of parallel incomplete LU factorization algorithms featuring dual threshold (ILUT) and level-based (ILUK) variants. The implementation is based on a domain decomposition framework for achieving distributed parallelism. ILU can be used as a standalone iterative solver (this is not recommended), preconditioner for Krylov subspace methods, or smoother for multigrid methods such as BoomerAMG and MGR.

### Note

ILU is currently only supported by the IJ interface.

### 5.13.1 Overview

ILU utilizes a domain decomposition framework. A basic block-Jacobi approach involves performing inexact solutions within the local domains owned by the processes, using parallel local ILU factorizations. In a more advanced approach, the unknowns are partitioned into interior and interface points, where the interface points separate the interior points in adjacent domains. In an algebraic context, this is equivalent to dividing the matrix rows into local (processor-owned) and external (off-processor-owned) blocks. The overall parallel ILU strategy is a two-level method that consists of ILU solves within the local blocks and a global solve involving the Schur complement system, which various iterative approaches in this framework can solve.

### 5.13.2 User-level functions

A list of user-level functions for configuring ILU is given below, where each block of functions is marked as *Required*, *Recommended*, *Optional*, or *Exclusively required*. Note that the last two blocks of function calls are exclusively required, i.e., the first block should be called only when ILU is used as a standalone solver, while the second block should be called only when it is used as a preconditioner to GMRES. In the last case, other Krylov methods can be chosen. We refer the reader to *Solvers and Preconditioners* for more information.

```

/* (Required) Create ILU solver */
HYPRE_ILUCreate(&ilu_solver);

/* (Recommended) General solver options */
HYPRE_ILUSetType(ilu_solver, ilu_type); /* 0, 1, 10, 11, 20, 21, 30, 31, 40, 41, 50 */
HYPRE_ILUSetMaxIter(ilu_solver, max_iter);
HYPRE_ILUSetTol(ilu_solver, tol);
HYPRE_ILUSetLocalReordering(ilu_solver, reordering); /* 0: none, 1: RCM */
HYPRE_ILUSetPrintLevel(ilu_solver, print_level);

/* (Optional) Function calls for ILUK variants */
HYPRE_ILUSetLevelOfFill(ilu_solver, fill);

/* (Optional) Function calls for ILUT variants */
HYPRE_ILUSetMaxNnzPerRow(ilu_solver, max_nnz_row);
HYPRE_ILUSetDropThreshold(ilu_solver, threshold);

/* (Optional) Function calls for GMRES-ILU or NSH-ILU */
HYPRE_ILUSetNSHDropThreshold(ilu_solver, threshold);
HYPRE_ILUSetSchurMaxIter(ilu_solver, schur_max_iter);

/* (Optional) Function calls for iterative ILU variants */
HYPRE_ILUSetTriSolve(ilu_solver, 0);
HYPRE_ILUSetLowerJacobiIters(ilu_solver, ljac_iters);
HYPRE_ILUSetUpperJacobiIters(ilu_solver, ujac_iters);

/* (Exclusively required) Function calls for using ILU as standalone solver */
HYPRE_ILUSetup(ilu_solver, parcsr_M, b, x);
HYPRE_ILUSolve(ilu_solver, parcsr_A, b, x);

/* (Exclusively required) Function calls for using ILU as preconditioner to GMRES */
HYPRE_GMRESSetup(gmres_solver, (HYPRE_Matrix)A, (HYPRE_Vector)b, (HYPRE_Vector)x);
HYPRE_GMRESSolve(gmres_solver, (HYPRE_Matrix)A, (HYPRE_Vector)b, (HYPRE_Vector)x);

/* (Required) Free memory */
HYPRE_ILUDestroy(ilu_solver);

```

A short explanation for each of those functions and its parameters is given next.

- HYPRE\_ILUCreate Create the hypre\_ILU solver object.
- HYPRE\_ILUDestroy Destroy the hypre\_ILU solver object.
- HYPRE\_ILUSetType Set the type of ILU factorization. Options are:
  - 0: Block-Jacobi ILUK (BJ-ILUK).
  - 1: Block-Jacobi ILUT (BJ-ILUT).
  - 10: GMRES with ILUK (GMRES-ILUK).

- 11: GMRES with ILUT (GMRES-ILUT).
  - 20: NSH with ILUK (NSH-ILUK).
  - 21: NSH with ILUT (NSH-ILUT).
  - 30: RAS with ILUK (RAS-ILUK).
  - 31: RAS with ILUT (RAS-ILUT).
  - 40: ddPQ-GMRES with ILUK (ddPQ-GMRES-ILUK).
  - 41: ddPQ-GMRES with ILUT (ddPQ-GMRES-ILUT).
  - 50: GMRES with RAP-ILU0 with modified ILU0 (GMRES-RAP-ILU0).
- `HYPRE_ILUSetMaxIter` Set the maximum number of ILU iterations. We recommend setting this value to one when ILU is used as a preconditioner or smoother.
  - `HYPRE_ILUSetTol` Set the convergence tolerance for ILU. We recommend setting this value to zero when ILU is used as a preconditioner or smoother.
  - `HYPRE_ILUSetLocalReordering` Set the local matrix reordering algorithm.
    - 0: No reordering.
    - 1: Reverse Cuthill–McKee (RCM).
  - `HYPRE_ILUSetPrintLevel` Set the verbosity level for algorithm statistics.
    - 0: No output.
    - 1: Print setup info.
    - 2: Print solve info.
    - 3: Print setup and solve info.
  - `HYPRE_ILUSetLevelOfFill` Set the level of fill used by the level-based ILUK strategy.
  - `HYPRE_ILUSetMaxNnzPerRow` Set the maximum number of nonzero entries per row in the triangular factors for ILUT.
  - `HYPRE_ILUSetDropThreshold` Set the threshold for dropping nonzero entries during the construction of the triangular factors for ILUT.
  - `HYPRE_ILUSetNSHDropThreshold` Set the threshold for dropping nonzero entries during the computation of the approximate inverse matrix via NSH-ILU.
  - `HYPRE_ILUSetSchurMaxIter` Set the maximum number of iterations for solving the Schur complement system (GMRES-ILU or NSH-ILU).
  - `HYPRE_ILUSetTriSolve` Set triangular solve method used in ILU's solve phase. Option zero refers to the iterative approach, which leads to good performance in GPUs, and option one refers to the direct (exact) approach.
  - `HYPRE_ILUSetLowerJacobiIters` Set the number of iterations for solving the lower triangular linear system. This option makes sense when enabling the iterative triangular solve approach.
  - `HYPRE_ILUSetUpperJacobiIters` Same as previous function, but for the upper triangular factor.
  - `HYPRE_ILUSetup` Setup a `hypr_ILU` solver object.
  - `HYPRE_ILUSolve` Solve the linear system with `hypr_ILU`.
  - `HYPRE_ILUDestroy` Destroy the `hypr_ILU` solver object.

**Note**

For more details about ILU options and parameters, including their default values, we refer the reader to hypr's reference manual or section *ParCSR Solvers*.

### 5.13.3 ILU as Smoother for BoomerAMG

The following functions can be used to configure ILU as a smoother to BoomerAMG:

```

/* (Required) Set ILU as smoother to BoomerAMG */
HYPRE_BoomerAMGSetSmoothType(amg_solver, 5);
HYPRE_BoomerAMGSetSmoothNumLevels(amg_solver, num_levels);

/* (Optional) General ILU configuration parameters */
HYPRE_BoomerAMGSetILUType(amg_solver, ilu_type);
HYPRE_BoomerAMGSetILUMaxIter(amg_solver, ilu_max_iter);
HYPRE_BoomerAMGSetILULocalReordering(amg_solver, ilu_reordering);

/* (Optional) Function calls for ILUK smoother variants */
HYPRE_BoomerAMGSetILULevel(amg_solver, ilu_fill);

/* (Optional) Function calls for ILUT smoother variants */
HYPRE_BoomerAMGSetILUDroptol(amg_solver, ilu_threshold);
HYPRE_BoomerAMGSetILUMaxRowNnz(amg_solver, ilu_max_nnz_row);

/* (Optional) Function calls for iterative ILU smoother variants */
HYPRE_BoomerAMGSetILUTriSolve(amg_solver, 0);
HYPRE_BoomerAMGSetILULowerJacobiIters(amg_solver, ilu_ljac_iters);
HYPRE_BoomerAMGSetILUUpperJacobiIters(amg_solver, ilu_ujac_iters);

```

where:

- `HYPRE_BoomerAMGSetSmoothNumLevels` Enable smoothing in the first `num_levels` levels of AMG.
- `HYPRE_BoomerAMGSetILUType` Set the type of ILU factorization. See `HYPRE_ILUSetType`.
- `HYPRE_BoomerAMGSetILUMaxIter` Set the number of ILU smoother sweeps.
- `HYPRE_BoomerAMGSetILULocalReordering` Set the local matrix reordering algorithm.
- `HYPRE_BoomerAMGSetILULevel` Set ILUK's fill level.
- `HYPRE_BoomerAMGSetILUDroptol` Set ILUT's threshold.
- `HYPRE_BoomerAMGSetILUMaxRowNnz` Set ILUT's maximum number of nonzero entries per row.
- `HYPRE_BoomerAMGSetILUTriSolve` Set triangular solve method. See `HYPRE_ILUSetTriSolve`.
- `HYPRE_BoomerAMGSetILULowerJacobiIters` Set the number of iterations for the L factor.
- `HYPRE_BoomerAMGSetILUUpperJacobiIters` Same as previous function, but for the U factor.

### 5.13.4 GPU support

The addition of GPU support to ILU is ongoing work. A few algorithm types have already been fully ported to the CUDA and HIP backends, i.e., both their setup (factorization) and solve phases are executed on the device. Below is a detailed list of which phases (setup and solve) of the various ILU algorithms have been ported to GPUs. In the table, *UVM-Setup* indicates that the setup phase is executed on the CPU (host); at the same time, the triangular factors are

stored in a memory space that is accessible from the GPU (device) via unified memory. This feature must be enabled during hypr's configuration.

	CUDA (NVIDIA GPUs)	HIP (AMD GPUs)	SYCL (Intel GPUs)
<b>BJ-ILU0</b>	Setup and Solve	Setup and Solve	UVM-Setup and Solve
<b>BJ-ILU(K/T)</b>	UVM-Setup and Solve	UVM-Setup and Solve	UVM-Setup and Solve
<b>GMRES-ILU0</b>	Setup and Solve	Setup and Solve	UVM-Setup and Solve
<b>GMRES-RAP-ILU0</b>	UVM-Setup and Solve	UVM-Setup and Solve	UVM-Setup and Solve
<b>GMRES-ILU(K/T)</b>	UVM-Setup and Solve	UVM-Setup and Solve	UVM-Setup and Solve
<b>ddPQ-GMRES-ILU(K/T)</b>	UVM-Setup and Solve	UVM-Setup and Solve	UVM-Setup and Solve
<b>NSH-ILU(K/T)</b>	UVM-Setup and Solve	UVM-Setup and Solve	UVM-Setup and Solve
<b>RAS-ILU(K/T)</b>	UVM-Setup and Solve	UVM-Setup and Solve	UVM-Setup and Solve

#### Hint

For better setup performance on GPUs, disable local reordering by passing option zero to `HYPRE_ILUSetLocalReordering` or `HYPRE_BoomerAMGSetILULocalReordering`. This may degrade convergence of the iterative solver.

#### Note

hypr must be built with cuSPARSE support when running ILU on NVIDIA GPUs, rocSPARSE when running on AMD GPUs, or oneMKL sparse when running on Intel GPUs.

## 5.14 Euclid

#### Warning

Euclid is not actively supported by the hypr development team. We recommend using *ILU* for parallel ILU algorithms. This new ILU implementation includes 64-bit integers support (for linear systems with more than 2,147,483,647 global unknowns) through both *mixedint* and *bigint* builds of hypr and NVIDIA/AMD GPUs support through the CUDA/HIP backends.

The Euclid library is a scalable implementation of the Parallel ILU algorithm that was presented at SC99 [HyPo1999], and published in expanded form in the SIAM Journal on Scientific Computing [HyPo2001]. By *scalable* we mean that the factorization (setup) and application (triangular solve) timings remain nearly constant when the global problem size is scaled in proportion to the number of processors. As with all ILU preconditioning methods, the number of iterations is expected to increase with global problem size.

Experimental results have shown that PILU preconditioning is in general more effective than Block Jacobi preconditioning for minimizing total solution time. For scaled problems, the relative advantage appears to increase as the number of processors is scaled upwards. Euclid may also be used to good advantage as a smoother within multigrid methods.

### 5.14.1 Overview

Euclid is best thought of as an “extensible ILU preconditioning framework.” *Extensible* means that Euclid can (and eventually will, time and contributing agencies permitting) support many variants of ILU( $k$ ) and ILUT preconditioning. (The current release includes Block Jacobi ILU( $k$ ) and Parallel ILU( $k$ ) methods.) Due to this extensibility, and also because Euclid was developed independently of the hypr project, the methods by which one passes runtime parameters to Euclid preconditioners differ in some respects from the hypr norm. While users can directly set options within their code, options can also be passed to Euclid preconditioners via command line switches and/or small text-based configuration files. The latter strategies have the advantage that users will not need to alter their codes as Euclid’s capabilities are extended.

The following fragment illustrates the minimum coding required to invoke Euclid preconditioning within hypr application contexts. The next subsection provides examples of the various ways in which Euclid’s options can be set. The final subsection lists the options, and provides guidance as to the settings that (in our experience) will likely prove effective for minimizing execution time.

```
#include "HYPRE_parcsr_ls.h"

HYPRE_Solver eu;
HYPRE_Solver pcg_solver;
HYPRE_ParVector b, x;
HYPRE_ParCSRMatrix A;

//Instantiate the preconditioner.
HYPRE_EuclidCreate(comm, &eu);

//Optionally use the following three methods to set runtime options.
// 1. pass options from command line or string array.
HYPRE_EuclidSetParams(eu, argc, argv);

// 2. pass options from a configuration file.
HYPRE_EuclidSetParamsFromFile(eu, "filename");

// 3. pass options using interface functions.
HYPRE_EuclidSetLevel(eu, 3);
...

//Set Euclid as the preconditioning method for some
//other solver, using the function calls HYPRE_EuclidSetup
//and HYPRE_EuclidSolve. We assume that the pcg_solver
//has been properly initialized.
HYPRE_PCGSetPrecond(pcg_solver,
                    (HYPRE_PtrToSolverFcn) HYPRE_EuclidSolve,
                    (HYPRE_PtrToSolverFcn) HYPRE_EuclidSetup,
                    eu);

//Solve the system by calling the Setup and Solve methods for,
//in this case, the HYPRE_PCG solver. We assume that A, b, and x
//have been properly initialized.
HYPRE_PCGSetup(pcg_solver, (HYPRE_Matrix)A, (HYPRE_Vector)b, (HYPRE_Vector)x);
HYPRE_PCGSolve(pcg_solver, (HYPRE_Matrix)parcsr_A, (HYPRE_Vector)b, (HYPRE_Vector)x);

//Destroy the Euclid preconditioning object.
HYPRE_EuclidDestroy(eu);
```

### 5.14.2 Setting Options: Examples

For expositional purposes, assume you wish to set the  $ILU(k)$  factorization level to the value  $k = 3$ . There are several methods of accomplishing this. Internal to Euclid, options are stored in a simple database that contains (name, value) pairs. Various of Euclid's internal (private) functions query this database to determine, at runtime, what action the user has requested. If you enter the option `-eu_stats 1`, a report will be printed when Euclid's destructor is called; this report lists (among other statistics) the options that were in effect during the factorization phase.

**Method 1.** By default, Euclid always looks for a file titled `database` in the working directory. If it finds such a file, it opens it and attempts to parse it as a configuration file. Configuration files should be formatted as follows.

```
>cat database
#this is an optional comment
-level 3
```

Any line in a configuration file that contains a “#” character in the first column is ignored. All other lines should begin with an option *name*, followed by one or more blanks, followed by the option *value*. Note that option names always begin with a - character. If you include an option name that is not recognized by Euclid, no harm should ensue.

**Method 2.** To pass options on the command line, call

```
HYPRE_EuclidSetParams(HYPRE_Solver solver, int argc, char *argv[]);
```

where `argc` and `argv` carry the usual connotation: `main(int argc, char *argv[])`. If your hypr application is called `phoo`, you can then pass options on the command line per the following example.

```
mpirun -np 2 phoo -level 3
```

Since Euclid looks for the `database` file when `HYPRE_EuclidCreate` is called, and parses the command line when `HYPRE_EuclidSetParams` is called, option values passed on the command line will override any similar settings that may be contained in the `database` file. Also, if same option name appears more than once on the command line, the final appearance determines the setting.

Some options, such as `-bj` (see next subsection) are boolean. Euclid always treats these options as the value 1 (true) or 0 (false). When passing boolean options from the command line the value may be committed, in which case it assumed to be 1. Note, however, that when boolean options are contained in a configuration file, either the 1 or 0 must stated explicitly.

**Method 3.** There are two ways in which you can read in options from a file whose name is other than `database`. First, you can call `HYPRE_EuclidSetParamsFromFile` to specify a configuration filename. Second, if you have passed the command line arguments as described above in Method 2, you can then specify the configuration filename on the command line using the `-db_filename filename` option, e.g.,

```
mpirun -np 2 phoo -db_filename ../myConfigFile
```

**Method 4.** One can also set parameters via interface functions, e.g

```
int HYPRE_EuclidSetLevel(HYPRE_Solver solver, int level);
```

For a full set of functions, see the reference manual.

### 5.14.3 Options Summary

- **-level** *<int>* Factorization level for  $ILU(k)$ . Default: 1. Guidance: for 2D convection-diffusion and similar problems, fastest solution time is typically obtained with levels 4 through 8. For 3D problems fastest solution time is typically obtained with level 1.

- **-bj** Use Block Jacobi ILU preconditioning instead of PILU. Default: 0 (false). Guidance: if subdomains contain relatively few nodes (less than 1,000), or the problem is not well partitioned, Block Jacobi ILU may give faster solution time than PILU.
- **-eu\_stats** When Euclid's destructor is called a summary of runtime settings and timing information is printed to stdout. Default: 0 (false). The timing marks in the report are the maximum over all processors in the MPI communicator.
- **-eu\_mem** When Euclid's destructor is called a summary of Euclid's memory usage is printed to stdout. Default: 0 (false). The statistics are for the processor whose rank in `MPI_COMM_WORLD` is 0.
- **-printTestData** This option is used in our autotest procedures, and should not normally be invoked by users.
- **-sparseA**  $\langle float \rangle$  Drop-tolerance for  $ILU(k)$  factorization. Default: 0 (no dropping). Entries are treated as zero if their absolute value is less than `sparseA * max`, where `max` is the largest absolute value of any entry in the row. Guidance: try this in conjunction with `-rowScale`. CAUTION: If the coefficient matrix  $A$  is symmetric, this setting is likely to cause the filled matrix,  $F = L + U - I$ , to be non-symmetric. This setting has no effect when ILUT factorization is selected.
- **-rowScale** Scale values prior to factorization such that the largest value in any row is +1 or -1. Default: 0 (false). CAUTION: If the coefficient matrix  $A$  is symmetric, this setting is likely to cause the filled matrix,  $F = L + U - I$ , to be non-symmetric. Guidance: if the matrix is poorly scaled, turning on row scaling may help convergence.
- **-ilut**  $\langle float \rangle$  Use ILUT factorization instead of the default,  $ILU(k)$ . Here,  $\langle float \rangle$  is the drop tolerance, which is relative to the largest absolute value of any entry in the row being factored. CAUTION: If the coefficient matrix  $A$  is symmetric, this setting is likely to cause the filled matrix,  $F = L + U - I$ , to be non-symmetric. NOTE: this option can only be used sequentially!

## 5.15 PILUT: Parallel Incomplete Factorization

### Warning

PILUT is not actively supported by the hypr development team. We recommend using *ILU* for parallel ILU algorithms. This new ILU implementation includes 64-bit integers support (for linear systems with more than 2,147,483,647 global unknowns) through both *mixedint* and *bigint* builds of hypr and NVIDIA/AMD GPUs support through the CUDA/HIP backends.

PILUT is a parallel preconditioner based on Saad's dual-threshold incomplete factorization algorithm. The original version of PILUT was done by Karypis and Kumar [KaKu1998] in terms of the Cray SHMEM library. The code was subsequently modified by the hypr team: SHMEM was replaced by MPI; some algorithmic changes were made; and it was software engineered to be interoperable with several matrix implementations, including hypr's ParCSR format, PETSc's matrices, and ISIS++ RowMatrix. The algorithm produces an approximate factorization  $LU$ , with the preconditioner  $M$  defined by  $M = LU$ .

**Note**

PILUT produces a nonsymmetric preconditioner even when the original matrix is symmetric. Thus, it is generally inappropriate for preconditioning symmetric methods such as Conjugate Gradient.

**5.15.1 Parameters:**

- `SetMaxNonzerosPerRow( int LFIL );` (Default: 20) Set the maximum number of nonzeros to be retained in each row of  $L$  and  $U$ . This parameter can be used to control the amount of memory that  $L$  and  $U$  occupy. Generally, the larger the value of `LFIL`, the longer it takes to calculate the preconditioner and to apply the preconditioner and the larger the storage requirements, but this trades off versus a higher quality preconditioner that reduces the number of iterations.
- `SetDropTolerance( double tol );` (Default: 0.0001) Set the tolerance (relative to the 2-norm of the row) below which entries in  $L$  and  $U$  are automatically dropped. PILUT first drops entries based on the drop tolerance, and then retains the largest `LFIL` elements in each row that remain. Smaller values of `tol` lead to more accurate preconditioners, but can also lead to increases in the time to calculate the preconditioner.

**5.16 LOBPCG Eigensolver**

LOBPCG (Locally Optimal Block Preconditioned Conjugate Gradient) is a simple, yet very efficient, algorithm suggested in [Knya2001], [KLAO2007], [BLOPEWeb] for computing several smallest eigenpairs of the symmetric generalized eigenvalue problem  $Ax = \lambda Bx$  with large, possibly sparse, symmetric matrix  $A$  and symmetric positive definite matrix  $B$ . The matrix  $A$  is not assumed to be positive, which also allows one to use LOBPCG to compute the largest eigenpairs of  $Ax = \lambda Bx$  simply by solving  $-Ax = \mu Bx$  for the smallest eigenvalues  $\mu = -\lambda$ .

LOBPCG simultaneously computes several eigenpairs together, which is controlled by the `blockSize` parameter, see example `ex11.c`. The LOBPCG also allows one to impose constraints on the eigenvectors of the form  $x^T B y_i = 0$  for a set of vectors  $y_i$  given to LOBPCG as input parameters. This makes it possible to compute, e.g., 50 eigenpairs by 5 subsequent calls to LOBPCG with the `blockSize=10`, using deflation. LOBPCG can use preconditioning in two different ways: by running an inner preconditioned PCG linear solver, or by applying the preconditioner directly to the eigenvector residual (option `-pcgitr 0`). In all other respects, LOBPCG is similar to the PCG linear solver.

The LOBPCG code is available for system interfaces: `Struct`, `SStruct`, and `IJ`. It is also used in the Auxiliary-space Maxwell Eigensolver (AME). The LOBPCG setup is similar to the setup for PCG.



## MIXED PRECISION

The hypre library has provided compile-time multi-precision support for many years. For example, the autotools option `--enable-single` or the CMake option `-DHYPRE_SINGLE=ON` will produce a single precision library.

Starting in hypre version 3.0, multiple precision and mixed precision support are provided at runtime. To turn this on, use

- `--enable-mixed-precision` (autotools)
- `-DHYPRE_ENABLE_MIXED_PRECISION=ON` (CMake)

With the above, users can compile, link, and run as before without changes to their code. To access runtime precision, there are several levels of support that can be used, outlined in the following sections. For clarity, a generic name `Foo` is used to represent a function in hypre, e.g., `HYPRE_PCGSolve`.

### 6.1 Calling functions with fixed precision

For every function `Foo` in hypre, the following fixed-precision versions are also available,

- `Foo_float`
- `Foo_double`
- `Foo_long_double`

where the precision of each function is determined by the C compiler and the respective C types `float`, `double`, and `long double`. The prototypes for these functions are exactly the same as for `Foo`, but with real-valued arguments like `HYPRE_Real` mapping to the specific C types (and precisions) indicated above.

### 6.2 Calling functions with multiple precisions

Every user-API function `Foo` in hypre (any function beginning with the upper case `HYPRE_` prefix) is also available in the mixed-precision configuration of the library, but its precision is determined by a global runtime precision that can be set by calling

```
HYPRE_Int HYPRE_SetGlobalPrecision(HYPRE_Precision precision)
```

where `precision` is either `HYPRE_REAL_SINGLE`, `HYPRE_REAL_DOUBLE`, or `HYPRE_REAL_LONGDOUBLE`. Real-valued arguments for `Foo` have different types from the functions described in Section *Calling functions with fixed precision* because they have to support all three precisions, but calling `Foo` in practice is much the same. Specifically, real arrays such as `HYPRE_Real *` become `void *`, and real values such as `HYPRE_Real` become `long double`. This prototyping enables multiple-precision functionality, although strong type checking at compile time is lost.



## GENERAL INFORMATION

In this and the following sections, we discuss how to obtain and build hypre, interpret error flags, report bugs, and call hypre from different programming languages. We provide instructions for two build systems: autotools (configure & make) and CMake. While autotools is traditionally used on Unix-like systems (Linux, macOS, etc.), CMake provides cross-platform support and is required for Windows builds. Both systems are actively maintained and supported.

### 7.1 Getting the Source Code

There are two ways to obtain hypre:

1. **Clone from GitHub (recommended):**

```
git clone https://github.com/hypre-space/hypre.git
cd hypre/src
```

2. **Download a Release**

Download the latest release from our [GitHub releases page](#). After extracting the archive, you'll find the source code in the src directory:

```
tar -xvf hypre-x.y.z.tar.gz
cd hypre-x.y.z/src
```

where `x.y.z` represents the version number (e.g., 2.29.0).

### 7.2 Building the Library

After obtaining the source code (see *Getting the Source Code*), there are three main ways to build hypre:

#### Tip

For the fastest build, use CMake with the Ninja generator (if Ninja is installed):

```
cmake -G Ninja
```

Ninja handles parallel builds efficiently, often completing faster than a traditional Makefile build, especially on multi-core systems.

## 7.2.1 1. Using autotools (Configure & Make)

The simplest method is to use the traditional configure and make:

```
cd ${HYPRE_HOME}/src    # Move to the source directory
./configure             # Configure the build system
make -j 4               # Use threads for a faster parallel build
make install            # (Optional) Install hypr on a user-specified path via --prefix=
↳ <path>
```

This will build and install hypr in the default locations:

- Libraries: `${HYPRE_HOME}/src/hypr/lib`
- Headers: `${HYPRE_HOME}/src/hypr/include`

There are many options to configure and make to customize such things as installation directories, compilers used, compile and load flags, etc. For more information on the configure options, see *Build System Options*.

Executing `configure` results in the creation of platform specific files that are used when building the library. The information may include such things as the system type being used for building and executing, compilers being used, libraries being searched, option flags being set, etc. When all of the searching is done two files are left in the `src` directory; `config.status` contains information to recreate the current configuration and `config.log` contains compiler messages which may help in debugging configure errors. Upon successful completion of configure the file `config/Makefile.config` is created from its template `config/Makefile.config.in` and hypr is ready to be built.

Executing `make`, with or without targets being specified, in the `src` directory initiates compiling of all of the source code and building of the hypr library. If any errors occur while compiling, the user can edit the file `config/Makefile.config` directly then run `make` again; without having to re-run `configure`. When building hypr without the `install` target, the libraries and include files will be copied into the default directories, `src/hypr/lib` and `src/hypr/include`, respectively. When building hypr using the `install` target, the libraries and include files will be copied into the directories that the user specified in the options to `configure`, e.g. `--prefix=/usr/apps`. If none were specified the default directories, `src/hypr/lib` and `src/hypr/include`, are used.

## 7.2.2 2. Using CMake (Windows, macOS, Linux, etc.)

CMake provides a modern, platform-independent build system. When using CMake to build hypr, several files and directories are created during the build process:

- `CMakeCache.txt` - Stores configuration options and settings
- `CMakeFiles/` - Contains intermediate build files and dependency information
- `cmake_install.cmake` - Instructions for installing the built library
- `Makefile` - Generated build instructions (on Unix-like systems)

The build process has three main steps:

1. **Configure:** CMake reads the `CMakeLists.txt` files and generates the build system
2. **Build:** The native build tool (`make`, Visual Studio, etc.) compiles the code
3. **Install:** Built libraries and headers are copied to their final location

The simplest way to build hypr using CMake is:

```
cd ${HYPRE_HOME}/build    # Use a separate build directory to keep source clean
cmake ../src              # Generate build system
make -j                   # Build the library in parallel
```

(continues on next page)

(continued from previous page)

```
make install # (Optional) Install to specified location via -DCMAKE_
↳INSTALL_PREFIX=<path>
```

During the configure step, CMake will detect your compiler and build tools, it will find required dependencies, set up platform-specific build flags, and generate native build files. If any errors occur during configuration, check CMakeCache.txt for current settings and CMakeFiles/CMakeError.log for detailed error messages. The build step will create:

- Static library: libHYPRE.a (Unix/macOS) or HYPRE.lib (Windows)
- Shared library: libHYPRE.so (Linux), libHYPRE.dylib (macOS), or HYPRE.dll (Windows) if enabled
- Object files in CMakeFiles/ subdirectories

By default, make will place the library file in `${HYPRE_HOME}/src/hypre/lib` and the header files in `${HYPRE_HOME}/src/hypre/include`. As with the autotools method, hypre's CMake build provides several options. For more information, see *Build System Options*.

### Note

CMake GUI (ccmake or cmake-gui) provides an interactive way to change build options:

- **Unix:** From the `${HYPRE_HOME}/build` directory:
  1. Run `ccmake ../src`
  2. Change options: - Press Enter to modify a variable - Boolean options (ON/OFF) toggle with Enter - String/file options allow text editing
  3. Press 'c' to configure
  4. Repeat until satisfied
  5. Press 'g' to generate
- **Windows:** Using Visual Studio:
  1. Change desired options
  2. Click "Configure"
  3. Click "Generate"

## 7.2.3 3. Using Spack (Recommended for HPC environments)

Spack is a package manager designed for supercomputers, Linux, and macOS. It makes installing scientific software easy and handles dependencies automatically. To build hypre using Spack:

```
# Install Spack if you haven't already
git clone -c feature.manyFiles=true --depth=2 https://github.com/spack/spack.git
. spack/share/spack/setup-env.sh

# Install hypre with default options
spack install hypre

# Or install with specific options (e.g., with CUDA support)
spack install hypre+cuda
```

Common Spack variants for hypre include:

- `+mpi` / `~mpi` - Enable/disable MPI support (default: `+mpi`)
- `+cuda` / `~cuda` - Enable/disable CUDA support (default: `~cuda`)
- `+openmp` / `~openmp` - Enable/disable OpenMP support (default: `~openmp`)
- `+shared` / `~shared` - Build shared libraries (default: `~shared`)
- `+debug` / `~debug` - Build with debug flags (default: `~debug`)

To see all available build options:

```
spack info hypr
```

### Note

Spack will automatically handle dependencies and choose appropriate versions based on your system and requirements. It's particularly useful in HPC environments where you need to manage multiple versions or build configurations of hypr and its dependencies.

## 7.3 Build System Options

The table below lists the most commonly used build options for both autotools and CMake build systems. Each option is shown with its default value (if applicable) and any relevant platform restrictions. For GPU-specific options, see the *GPU Build Options* section below.

Table 7.1: Build Configuration Options

Feature	Autotools (configure)	CMake
Install Path	<code>--prefix=&lt;path&gt;</code>	<code>-DCMAKE_INSTALL_PREFIX=&lt;path&gt;</code>
Debug Build (default is off)	<code>--enable-debug</code>	<code>-DCMAKE_BUILD_TYPE=Debug</code>
Memory tracker (default is off)	<code>--with-memory-tracker</code>	<code>-DHYPRE_ENABLE_MEMORY_TRACKER=ON</code>
Print Errors (default is off)	<code>--with-print-errors</code>	<code>-DHYPRE_ENABLE_PRINT_ERRORS=ON</code>
Floating-point exception trap (default is off)	<code>--enable-fpe-trap</code>	<code>-DHYPRE_ENABLE_FPE_TRAP=ON</code>
Shared Library (default is off)	<code>--enable-shared</code>	<code>-DBUILD_SHARED_LIBS=ON</code>
64-bit integers (default is off, no GPU support)	<code>--enable-bigint</code>	<code>-DHYPRE_ENABLE_BIGINT=ON</code>
Mixed 32/64-bit integers (default is off)	<code>--enable-mixedint</code>	<code>-DHYPRE_ENABLE_MIXEDINT=ON</code>
Single FP precision (default is off)	<code>--enable-single</code>	<code>-DHYPRE_ENABLE_SINGLE=ON</code>
Long double precision (default is off, no GPU support)	<code>--enable-long-double</code>	<code>-DHYPRE_ENABLE_LONG_DOUBLE=ON</code>
Link-time optimization (default is off)	N/A	<code>-DHYPRE_ENABLE_LTO=ON</code>
MPI Support (default is on)	<code>--enable-mpi</code>	<code>-DHYPRE_ENABLE_MPI=ON</code>
MPI Persistent (default is off)	<code>--enable-persistent</code>	<code>-DHYPRE_ENABLE_PERSISTENT_COMM=ON</code>
OpenMP Support (default is off)	<code>--with-openmp</code>	<code>-DHYPRE_ENABLE_OPENMP=ON</code>
Hopscotch hashing (requires OpenMP) (default is off)	<code>--enable-hopscotch</code>	<code>-DHYPRE_ENABLE_HOPSCOTCH=ON</code>
Fortran Support (default is on)	<code>--enable-fortran</code>	<code>-DHYPRE_ENABLE_FORTRAN=ON</code>
Fortran mangling	<code>--with-fmangle</code>	<code>-DHYPRE_ENABLE_FMANGLE=0</code>

**Note**

- CMake options are case-sensitive
- Boolean CMake options accept ON/OFF values
- Executables located under `src/test` and `src/examples` are built separately when using the autotools build system
- For a complete list of options:
  - **Autotools:** Run `./configure --help`
  - **CMake:** See `CMakeLists.txt` or run `cmake -LAH`
- For third-party libraries (TPLs), hypr supports two methods:
  1. **CMake Package Config (recommended):** Use `-DPackage_ROOT=/path/to/package` to help CMake find package configuration files
  2. **Manual specification:**
    - a. **Autotools:**

```
--with-pkg-include=/path/to/pkg-include
--with-pkg-lib=/path/to/pkg-lib
```
    - b. **CMake:**

```
-DTPL_PACKAGE_INCLUDE_DIRS=/path/to/pkg-include
-DTPL_PACKAGE_LIBRARIES=/path/to/pkg-lib/libpackage.so
```

## 7.4 GPU Build Options

The hypr library provides support for multiple GPU architectures through different programming models: CUDA (for NVIDIA GPUs), HIP (for AMD GPUs), and SYCL (for Intel GPUs). Each model has its own set of build options and requirements. Some solvers and features may have different levels of support across these platforms. Key considerations when building for GPUs are:

1. Only one GPU backend can be enabled at a time (CUDA, HIP, or SYCL)
2. Some features like full support for 64-bit integers (*BigInt*) are not available
3. Memory management options (device vs unified memory) affect solver availability
4. Umpire is implicitly enabled by default when building with CUDA or HIP support

The table below lists the available GPU-specific build options for both autotools and CMake build systems.

Table 7.2: GPU Configuration Options

Feature	Autotools (configure)	CMake
CUDA Support (default is off)	<code>--with-cuda</code>	<code>-DHYPRE_ENABLE_CUDA=ON</code>
HIP Support (default is off)	<code>--with-hip</code>	<code>-DHYPRE_ENABLE_HIP=ON</code>
SYCL Support (default is off)	<code>--with-sycl</code>	<code>-DHYPRE_ENABLE_SYCL=ON</code>
SYCL Target (default is empty, <b>SYCL</b> only)	<code>--with-sycl-target=ARG</code>	<code>-DHYPRE_SYCL_TARGET=ARG</code>
SYCL Target Backend (default is empty, <b>SYCL</b> only)	<code>--with-sycl-target-backend=ARG</code>	<code>-DHYPRE_SYCL_TARGET_BACKEND=ARG</code>
GPU architecture (determined automatically)	<code>--with-gpu-arch=ARG</code>	<code>-DCMAKE_CUDA_ARCHITECTURES=ARG</code> <code>-DCMAKE_HIP_ARCHITECTURES=ARG</code>
GPU Profiling (default is off)	<code>--enable-gpu-profiling</code>	<code>-DHYPRE_ENABLE_GPU_PROFILING=ON</code>
GPU-aware MPI (default is off)	<code>--enable-gpu-aware-mpi</code>	<code>-DHYPRE_ENABLE_GPU_AWARE_MPI=ON</code>
Unified Memory (default is off)	<code>--enable-unified-memory</code>	<code>-DHYPRE_ENABLE_UNIFIED_MEMORY=ON</code>
Device async malloc (default is off)	<code>--enable-device-malloc-async</code>	<code>-DHYPRE_ENABLE_DEVICE_MALLOC_ASYNC=ON</code>
Thrust async execution (default is off)	<code>--enable-thrust-async</code>	<code>-DHYPRE_ENABLE_THRUST_ASYNC=ON</code>
cuSPARSE Support (default is on, <b>CUDA</b> only)	<code>--enable-cusparse</code>	<code>-DHYPRE_ENABLE_CUSPARSE=ON</code>
cuSOLVER Support (default is on, <b>CUDA</b> only)	<code>--enable-cusolver</code>	<code>-DHYPRE_ENABLE_CUSOLVER=ON</code>
cuBLAS Support (default is on, <b>CUDA</b> only)	<code>--enable-cublas</code>	<code>-DHYPRE_ENABLE_CUBLAS=ON</code>
cuRAND Support (default is on, <b>CUDA</b> only)	<code>--enable-curand</code>	<code>-DHYPRE_ENABLE_CURAND=ON</code>

#### 7.4. GPU Build Options

**Note**

Allocations and deallocations of GPU memory can be slow. Memory pooling is a common approach to reduce such overhead and improve performance. For better performance, [Umpire] is enabled by default for CUDA and HIP builds and provides robust pooling capabilities for both device and unified memory.

For SYCL builds, Umpire remains optional and must be enabled explicitly.

See *Building Umpire* for detailed instructions on how to build and use Umpire with HYPRE.

**Note**

When hypr is configured with device support, but without unified memory, the memory allocated on the GPUs, by default, is the GPU device memory, which is not accessible from the CPUs. Hypr's structured solvers can run with device memory, whereas only selected unstructured solvers can run with device memory. See *GPU-supported Options* for details. Some solver options for BoomerAMG require unified (managed) memory.

## 7.4.1 Building Umpire

HYPRE provides three primary methods for integrating Umpire for GPU memory management: the recommended Automatic Build, where HYPRE handles the download and configuration process automatically; a Manual Build from Source, which offers maximum control for advanced users; and installation via Package Managers like Spack, which is convenient for managing Umpire within an existing software environment.

### Automatic Umpire Build

The easiest way to use Umpire with HYPRE is to enable the automatic build feature. This is recommended if you don't have Umpire pre-installed or want to ensure maximum compatibility. When enabled, HYPRE handles the entire process: it automatically downloads a compatible Umpire version, detects your GPU backend (CUDA, HIP, or SYCL), and configures Umpire with settings optimized for HYPRE's specific memory management needs. This approach simplifies setup by eliminating the need to manually manage Umpire's installation, dependencies, or build configuration.

To enable this feature, set `HYPRE_BUILD_UMPIRE=ON` during CMake configuration:

```
# Enable automatic Umpire build and CUDA
cmake -DHYPRE_BUILD_UMPIRE=ON \
      -DHYPRE_ENABLE_CUDA=ON \
      ../src

# Build HYPRE (Umpire will be built automatically)
make -j
```

### Manual Umpire Build

If you prefer to build Umpire separately or need specific Umpire configurations, you can build it manually from source:

```
git clone --recursive https://github.com/LLNL/Umpire.git

cd Umpire
cmake -S . -B build \
      -DUMPIRE_ENABLE_C=ON \
      -DUMPIRE_ENABLE_TOOLS=OFF \
      -DENABLE_CUDA=${ENABLE_CUDA} \
```

(continues on next page)

(continued from previous page)

```

-DENABLE_HIP=${ENABLE_HIP} \
-DENABLE_SYCL=${ENABLE_SYCL} \
-DENABLE_BENCHMARKS=OFF \
-DENABLE_EXAMPLES=OFF \
-DENABLE_DOCS=OFF \
-DENABLE_TESTS=OFF \
-DCMAKE_BUILD_TYPE=Release \
-DCMAKE_INSTALL_LIBDIR=/path-to-umpire-install/lib \
-DCMAKE_INSTALL_PREFIX=/path-to-umpire-install

cmake --build build -j
cmake --install build

```

Enable either CUDA, HIP, or SYCL by setting the corresponding flag to ON and the others to OFF.

### Using Spack Package Manager

```

# Install Umpire with Spack
spack install umpire+cuda # or +hip, +sycl

```

### Linking with Pre-built Umpire

After building or installing Umpire manually, configure HYPRE to use it:

#### Autotools:

```

./configure --with-umpire-include=/path-to-umpire-install/include \
            --with-umpire-lib-dirs=/path-to-umpire-install/lib \
            --with-umpire-libs="umpire camp" \

```

#### CMake:

```

# Method 1: Using Umpire's CMake config
cmake -DHYPRE_ENABLE_UMPIRE=ON \
      -Dumpire_DIR=/path-to-umpire-install/lib/cmake/umpire \
      ../src

# Method 2: Using umpire_ROOT pattern
cmake -DHYPRE_ENABLE_UMPIRE=ON \
      -Dumpire_ROOT=/path-to-umpire-install \
      ../src

# Method 3: Manual specification
cmake -DHYPRE_ENABLE_UMPIRE=ON \
      -DTPL_UMPIRE_INCLUDE_DIRS=/path-to-umpire-install/include \
      -DTPL_UMPIRE_LIBRARIES=/path-to-umpire-install/lib/libumpire.so \
      ../src

```

## 7.5 Make Targets

The make step in building hypr is where the compiling, loading and creation of libraries occurs. Make has several options that are called targets. These include:

help	prints the details of each target
all	default target in all directories compile the entire library does NOT rebuild documentation
clean	deletes all files from the current directory that are created by building the library
distclean	deletes all files from the current directory that are created by configuring or building the library
install	compile the source code, build the library and copy executables, libraries, etc to the appropriate directories for user access
uninstall	deletes all files that the install target created
tags	runs etags to create a tags table file is named TAGS and is saved in the top-level directory
test	depends on the all target to be completed removes existing temporary installation directories creates temporary installation directories copies all libHYPRE* and *.h files to the temporary locations builds the test drivers; linking to the temporary locations to simulate how application codes will link to HYPRE

## 7.6 Using the Library

The `examples` subdirectory contains several codes that demonstrate hypr's features and can be used to test the library. These examples can be built in two ways:

1. **Using CMake:** Enable the `HYPRE_BUILD_EXAMPLES` option during configuration:

```
cmake -DHYPRE_BUILD_EXAMPLES=ON ..  
make
```

2. **Using Makefiles:** Navigate to the `examples` subdirectory and build directly:

```
cd examples  
make
```

Each example contains detailed comments at the beginning of its source file explaining its purpose and how to run it. The examples demonstrate various interfaces, solvers, and problem types. For a categorized list of examples and their features, see the HTML documentation in the `examples/docs` directory.

**Note**

The examples are designed to mimic real application codes and can serve as templates for your own implementations.

## 7.7 Testing the Library

hyre provides several approaches to test the library, in increasing order of comprehensiveness:

1. **Basic Tests** (Recommended first step): Quick tests to check library functionality (CMake requires `-DBUILD_TESTING=ON`):

```
# Single test for each linear system interface
make check

# Test IJ, Struct and SStruct linear solvers in parallel
make checkpar
```

2. **Comprehensive Tests** (CMake only): Test linear solvers for all linear system interfaces (linear-algebraic, Struct and SStruct):

```
cmake -DBUILD_TESTING=ON ..
make -j
make test # or ctest
```

3. **Automated Testing** (Developers only): For thorough testing across different configurations and machines including regression tests, and performance benchmarks, with support for both CPU and GPU executions. Test results are automatically compared against saved baseline outputs, with the ability to update these baselines when legitimate changes occur. The automated testing infrastructure is particularly focused on ensuring consistency across different build configurations and execution environments. For more information, see the [README](#) file.

**Note**

- Test tolerance can be adjusted using `-DHYPRE_CHECK_TOL=<value>` during CMake configuration. Default tolerance is 1.0e-6
- Test output files with `.err` extension contain error messages and diagnostics
- AUTOTEST configurations can be customized by modifying machine-specific files in the AUTOTEST directory

For detailed test results and logs:

- Make check results: `build/test/*.err` (CMake) or `src/test/TEST_(ij|struct|ssstruct)/*.err` (Autotools)
- CTest results: `build/Testing/Temporary/LastTest.log`
- AUTOTEST results: `src/AUTOTEST/machine_name.dir/machine_name.err`

## 7.8 Linking to the Library

There are two main approaches to link your application with hypr:

### 7.8.1 Using CMake

The hypr library provides CMake configuration files that enable easy integration. Create a `CMakeLists.txt` with:

```
cmake_minimum_required(VERSION 3.21)
project(MyApp LANGUAGES C)

find_package(HYPRE REQUIRED)

add_executable(myapp main.c)
target_link_libraries(myapp PUBLIC HYPRE::HYPRE lm)
```

If hypr is not in a standard location, specify its path:

```
cmake -DHYPRE_ROOT=/path/to/hypr-install-directory ..
```

### 7.8.2 Using Autotools

For non-CMake builds, manually specify compilation and linking flags:

```
# Compilation
-I${HYPRE_INSTALL_DIR}/include

# Linking
-L${HYPRE_INSTALL_DIR}/lib -lHYPRE -lm
```

Where `${HYPRE_INSTALL_DIR}` is your hypr installation directory (default is `${HYPRE_HOME}/src/hypr`, or as specified by `--prefix=PREFIX` during configuration).

### 7.8.3 Shared Library Considerations

If hypr was built as a shared library, you have several options:

1. **Environment Variables:** Add hypr's library location to your system's library path:

```
# Linux/Unix
export LD_LIBRARY_PATH=${HYPRE_INSTALL_DIR}/lib:${LD_LIBRARY_PATH}

# macOS
export DYLD_LIBRARY_PATH=${HYPRE_INSTALL_DIR}/lib:${DYLD_LIBRARY_PATH}

# Windows
set PATH=%HYPRE_INSTALL_DIR%\lib;%PATH%
```

2. **RPATH/RUNPATH:** Set the runtime search path during linking. With CMake:

```
# Use RPATH (searched before LD_LIBRARY_PATH)
set(CMAKE_INSTALL_RPATH "${HYPRE_INSTALL_DIR}/lib")
set(CMAKE_INSTALL_RPATH_USE_LINK_PATH TRUE)
```

(continues on next page)

(continued from previous page)

```
# Or use RUNPATH (searched after LD_LIBRARY_PATH)
set(CMAKE_SHARED_LINKER_FLAGS "-Wl,--enable-new-dtags")
set(CMAKE_INSTALL_RPATH "${HYPRE_INSTALL_DIR}/lib")
set(CMAKE_INSTALL_RPATH_USE_LINK_PATH TRUE)
```

Or with manual linking:

```
# RPATH
-Wl,-rpath,${HYPRE_INSTALL_DIR}/lib

# RUNPATH
-Wl,--enable-new-dtags,-rpath,${HYPRE_INSTALL_DIR}/lib
```

RPATH is searched before LD\_LIBRARY\_PATH while RUNPATH is searched after, giving you flexibility in controlling library resolution precedence.

### **Note**

For examples of linking applications with hypr, refer to the `examples` subdirectory.

## 7.9 Error Flags

Every hypr function returns an integer, which is used to indicate errors during execution. Note that the error flag returned by a given function reflects the errors from *all* previous calls to hypr functions. In particular, a value of zero means that all hypr functions up to (and including) the current one have completed successfully. This new error flag system is being implemented throughout the library, but currently there are still functions that do not support it. The error flag value is a combination of one or a few of the following error codes:

1. HYPRE\_ERROR\_GENERIC – describes a generic error
2. HYPRE\_ERROR\_MEMORY – hypr was unable to allocate memory
3. HYPRE\_ERROR\_ARG – error in one of the arguments of a hypr function
4. HYPRE\_ERROR\_CONV – a hypr solver did not converge as expected

One can use the HYPRE\_CheckError function to determine exactly which errors have occurred:

```
/* call some HYPRE functions */
int hypr_ierr;
hypr_ierr = HYPRE_Function();

/* check if the previously called hypr functions returned error(s) */
if (hypr_ierr)
  /* check if the error with code HYPRE_ERROR_CODE has occurred */
  if (HYPRE_CheckError(hypr_ierr,HYPRE_ERROR_CODE))
```

The corresponding FORTRAN code is

```
! header file with hypr error codes
include 'HYPRE_error_f.h'

! call some HYPRE functions
```

(continues on next page)

(continued from previous page)

```

integer hypr_ierr
call HYPRE_Function(hypr_ierr)

! check if the previously called hypr functions returned error(s)
if (hypr_ierr .ne. 0) then
  ! check if the error with code HYPRE_ERROR_CODE has occurred
  call HYPRE_CheckError(hypr_ierr, HYPRE_ERROR_CODE, check)
  if (check .ne. 0) then

```

The global error flag can also be obtained directly, between calls to other hypr functions, by calling `HYPRE_GetError()`. If an argument error (`HYPRE_ERROR_ARG`) has occurred, the argument index (counting from 1) can be obtained from `HYPRE_GetErrorArg()`. To get a character string with a description of all errors in a given error flag, use

```
HYPRE_DescribeError(int hypr_ierr, char *descr);
```

The global error flag can be cleared manually by calling `HYPRE_ClearAllErrors()`, which will essentially ignore all previous hypr errors. To only clear a specific error code, the user can call `HYPRE_ClearError(HYPRE_ERROR_CODE)`. Finally, if hypr was configured with `--with-print-errors` or `-DHYPRE_ENABLE_PRINT_ERRORS=ON`, additional error information will be printed to the standard error during execution.

## 7.10 Bug Reporting and General Support

For bug reports, feature requests, and general usage questions, please create an issue on [GitHub issues](#). You can also browse existing issues to see if your question has already been addressed. To help us address your issue effectively, please include:

### Required Information:

- hypr version number
- Description of the problem or feature request
- Minimal example demonstrating the issue (if applicable)

### For Build Issues:

- Build system used (CMake or autotools)
- Build configuration options
- Complete build output showing the error
- Operating system and version
- Compiler and version
- MPI implementation and version

### For Runtime Issues:

- Command line arguments used
- Problem size and configuration
- Number of processes/threads
- Complete error messages or stack traces
- Information about the computing environment:

- GPU type and driver version (for GPU builds)
- Relevant environment variables
- System architecture (CPU type, memory)

#### For Performance Issues:

- Performance measurements or profiling data
- Comparison with previous versions (if applicable)
- Problem size and scaling information
- Hardware configuration details

## 7.11 Calling HYPRE from Other Languages

The hypr library currently supports two languages: C (native) and Fortran (in version 2.10.1 and earlier, additional language interfaces were also provided through a tool called Babel). The Fortran interface is manually supported to mirror the “native” C interface used throughout most of this manual. We describe this interface next.

Typically, the Fortran subroutine name is the same as the C name, unless it is longer than 31 characters. In these situations, the name is condensed to 31 characters, usually by simple truncation. For now, users should look at the Fortran test drivers (\*.f codes) in the `test` directory for the correct condensed names. In the future, this aspect of the interface conversion will be made consistent and straightforward.

The Fortran subroutine argument list is always the same as the corresponding C routine, except that the error return code `ierr` is always last. Conversion from C parameter types to Fortran argument type is summarized in following table:

C parameter	Fortran argument
<code>int i</code>	<code>integer i</code>
<code>double d</code>	<code>double precision d</code>
<code>int *array</code>	<code>integer array(*)</code>
<code>double *array</code>	<code>double precision array(*)</code>
<code>char *string</code>	<code>character string(*)</code>
<code>HYPRE_Type object</code>	<code>integer*8 object</code>
<code>HYPRE_Type *object</code>	<code>integer*8 object</code>

Array arguments in hypr are always of type `(int *)` or `(double *)`, and the corresponding Fortran types are simply `integer` or `double precision` arrays. Note that the Fortran arrays may be indexed in any manner. For example, an integer array of length `N` may be declared in fortran as either of the following:

```
integer array(N)
integer array(0:N-1)
```

hypr objects can usually be declared as in the table because `integer*8` usually corresponds to the length of a pointer. However, there may be some machines where this is not the case. On such machines, the Fortran type for a hypr object should be an `integer` of the appropriate length.

This simple example illustrates the above information:

C prototype:

```
int HYPRE_IJMatrixSetValues(HYPRE_IJMatrix matrix,
                           int nrows, int *ncols,
```

(continues on next page)

(continued from previous page)

```
const int *rows, const int *cols,  
const double *values);
```

The corresponding Fortran code for calling this routine is as follows:

```
integer*8      matrix  
integer       nrows, ncols(MAX_NCOLS)  
integer       rows(MAX_ROWS), cols(MAX_COLS)  
double precision values(MAX_COLS)  
integer       ierr  
  
call HYPRE_IJMatrixSetValues(matrix, nrows, ncols, rows, cols, values, ierr)
```

## 8.1 Struct System and Object Interface

### *group* Struct System and Object Interface

A structured-grid conceptual interface.

This interface represents a structured-grid conceptual view of a linear system.

#### Struct Grids

typedef struct hypre\_StructGrid\_struct **HYPRE\_StructGrid**

A grid object is constructed out of several “boxes”, defined on a global abstract index space.

**HYPRE\_Int HYPRE\_StructGridCreate**(MPI\_Comm comm, HYPRE\_Int ndim, *HYPRE\_StructGrid* \*grid)  
Create an *ndim*-dimensional grid object.

**HYPRE\_Int HYPRE\_StructGridDestroy**(*HYPRE\_StructGrid* grid)  
Destroy a grid object.

An object should be explicitly destroyed using this destructor when the user’s code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

**HYPRE\_Int HYPRE\_StructGridSetExtents**(*HYPRE\_StructGrid* grid, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper)

Set the extents for a box on the grid.

**HYPRE\_Int HYPRE\_StructGridAssemble**(*HYPRE\_StructGrid* grid)  
Finalize the construction of the grid before using.

**HYPRE\_Int HYPRE\_StructGridPrintVTK**(const char \*filename, *HYPRE\_StructGrid* grid)  
Prints a grid in VTK format.

**HYPRE\_Int HYPRE\_StructGridSetPeriodic**(*HYPRE\_StructGrid* grid, HYPRE\_Int \*periodic)  
Set the periodicity for the grid.

The argument *periodic* is an *ndim*-dimensional integer array that contains the periodicity for each dimension. A zero value for a dimension means non-periodic, while a nonzero value means periodic and contains the actual period. For example, periodicity in the first and third dimensions for a 10x11x12 grid is indicated by the array [10,0,12].

NOTE: Some of the solvers in hypre have power-of-two restrictions on the size of the periodic dimensions.

HYPRE\_Int **HYPRE\_StructGridSetNumGhost**(*HYPRE\_StructGrid* grid, HYPRE\_Int \*num\_ghost)

Set the ghost layer in the grid object.

HYPRE\_Int **HYPRE\_StructGridCoarsen**(*HYPRE\_StructGrid* grid, HYPRE\_Int \*stride, *HYPRE\_StructGrid* \*cgrid)

Coarsen *grid* by factor *stride* to create *cgrid*.

HYPRE\_Int **HYPRE\_StructGridProjectBox**(*HYPRE\_StructGrid* grid, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int \*origin, HYPRE\_Int \*stride)

Project the box described by *ilower* and *iupper* onto the strided index space that contains the index *origin* and has stride *stride*.

This routine is useful in combination with *HYPRE\_StructGridCoarsen* when dealing with rectangular matrices.

## Struct Stencils

typedef struct hypr\_StructStencil\_struct \***HYPRE\_StructStencil**

The stencil object.

HYPRE\_Int **HYPRE\_StructStencilCreate**(HYPRE\_Int ndim, HYPRE\_Int size, *HYPRE\_StructStencil* \*stencil)

Create a stencil object for the specified number of spatial dimensions and stencil entries.

HYPRE\_Int **HYPRE\_StructStencilDestroy**(*HYPRE\_StructStencil* stencil)

Destroy a stencil object.

HYPRE\_Int **HYPRE\_StructStencilSetEntry**(*HYPRE\_StructStencil* stencil, HYPRE\_Int entry, HYPRE\_Int \*offset)

Set a stencil entry.

HYPRE\_Int **HYPRE\_StructStencilSetElement**(*HYPRE\_StructStencil* stencil, HYPRE\_Int entry, HYPRE\_Int \*offset)

## Struct Matrices

typedef struct hypr\_StructMatrix\_struct \***HYPRE\_StructMatrix**

The matrix object.

HYPRE\_Int **HYPRE\_StructMatrixCreate**(MPI\_Comm comm, *HYPRE\_StructGrid* grid, *HYPRE\_StructStencil* stencil, *HYPRE\_StructMatrix* \*matrix)

Create a matrix object.

Matrices may have different range and domain grids, that is, they need not be square. By default, the range and domain grids are the same as *grid*. In general, the range is a coarsening of *grid* as specified in *HYPRE\_StructMatrixSetRangeStride*, and similarly for the domain. Note that the range index space must either be a subspace of the domain index space or vice versa. Also, (currently) either the range or domain coarsening factor (or both) must be all ones (i.e., no coarsening).

HYPRE\_Int **HYPRE\_StructMatrixDestroy**(*HYPRE\_StructMatrix* matrix)

Destroy a matrix object.

HYPRE\_Int **HYPRE\_StructMatrixSetRangeStride**(*HYPRE\_StructMatrix* matrix, HYPRE\_Int \*range\_stride)

(Optional) Set the range coarsening stride.

For more information, see *HYPRE\_StructMatrixCreate*.

`HYPRE_Int HYPRE_StructMatrixSetDomainStride`(*HYPRE\_StructMatrix* matrix, `HYPRE_Int` \*domain\_stride)

(Optional) Set the domain coarsening stride.

For more information, see *HYPRE\_StructMatrixCreate*.

`HYPRE_Int HYPRE_StructMatrixInitialize`(*HYPRE\_StructMatrix* matrix)

Prepare a matrix object for setting coefficient values.

`HYPRE_Int HYPRE_StructMatrixSetValues`(*HYPRE\_StructMatrix* matrix, `HYPRE_Int` \*index, `HYPRE_Int` nentries, `HYPRE_Int` \*entries, `HYPRE_Complex` \*values)

Set matrix coefficients index by index.

The *values* array is of length *nentries*.

NOTE: For better efficiency, use *HYPRE\_StructMatrixSetBoxValues* to set coefficients a box at a time.

`HYPRE_Int HYPRE_StructMatrixAddToValues`(*HYPRE\_StructMatrix* matrix, `HYPRE_Int` \*index, `HYPRE_Int` nentries, `HYPRE_Int` \*entries, `HYPRE_Complex` \*values)

Add to matrix coefficients index by index.

The *values* array is of length *nentries*.

NOTE: For better efficiency, use *HYPRE\_StructMatrixAddToBoxValues* to set coefficients a box at a time.

`HYPRE_Int HYPRE_StructMatrixSetConstantValues`(*HYPRE\_StructMatrix* matrix, `HYPRE_Int` nentries, `HYPRE_Int` \*entries, `HYPRE_Complex` \*values)

Set matrix coefficients which are constant over the grid.

The *values* array is of length *nentries*.

`HYPRE_Int HYPRE_StructMatrixAddToConstantValues`(*HYPRE\_StructMatrix* matrix, `HYPRE_Int` nentries, `HYPRE_Int` \*entries, `HYPRE_Complex` \*values)

Add to matrix coefficients which are constant over the grid.

The *values* array is of length *nentries*.

`HYPRE_Int HYPRE_StructMatrixSetBoxValues`(*HYPRE\_StructMatrix* matrix, `HYPRE_Int` \*ilower, `HYPRE_Int` \*iupper, `HYPRE_Int` nentries, `HYPRE_Int` \*entries, `HYPRE_Complex` \*values)

Set matrix coefficients a box at a time.

The data in *values* is ordered as follows:

```
m = 0;
for (k = ilower[2]; k <= iupper[2]; k++)
  for (j = ilower[1]; j <= iupper[1]; j++)
    for (i = ilower[0]; i <= iupper[0]; i++)
      for (entry = 0; entry < nentries; entry++)
        {
```

(continues on next page)

(continued from previous page)

```

        values[m] = ...;
        m++;
    }

```

**HYPRE\_Int HYPRE\_StructMatrixAddToBoxValues**(*HYPRE\_StructMatrix* matrix, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Complex \*values)

Add to matrix coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_StructMatrixSetBoxValues*.

**HYPRE\_Int HYPRE\_StructMatrixSetBoxValues2**(*HYPRE\_StructMatrix* matrix, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Int \*vilower, HYPRE\_Int \*viupper, HYPRE\_Complex \*values)

Set matrix coefficients a box at a time.

The *values* array is logically box shaped with value-box extents *vilower* and *viupper* that must contain the set-box extents *ilower* and *iupper* . The data in the *values* array is ordered as in *HYPRE\_StructMatrixSetBoxValues*, but based on the value-box extents.

**HYPRE\_Int HYPRE\_StructMatrixAddToBoxValues2**(*HYPRE\_StructMatrix* matrix, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Int \*vilower, HYPRE\_Int \*viupper, HYPRE\_Complex \*values)

Add to matrix coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_StructMatrixSetBoxValues2*.

**HYPRE\_Int HYPRE\_StructMatrixAssemble**(*HYPRE\_StructMatrix* matrix)

Finalize the construction of the matrix before using.

**HYPRE\_Int HYPRE\_StructMatrixGetValues**(*HYPRE\_StructMatrix* matrix, HYPRE\_Int \*index, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Complex \*values)

Get matrix coefficients index by index.

The *values* array is of length *nentries*.

NOTE: For better efficiency, use *HYPRE\_StructMatrixGetBoxValues* to get coefficients a box at a time.

**HYPRE\_Int HYPRE\_StructMatrixGetBoxValues**(*HYPRE\_StructMatrix* matrix, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Complex \*values)

Get matrix coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_StructMatrixSetBoxValues*.

**HYPRE\_Int HYPRE\_StructMatrixGetBoxValues2**(*HYPRE\_StructMatrix* matrix, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Int \*vilower, HYPRE\_Int \*viupper, HYPRE\_Complex \*values)

Get matrix coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_StructMatrixSetBoxValues2*.

HYPRE\_Int **HYPRE\_StructMatrixSetSymmetric**(*HYPRE\_StructMatrix* matrix, HYPRE\_Int symmetric)

Define symmetry properties of the matrix.

By default, matrices are assumed to be nonsymmetric. Significant storage savings can be made if the matrix is symmetric.

HYPRE\_Int **HYPRE\_StructMatrixSetConstantEntries**(*HYPRE\_StructMatrix* matrix, HYPRE\_Int nentries, HYPRE\_Int \*entries)

Specify which stencil entries are constant over the grid.

Declaring entries to be “constant over the grid” yields significant memory savings because the value for each declared entry will only be stored once. However, not all solvers are able to utilize this feature.

Presently supported:

- no entries constant (this function need not be called)
- all entries constant
- all but the diagonal entry constant

HYPRE\_Int **HYPRE\_StructMatrixSetTranspose**(*HYPRE\_StructMatrix* matrix, HYPRE\_Int transpose)

Indicate whether the transpose coefficients should also be stored.

HYPRE\_Int **HYPRE\_StructMatrixSetNumGhost**(*HYPRE\_StructMatrix* matrix, HYPRE\_Int \*num\_ghost)

Set the ghost layer in the matrix.

HYPRE\_Int **HYPRE\_StructMatrixPrint**(const char \*filename, *HYPRE\_StructMatrix* matrix, HYPRE\_Int all)

Print the matrix to file.

This is mainly for debugging purposes.

HYPRE\_Int **HYPRE\_StructMatrixRead**(MPI\_Comm comm, const char \*filename, HYPRE\_Int \*num\_ghost, *HYPRE\_StructMatrix* \*matrix)

Read the matrix from file.

This is mainly for debugging purposes.

## Struct Vectors

HYPRE\_Int **HYPRE\_StructVectorCreate**(MPI\_Comm comm, *HYPRE\_StructGrid* grid, HYPRE\_StructVector \*vector)

The vector object.

Create a vector object.

HYPRE\_Int **HYPRE\_StructVectorDestroy**(HYPRE\_StructVector vector)

Destroy a vector object.

HYPRE\_Int **HYPRE\_StructVectorInitialize**(HYPRE\_StructVector vector)

Prepare a vector object for setting coefficient values.

HYPRE\_Int **HYPRE\_StructVectorSetValues**(HYPRE\_StructVector vector, HYPRE\_Int \*index, HYPRE\_Complex \*values)

Set vector coefficients index by index.

NOTE: For better efficiency, use *HYPRE\_StructVectorSetBoxValues* to set coefficients a box at a time.

HYPRE\_Int **HYPRE\_StructVectorSetConstantValues**(HYPRE\_StructVector vector, HYPRE\_Complex value)

Set vector coefficients to a constant value over the grid.

HYPRE\_Int **HYPRE\_StructVectorSetRandomValues**(HYPRE\_StructVector vector, HYPRE\_Int seed)

Set vector coefficients to random values between -1.0 and 1.0 over the grid.

The parameter *seed* controls the generation of random numbers.

HYPRE\_Int **HYPRE\_StructVectorAddToValues**(HYPRE\_StructVector vector, HYPRE\_Int \*index, HYPRE\_Complex \*values)

Add to vector coefficients index by index.

NOTE: For better efficiency, use *HYPRE\_StructVectorAddToBoxValues* to set coefficients a box at a time.

HYPRE\_Int **HYPRE\_StructVectorSetBoxValues**(HYPRE\_StructVector vector, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Complex \*values)

Set vector coefficients a box at a time.

The data in *values* is ordered as follows:

```
m = 0;
for (k = ilower[2]; k <= iupper[2]; k++)
  for (j = ilower[1]; j <= iupper[1]; j++)
    for (i = ilower[0]; i <= iupper[0]; i++)
    {
      values[m] = ...;
      m++;
    }
```

HYPRE\_Int **HYPRE\_StructVectorAddToBoxValues**(HYPRE\_StructVector vector, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Complex \*values)

Add to vector coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_StructVectorSetBoxValues*.

HYPRE\_Int **HYPRE\_StructVectorSetBoxValues2**(HYPRE\_StructVector vector, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int \*vilower, HYPRE\_Int \*viupper, HYPRE\_Complex \*values)

Set vector coefficients a box at a time.

The *values* array is logically box shaped with value-box extents *vilower* and *viupper* that must contain the set-box extents *ilower* and *iupper* . The data in the *values* array is ordered as in *HYPRE\_StructVectorSetBoxValues*, but based on the value-box extents.

HYPRE\_Int **HYPRE\_StructVectorAddToBoxValues2**(HYPRE\_StructVector vector, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int \*vilower, HYPRE\_Int \*viupper, HYPRE\_Complex \*values)

Add to vector coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_StructVectorSetBoxValues2*.

HYPRE\_Int **HYPRE\_StructVectorAssemble**(HYPRE\_StructVector vector)

Finalize the construction of the vector before using.

HYPRE\_Int **HYPRE\_StructVectorGetValues**(HYPRE\_StructVector vector, HYPRE\_Int \*index, HYPRE\_Complex \*value)

Get vector coefficients index by index.

NOTE: For better efficiency, use *HYPRE\_StructVectorGetBoxValues* to get coefficients a box at a time.

`HYPRE_Int HYPRE_StructVectorGetBoxValues`(`HYPRE_StructVector` vector, `HYPRE_Int` \*ilower,  
`HYPRE_Int` \*iupper, `HYPRE_Complex` \*values)

Get vector coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_StructVectorSetBoxValues*.

`HYPRE_Int HYPRE_StructVectorGetBoxValues2`(`HYPRE_StructVector` vector, `HYPRE_Int` \*ilower,  
`HYPRE_Int` \*iupper, `HYPRE_Int` \*vilower,  
`HYPRE_Int` \*viupper, `HYPRE_Complex` \*values)

Get vector coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_StructVectorSetBoxValues2*.

`HYPRE_Int HYPRE_StructVectorPrint`(`const char` \*filename, `HYPRE_StructVector` vector, `HYPRE_Int`  
all)

Print the vector to file.

This is mainly for debugging purposes.

`HYPRE_Int HYPRE_StructVectorRead`(`MPI_Comm` comm, `const char` \*filename, `HYPRE_Int`  
\*num\_ghost, `HYPRE_StructVector` \*vector)

Read the vector from file.

This is mainly for debugging purposes.

`HYPRE_Int HYPRE_StructVectorClone`(`HYPRE_StructVector` x, `HYPRE_StructVector` \*y\_ptr)

Clone a vector x.

### Basic Matrix/vector routines

`HYPRE_Int HYPRE_StructVectorCopy`(`HYPRE_StructVector` x, `HYPRE_StructVector` y)

Copy vector x into y ( $y \leftarrow x$ ).

`HYPRE_Int HYPRE_StructVectorScale`(`HYPRE_Complex` alpha, `HYPRE_StructVector` y)

Scale a vector by *alpha* ( $y \leftarrow \alpha y$ ).

`HYPRE_Int HYPRE_StructVectorAxpv`(`HYPRE_Complex` alpha, `HYPRE_StructVector` x,  
`HYPRE_Complex` beta, `HYPRE_StructVector` y)

Compute  $y = y + \alpha x$ .

`HYPRE_Int HYPRE_StructVectorInnerProd`(`HYPRE_StructVector` x, `HYPRE_StructVector` y,  
`HYPRE_Real` \*result)

Compute *result*, the inner product of vectors x and y.

`HYPRE_Int HYPRE_StructMatrixMatvec`(`HYPRE_Complex` alpha, *HYPRE\_StructMatrix* A,  
`HYPRE_StructVector` x, `HYPRE_Complex` beta,  
`HYPRE_StructVector` y)

Matvec operator.

This operation is  $y = \alpha Ax + \beta y$ . Note that you can do a simple matrix-vector multiply by setting  $\alpha = 1$  and  $\beta = 0$ .

HYPRE\_Int **HYPRE\_StructMatrixMatvecT**(HYPRE\_Complex alpha, *HYPRE\_StructMatrix* A, HYPRE\_StructVector x, HYPRE\_Complex beta, HYPRE\_StructVector y)

Matvec transpose operation.

This operation is  $y = \alpha A^T x + \beta y$ . Note that you can do a simple matrix-vector multiply by setting  $\alpha = 1$  and  $\beta = 0$ .

HYPRE\_Int **HYPRE\_StructMatrixMatmat**(*HYPRE\_StructMatrix* A, HYPRE\_Int Atranspose, *HYPRE\_StructMatrix* B, HYPRE\_Int Btranspose, *HYPRE\_StructMatrix* \*C)

Matrix-matrix multiply.

Returns  $C = AB$ ,  $C = A^T B$ ,  $C = AB^T$ , or  $C = A^T B^T$ , depending on the boolean arguments *Atranspose* and *Btranspose*.

## 8.2 SStruct System and Object Interface

### group SStruct System and Object Interface

A semi-structured-grid conceptual interface.

This interface represents a semi-structured-grid conceptual view of a linear system.

### SStruct Grids

typedef struct hypr\_SStructGrid\_struct \***HYPRE\_SStructGrid**

A grid object is constructed out of several structured “parts” and an optional unstructured “part”.

Each structured part has its own abstract index space.

typedef HYPRE\_Int **HYPRE\_SStructVariable**

An enumerated type that supports cell centered, node centered, face centered, and edge centered variables.

Face centered variables are split into x-face, y-face, and z-face variables, and edge centered variables are split into x-edge, y-edge, and z-edge variables. The edge centered variable types are only used in 3D. In 2D, edge centered variables are handled by the face centered types.

Variables are referenced relative to an abstract (cell centered) index in the following way:

- cell centered variables are aligned with the index;
- node centered variables are aligned with the cell corner at relative index (1/2, 1/2, 1/2);
- x-face, y-face, and z-face centered variables are aligned with the faces at relative indexes (1/2, 0, 0), (0, 1/2, 0), and (0, 0, 1/2), respectively;
- x-edge, y-edge, and z-edge centered variables are aligned with the edges at relative indexes (0, 1/2, 1/2), (1/2, 0, 1/2), and (1/2, 1/2, 0), respectively.

The supported identifiers are:

- HYPRE\_SSTRUCT\_VARIABLE\_CELL
- HYPRE\_SSTRUCT\_VARIABLE\_NODE
- HYPRE\_SSTRUCT\_VARIABLE\_XFACE
- HYPRE\_SSTRUCT\_VARIABLE\_YFACE

- HYPRE\_SSTRUCT\_VARIABLE\_ZFACE
- HYPRE\_SSTRUCT\_VARIABLE\_XEDGE
- HYPRE\_SSTRUCT\_VARIABLE\_YEDGE
- HYPRE\_SSTRUCT\_VARIABLE\_ZEDGE

NOTE: Although variables are referenced relative to a unique abstract cell-centered index, some variables are associated with multiple grid cells. For example, node centered variables in 3D are associated with 8 cells (away from boundaries). Although grid cells are distributed uniquely to different processes, variables may be owned by multiple processes because they may be associated with multiple cells.

HYPRE\_Int **HYPRE\_SStructGridCreate**(MPI\_Comm comm, HYPRE\_Int ndim, HYPRE\_Int nparts, *HYPRE\_SStructGrid \*grid*)

Create an *ndim*-dimensional grid object with *nparts* structured parts.

HYPRE\_Int **HYPRE\_SStructGridDestroy**(*HYPRE\_SStructGrid grid*)

Destroy a grid object.

An object should be explicitly destroyed using this destructor when the user's code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

HYPRE\_Int **HYPRE\_SStructGridSetExtents**(*HYPRE\_SStructGrid grid*, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper)

Set the extents for a box on a structured part of the grid.

HYPRE\_Int **HYPRE\_SStructGridSetVariables**(*HYPRE\_SStructGrid grid*, HYPRE\_Int part, HYPRE\_Int nvars, *HYPRE\_SStructVariable \*vartypes*)

Describe the variables that live on a structured part of the grid.

HYPRE\_Int **HYPRE\_SStructGridAddVariables**(*HYPRE\_SStructGrid grid*, HYPRE\_Int part, HYPRE\_Int \*index, HYPRE\_Int nvars, *HYPRE\_SStructVariable \*vartypes*)

Describe additional variables that live at a particular index.

These variables are appended to the array of variables set in *HYPRE\_SStructGridSetVariables*, and are referenced as such.

NOTE: This routine is not yet supported.

HYPRE\_Int **HYPRE\_SStructGridSetFEMOrdering**(*HYPRE\_SStructGrid grid*, HYPRE\_Int part, HYPRE\_Int \*ordering)

Set the ordering of variables in a finite element problem.

This overrides the default ordering described below.

Array *ordering* is composed of blocks of size  $(1 + ndim)$ . Each block indicates a specific variable in the element and the ordering of the blocks defines the ordering of the variables. A block contains a variable number followed by an offset direction relative to the element's center. For example, a block containing  $(2, 1, -1, 0)$  means variable 2 on the edge located in the  $(1, -1, 0)$  direction from the center of the element. Note that here variable 2 must be of type ZEDGE for this to make sense. The *ordering* array must account for all variables in the element. This routine can only be called after *HYPRE\_SStructGridSetVariables*.

The default ordering for element variables (var, i, j, k) varies fastest in index i, followed by j, then k, then var. For example, if var 0, var 1, and var 2 are declared to be XFACE, YFACE, and NODE variables, respectively, then the default ordering (in 2D) would first list the two XFACE variables, then the two YFACE variables, then the four NODE variables as follows:

(0,-1,0), (0,1,0), (1,0,-1), (1,0,1), (2,-1,-1), (2,1,-1), (2,-1,1), (2,1,1)

HYPRE\_Int HYPRE\_SStructGridSetNeighborPart(*HYPRE\_SStructGrid* grid, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int nbor\_part, HYPRE\_Int \*nbor\_ilower, HYPRE\_Int \*nbor\_iupper, HYPRE\_Int \*index\_map, HYPRE\_Int \*index\_dir)

Describe how regions just outside of a part relate to other parts.

This is done a box at a time.

Parts *part* and *nbor\_part* must be different, except in the case where only cell-centered data is used.

Indexes should increase from *ilower* to *iupper*. It is not necessary that indexes increase from *nbor\_ilower* to *nbor\_iupper*.

The *index\_map* describes the mapping of indexes 0, 1, and 2 on part *part* to the corresponding indexes on part *nbor\_part*. For example, triple (1, 2, 0) means that indexes 0, 1, and 2 on part *part* map to indexes 1, 2, and 0 on part *nbor\_part*, respectively.

The *index\_dir* describes the direction of the mapping in *index\_map*. For example, triple (1, 1, -1) means that for indexes 0 and 1, increasing values map to increasing values on *nbor\_part*, while for index 2, decreasing values map to increasing values.

NOTE: All parts related to each other via this routine must have an identical list of variables and variable types. For example, if part 0 has only two variables on it, a cell centered variable and a node centered variable, and we declare part 1 to be a neighbor of part 0, then part 1 must also have only two variables on it, and they must be of type cell and node. In addition, variables associated with FACES or EDGES must be grouped together and listed in X, Y, Z order. This is to enable the code to correctly associate variables on one part with variables on its neighbor part when a coordinate transformation is specified. For example, an XFACE variable on one part may correspond to a YFACE variable on a neighbor part under a particular tranformation, and the code determines this association by assuming that the variable lists are as noted here.

HYPRE\_Int HYPRE\_SStructGridSetSharedPart(*HYPRE\_SStructGrid* grid, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int \*offset, HYPRE\_Int shared\_part, HYPRE\_Int \*shared\_ilower, HYPRE\_Int \*shared\_iupper, HYPRE\_Int \*shared\_offset, HYPRE\_Int \*index\_map, HYPRE\_Int \*index\_dir)

Describe how regions inside a part are shared with regions in other parts.

Parts *part* and *shared\_part* must be different.

Indexes should increase from *ilower* to *iupper*. It is not necessary that indexes increase from *shared\_ilower* to *shared\_iupper*. This is to maintain consistency with the `SetNeighborPart` function, which is also able to describe shared regions but in a more limited fashion.

The *offset* is a triple (in 3D) used to indicate the dimensionality of the shared set of data and its position with respect to the box extents *ilower* and *iupper* on part *part*. The dimensionality is given by the number of 0's in the triple, and the position is given by plus or minus 1's. For example: (0, 0, 0) indicates sharing of all data in the given box; (1, 0, 0) indicates sharing of data on the faces in the (1, 0, 0) direction; (1, 0, -1) indicates sharing of data on the edges in the (1, 0, -1) direction; and (1, -1, 1) indicates sharing of data on the nodes in the (1, -1, 1) direction. To ensure the dimensionality, it is required that for every nonzero entry, the corresponding extents of the box are the same. For example, if *offset* is (0, 1, 0), then (2, 1, 3) and (10, 1, 15) are valid box extents, whereas (2, 1, 3) and (10, 7, 15) are invalid (because 1 and 7 are not the same).

The *shared\_offset* is used in the same way as *offset*, but with respect to the box extents *shared\_ilower* and *shared\_iupper* on part *shared\_part*.

The *index\_map* describes the mapping of indexes 0, 1, and 2 on part *part* to the corresponding indexes on part *shared\_part*. For example, triple (1, 2, 0) means that indexes 0, 1, and 2 on part *part* map to indexes 1, 2, and 0 on part *shared\_part*, respectively.

The *index\_dir* describes the direction of the mapping in *index\_map*. For example, triple (1, 1, -1) means that for indexes 0 and 1, increasing values map to increasing values on *shared\_part*, while for index 2, decreasing values map to increasing values.

NOTE: All parts related to each other via this routine must have an identical list of variables and variable types. For example, if part 0 has only two variables on it, a cell centered variable and a node centered variable, and we declare part 1 to have shared regions with part 0, then part 1 must also have only two variables on it, and they must be of type cell and node. In addition, variables associated with FACES or EDGES must be grouped together and listed in X, Y, Z order. This is to enable the code to correctly associate variables on one part with variables on a shared part when a coordinate transformation is specified. For example, an XFACE variable on one part may correspond to a YFACE variable on a shared part under a particular transformation, and the code determines this association by assuming that the variable lists are as noted here.

```
HYPRE_Int HYPRE_SStructGridSetAMRPart(HYPRE_SStructGrid grid, HYPRE_Int coarse_part,
                                       HYPRE_Int fine_part, HYPRE_Int *coarse_index,
                                       HYPRE_Int *fine_index, HYPRE_Int *rfactors)
```

AMRNEW.

Declare a part to be a refinement of another part in an AMR hierarchy.

The index space of *fine\_part* is defined to be a refinement of the index space of *coarse\_part* by a refinement factor in each dimension given by *rfactors*. The two index spaces are aligned based on *coarse\_index* and *fine\_index*, which specifies the fine index of the lower left cell of the given coarse index. This induces notions of *real* and *slave* variables and interpolation between them, which impacts how vector and matrix values are set.

By default, real variables are defined as follows, and the remaining variables are slave variables (this may be changed by the user through the routine *HYPRE\_SStructGridSetAMRRefSlaves*):

- fine variables on the interior of overlapping coarse-fine regions;
- coarse variables on the boundary of overlapping coarse-fine regions;
- all variables in non-overlapping regions.

By default, interpolation (and restriction) is defined to be the natural finite element interpolation corresponding to each variable type, but this may be changed by the user via the other *SStructGridSetAMR* routines.

There are two basic steps for changing interpolation, both optional. The first is through the *SStructGridSetAMRRef* routines using a reference overlapping coarse-fine patch. This reference patch is applied throughout the entire part to define the global interpolation operator. Interpolation may then be changed at individual locations in the grid through the *HYPRE\_SStructGridSetAMRInterp* routine.

The reference coarse-fine patch consists of a single coarse cell and its refinement, where the coarse and fine reference patches are assumed to have a lower left index of zero. Coarse and fine variables are referenced in the patch by their associated cell indexes (in the same way that variables are referenced on the grid). For example, in 2D, for a refinement factor of two in both directions and a nodal variable type, the lower left fine variable would be referenced by the index (-1,-1) and the upper right variable by index (1,1). Similarly, the lower left coarse variable would be referenced by index (-1,-1) and the upper right with index (0,0).

```
HYPRE_Int HYPRE_SStructGridSetAMRRefSlaves(HYPRE_SStructGrid grid, HYPRE_Int coarse_part,
                                             HYPRE_Int var, HYPRE_Int nslaves, HYPRE_Int
                                             *slaves)
```

AMRNEW.

Define the slave variables in the reference coarse-fine patch. The argument *slaves* is an array of blocks of size *ndim* containing the associated cell indexes for the slave variables.

See *HYPRE\_SStructGridSetAMRPart* for details on the coarse-fine patch. This routine must be called after *HYPRE\_SStructGridSetAMRPart* and before any other *SStructGridSetAMR* routines.

```
HYPRE_Int HYPRE_SStructGridSetAMRRefInterp(HYPRE_SStructGrid grid, HYPRE_Int coarse_part,
                                             HYPRE_Int scf, HYPRE_Int svar, HYPRE_Int
                                             *sindex, HYPRE_Int nvalues, HYPRE_Int *vars,
                                             HYPRE_Int *cf, HYPRE_Int *indexes,
                                             HYPRE_Complex *values)
```

AMRNEW.

Set interpolation on the reference coarse-fine patch. Interpolation maps real variables to all variables (real and slave). Real variables are mapped to real variables identically. Slave coarse variables are not interpolated at all. Users may only change interpolation from real variables (coarse and fine) to slave fine variables *scf=1* (see special cell-centered case below). The argument *sindex* is a fine reference patch index for slave variable *svar*. The array *indexes* contains both coarse and fine reference patch indexes for variables *vars* as specified in *cf* by a 0 (coarse) or a 1 (fine).

Cell-centered variables are treated as a special case, where coupling occurs through fictitious slave face variables. Although these variables are associated with faces, they should actually be thought of as being centered either at cells just outside of the patch for fine faces (*scf=1*) or at the patch center for coarse faces (*scf=0*). These fictitious variables are also referenced differently from other variables, where, for convenience, *sindex* always specifies cells just outside of the patch (in the coarse case in particular, this approach enables a distinction between the multiple coarse faces and interpolation formulas). Interpolation may be from any neighboring real variables, inside or outside of the patch.

Idea for the cell-centered case: To reduce the need to fix up interpolation at individual grid locations, consider adding a rule to replace interpolation from nonexistent coarse variables with interpolation from existing underlying fine variables. For example, in those cases, take the average of the closest fine variables (in the odd refinement case, this is just injection from the underlying fine variable).

See *HYPRE\_SStructGridSetAMRPart* for details on the coarse-fine patch. Note: Currently, each call must set an entire row of interpolation.

```
HYPRE_Int HYPRE_SStructGridSetAMRRefRestrictT(HYPRE_SStructGrid grid, HYPRE_Int
                                                coarse_part, HYPRE_Int scf, HYPRE_Int svar,
                                                HYPRE_Int *sindex, HYPRE_Int nvalues,
                                                HYPRE_Int *vars, HYPRE_Int *cf, HYPRE_Int
                                                *indexes, HYPRE_Complex *values)
```

AMRNEW.

Set (the transpose of) restriction on the reference coarse-fine patch. By default, restriction is the transpose of interpolation. Usually, this routine should only be called to set up a nonsymmetric operator. Usage is the same as for setting up interpolation.

```
HYPRE_Int HYPRE_SStructGridSetAMRInterp(HYPRE_SStructGrid grid, HYPRE_Int coarse_part,
                                           HYPRE_Int *coarse_index, HYPRE_Int scf, HYPRE_Int
                                           svar, HYPRE_Int *sindex, HYPRE_Int nvalues,
                                           HYPRE_Int *vars, HYPRE_Int *cf, HYPRE_Int *indexes,
                                           HYPRE_Complex *values)
```

AMRNEW.

Set interpolation at a specific grid index on a refined part. Usage is the same as for setting interpolation on the reference coarse-fine patch, except that an additional argument *coarse\_index* is given.

HYPRE\_Int **HYPRE\_SStructGridSetAMRRestrictT**(*HYPRE\_SStructGrid* grid, HYPRE\_Int coarse\_part, HYPRE\_Int \*coarse\_index, HYPRE\_Int scf, HYPRE\_Int svar, HYPRE\_Int \*sindex, HYPRE\_Int nvalues, HYPRE\_Int \*vars, HYPRE\_Int \*cf, HYPRE\_Int \*indexes, HYPRE\_Complex \*values)

AMRNEW.

Set restriction at a specific grid index on a refined part. Usage is the same as for setting restriction on the reference coarse-fine patch, except that an additional argument *coarse\_index* is given.

HYPRE\_Int **HYPRE\_SStructGridAddUnstructuredPart**(*HYPRE\_SStructGrid* grid, HYPRE\_Int ilower, HYPRE\_Int iupper)

Add an unstructured part to the grid.

The variables in the unstructured part of the grid are referenced by a global rank between 0 and the total number of unstructured variables minus one. Each process owns some unique consecutive range of variables, defined by *ilower* and *iupper*.

NOTE: This is just a placeholder. This part of the interface is not finished.

HYPRE\_Int **HYPRE\_SStructGridAssemble**(*HYPRE\_SStructGrid* grid)

Finalize the construction of the grid before using.

HYPRE\_Int **HYPRE\_SStructGridSetObjectType**(*HYPRE\_SStructGrid* grid, HYPRE\_Int type)

AMRNEW.

Set the storage type of associated matrix and vector objects. This object type is inherited by *SStructGraph*, *SStructMatrix*, and *SStructVector* at their creation, but can also be changed through the corresponding *SetObjectType* routines. Setting an object type with this routine helps to minimize the number of overall calls to *SetObjectType*, but is particularly useful when grid-based operators are needed, such as a discrete gradient (*HYPRE\_SStructGridGetGradient*).

Currently, *type* can be HYPRE\_SSTRUCT, HYPRE\_STRUCTURE, or HYPRE\_PARCSR. The default is HYPRE\_SSTRUCT.

HYPRE\_Int **HYPRE\_SStructGridGetGradient**(*HYPRE\_SStructGrid* grid, void \*\*gradient)

AMRNEW.

Get a reference to a discrete gradient matrix object.

RDF: Not sure yet if this routine really needs to exist.

HYPRE\_Int **HYPRE\_SStructGridSetPeriodic**(*HYPRE\_SStructGrid* grid, HYPRE\_Int part, HYPRE\_Int \*periodic)

Set the periodicity on a particular part.

The argument *periodic* is an *ndim-dimensional* integer array that contains the periodicity for each dimension. A zero value for a dimension means non-periodic, while a nonzero value means periodic and contains the actual period. For example, periodicity in the first and third dimensions for a 10x11x12 part is indicated by the array [10,0,12].

NOTE: Some of the solvers in hypr have power-of-two restrictions on the size of the periodic dimensions.

HYPRE\_Int **HYPRE\_SStructGridSetNumGhost**(*HYPRE\_SStructGrid* grid, HYPRE\_Int \*num\_ghost)

Setting ghost in the sgrids.

HYPRE\_Int **HYPRE\_SStructGridGetVariableBox**(*HYPRE\_SStructGrid* grid, HYPRE\_Int part, HYPRE\_Int var, HYPRE\_Int \*cell\_ilower, HYPRE\_Int \*cell\_iupper, HYPRE\_Int \*var\_ilower, HYPRE\_Int \*var\_iupper)

Helper function for returning the box dimensions of a (part,var)-structured grid.

HYPRE\_Int **HYPRE\_SStructGridProjectBox**(*HYPRE\_SStructGrid* grid, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int \*origin, HYPRE\_Int \*stride)

Project the box described by *ilower* and *iupper* onto the strided index space that contains the index *origin* and has stride *stride*.

This routine is useful when dealing with rectangular matrices.

HYPRE\_Int **HYPRE\_SStructGridCoarsen**(*HYPRE\_SStructGrid* fgrid, HYPRE\_Index \*strides, *HYPRE\_SStructGrid* \*cgrid)

Coarsen *grid* by factor *strides* to create *cgrid*.

HYPRE\_Int **HYPRE\_SStructGridPrintGLVis**(*HYPRE\_SStructGrid* grid, const char \*meshprefix, HYPRE\_Real \*trans, HYPRE\_Real \*origin)

HYPRE\_SSTRUCT\_VARIABLE\_UNDEFINED

HYPRE\_SSTRUCT\_VARIABLE\_CELL

HYPRE\_SSTRUCT\_VARIABLE\_NODE

HYPRE\_SSTRUCT\_VARIABLE\_XFACE

HYPRE\_SSTRUCT\_VARIABLE\_YFACE

HYPRE\_SSTRUCT\_VARIABLE\_ZFACE

HYPRE\_SSTRUCT\_VARIABLE\_XEDGE

HYPRE\_SSTRUCT\_VARIABLE\_YEDGE

HYPRE\_SSTRUCT\_VARIABLE\_ZEDGE

### SStruct Stencils

typedef struct hypr\_SStructStencil\_struct \***HYPRE\_SStructStencil**

The stencil object.

HYPRE\_Int **HYPRE\_SStructStencilCreate**(HYPRE\_Int ndim, HYPRE\_Int size, *HYPRE\_SStructStencil* \*stencil)

Create a stencil object for the specified number of spatial dimensions and stencil entries.

HYPRE\_Int **HYPRE\_SStructStencilDestroy**(*HYPRE\_SStructStencil* stencil)

Destroy a stencil object.

HYPRE\_Int **HYPRE\_SStructStencilSetEntry**(*HYPRE\_SStructStencil* stencil, HYPRE\_Int entry, HYPRE\_Int \*offset, HYPRE\_Int var)

Set a stencil entry.

HYPRE\_Int **HYPRE\_SStructStencilPrint**(FILE \*file, *HYPRE\_SStructStencil* stencil)

HYPRE\_Int **HYPRE\_SStructStencilRead**(FILE \*file, *HYPRE\_SStructStencil* \*stencil\_ptr)

## SStruct Graphs

typedef struct hypre\_SStructGraph\_struct **\*HYPRE\_SStructGraph**

The graph object is used to describe the nonzero structure of a matrix.

HYPRE\_Int **HYPRE\_SStructGraphCreate**(MPI\_Comm comm, *HYPRE\_SStructGrid* grid,  
*HYPRE\_SStructGraph* \*graph)

Create a graph object.

HYPRE\_Int **HYPRE\_SStructGraphDestroy**(*HYPRE\_SStructGraph* graph)

Destroy a graph object.

HYPRE\_Int **HYPRE\_SStructGraphSetDomainGrid**(*HYPRE\_SStructGraph* graph, *HYPRE\_SStructGrid*  
domain\_grid)

Set the domain grid.

HYPRE\_Int **HYPRE\_SStructGraphSetStencil**(*HYPRE\_SStructGraph* graph, HYPRE\_Int part,  
HYPRE\_Int var, *HYPRE\_SStructStencil* stencil)

Set the stencil for a variable on a structured part of the grid.

HYPRE\_Int **HYPRE\_SStructGraphSetFEM**(*HYPRE\_SStructGraph* graph, HYPRE\_Int part)

Indicate that an FEM approach will be used to set matrix values on this part.

HYPRE\_Int **HYPRE\_SStructGraphSetFEMSparsity**(*HYPRE\_SStructGraph* graph, HYPRE\_Int part,  
HYPRE\_Int nparse, HYPRE\_Int \*sparsity)

Set the finite element stiffness matrix sparsity.

This overrides the default full sparsity pattern described below.

Array *sparsity* contains *nparse* row/column tuples (I,J) that indicate the nonzeros of the local stiffness matrix. The layout of the values passed into the routine *HYPRE\_SStructMatrixAddFEMValues* is determined here.

The default sparsity is full (each variable is coupled to all others), and the values passed into the routine *HYPRE\_SStructMatrixAddFEMValues* are assumed to be by rows (that is, column indices vary fastest).

HYPRE\_Int **HYPRE\_SStructGraphAddEntries**(*HYPRE\_SStructGraph* graph, HYPRE\_Int part,  
HYPRE\_Int \*index, HYPRE\_Int var, HYPRE\_Int to\_part,  
HYPRE\_Int \*to\_index, HYPRE\_Int to\_var)

Add a non-stencil graph entry at a particular index.

This graph entry is appended to the existing graph entries, and is referenced as such.

NOTE: Users are required to set graph entries on all processes that own the associated variables. This means that some data will be multiply defined.

HYPRE\_Int **HYPRE\_SStructGraphAssemble**(*HYPRE\_SStructGraph* graph)

Finalize the construction of the graph before using.

HYPRE\_Int **HYPRE\_SStructGraphSetObjectType**(*HYPRE\_SStructGraph* graph, HYPRE\_Int type)

Set the storage type of the associated matrix object.

It is used before AddEntries and Assemble to compute the right ranks in the graph.

NOTE: This routine is only necessary for implementation reasons, and will eventually be removed.

➔ See also

*HYPRE\_SStructMatrixSetObjectType*

HYPRE\_Int **HYPRE\_SStructGraphPrint**(FILE \*file, *HYPRE\_SStructGraph* graph)

HYPRE\_Int **HYPRE\_SStructGraphRead**(FILE \*file, *HYPRE\_SStructGrid* grid, *HYPRE\_SStructStencil* \*\*stencils, *HYPRE\_SStructGraph* \*graph\_ptr)

## SStruct Matrices

```
typedef struct hypre_SStructMatrix_struct *HYPRE_SStructMatrix
```

The matrix object.

HYPRE\_Int **HYPRE\_SStructMatrixCreate**(MPI\_Comm comm, *HYPRE\_SStructGraph* graph, *HYPRE\_SStructMatrix* \*matrix)

Create a matrix object.

Matrices may have different range and domain grids, that is, they need not be square. By default, the range and domain grids are the same as the argument *grid* in *HYPRE\_SStructMatrixCreate*. Both grids live on a common fine index space and should have the same number of boxes. The actual range is a coarsening of the range grid with coarsening factor *ran\_stride* specified in *HYPRE\_SStructMatrixSetRangeStride*. Similarly, the actual domain is a coarsening of the domain grid with factor *dom\_stride* specified in *HYPRE\_SStructMatrixSetDomainStride*. Currently, either *ran\_stride* or *dom\_stride* or both must be all ones (i.e., no coarsening).

HYPRE\_Int **HYPRE\_SStructMatrixDestroy**(*HYPRE\_SStructMatrix* matrix)

Destroy a matrix object.

HYPRE\_Int **HYPRE\_SStructMatrixInitialize**(*HYPRE\_SStructMatrix* matrix)

Prepare a matrix object for setting coefficient values.

HYPRE\_Int **HYPRE\_SStructMatrixSetDomainStride**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*dom\_stride)

(Optional) Set matrix domain stride part by part.

For more information, see *HYPRE\_SStructMatrixCreate*.

HYPRE\_Int **HYPRE\_SStructMatrixSetRangeStride**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*ran\_stride)

(Optional) Set matrix range stride part by part.

For more information, see *HYPRE\_SStructMatrixCreate*.

HYPRE\_Int **HYPRE\_SStructMatrixSetConstantEntries**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int var, HYPRE\_Int to\_var, HYPRE\_Int nentries, HYPRE\_Int \*centries)

(Optional) Set matrix coefficients which are constant over the grid.

This considers a block matrix associated with the tuple (part, var, to\_var). The *centries* array is of length *nentries*.

HYPRE\_Int **HYPRE\_SStructMatrixSetEarlyAssemble**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int early\_assemble)

(Optional, GPU only) Sets if the matrix assemble routine does reductions of the IJ part before calling HYPRE\_SStructMatrixAssemble.

See also the comments of HYPRE\_IJMatrixSetEarlyAssemble. This early assemble feature may save the peak memory usage but requires extra work.

HYPRE\_Int **HYPRE\_SStructMatrixSetValues**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*index, HYPRE\_Int var, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Complex \*values)

Set matrix coefficients index by index.

The *values* array is of length *nentries*.

NOTE: For better efficiency, use *HYPRE\_SStructMatrixSetBoxValues* to set coefficients a box at a time.

NOTE: Users are required to set values on all processes that own the associated variables. This means that some data will be multiply defined.

NOTE: The entries in this routine must all be of the same type: either stencil or non-stencil, but not both. Also, if they are stencil entries, they must all represent couplings to the same variable type (there are no such restrictions for non-stencil entries).

HYPRE\_Int **HYPRE\_SStructMatrixAddToValues**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*index, HYPRE\_Int var, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Complex \*values)

Add to matrix coefficients index by index.

The *values* array is of length *nentries*.

NOTE: For better efficiency, use *HYPRE\_SStructMatrixAddToBoxValues* to set coefficients a box at a time.

NOTE: Users are required to set values on all processes that own the associated variables. This means that some data will be multiply defined.

NOTE: The entries in this routine must all be of the same type: either stencil or non-stencil, but not both. Also, if they are stencil entries, they must all represent couplings to the same variable type.

HYPRE\_Int **HYPRE\_SStructMatrixAddFEMValues**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*index, HYPRE\_Complex \*values)

Add finite element stiffness matrix coefficients index by index.

The layout of the data in *values* is determined by the routines *HYPRE\_SStructGridSetFEMOrdering* and *HYPRE\_SStructGraphSetFEMSparsity*.

NOTE: For better efficiency, use *HYPRE\_SStructMatrixAddFEMBoxValues* to set coefficients a box at a time.

HYPRE\_Int **HYPRE\_SStructMatrixGetValues**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*index, HYPRE\_Int var, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Complex \*values)

Get matrix coefficients index by index.

The *values* array is of length *nentries*.

NOTE: For better efficiency, use *HYPRE\_SStructMatrixGetBoxValues* to get coefficients a box at a time.

NOTE: Users may get values on any process that owns the associated variables.

NOTE: The entries in this routine must all be of the same type: either stencil or non-stencil, but not both. Also, if they are stencil entries, they must all represent couplings to the same variable type (there are no such restrictions for non-stencil entries).

HYPRE\_Int **HYPRE\_SStructMatrixGetFEMValues**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*index, HYPRE\_Complex \*values)

Get finite element stiffness matrix coefficients index by index.

The layout of the data in *values* is determined by the routines *HYPRE\_SStructGridSetFEMOrdering* and *HYPRE\_SStructGraphSetFEMSparsity*.

HYPRE\_Int **HYPRE\_SStructMatrixSetBoxValues**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Complex \*values)

Set matrix coefficients a box at a time.

The data in *values* is ordered as follows:

```
m = 0;
for (k = ilower[2]; k <= iupper[2]; k++)
  for (j = ilower[1]; j <= iupper[1]; j++)
    for (i = ilower[0]; i <= iupper[0]; i++)
      for (entry = 0; entry < nentries; entry++)
        {
          values[m] = ...;
          m++;
        }
```

NOTE: Users are required to set values on all processes that own the associated variables. This means that some data will be multiply defined.

NOTE: The entries in this routine must all be of the same type: either stencil or non-stencil, but not both. Also, if they are stencil entries, they must all represent couplings to the same variable type (there are no such restrictions for non-stencil entries).

HYPRE\_Int **HYPRE\_SStructMatrixAddToBoxValues**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Complex \*values)

Add to matrix coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_SStructMatrixSetBoxValues*.

NOTE: Users are required to set values on all processes that own the associated variables. This means that some data will be multiply defined.

NOTE: The entries in this routine must all be of stencil type. Also, they must all represent couplings to the same variable type.

HYPRE\_Int **HYPRE\_SStructMatrixSetBoxValues2**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Int \*vilower, HYPRE\_Int \*viupper, HYPRE\_Complex \*values)

Set matrix coefficients a box at a time.

The *values* array is logically box shaped with value-box extents *vilower* and *viupper* that must contain the set-box extents *ilower* and *iupper* . The data in the *values* array is ordered as in *HYPRE\_SStructMatrixSetBoxValues*, but based on the value-box extents.

HYPRE\_Int **HYPRE\_SStructMatrixAddToBoxValues2**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Int \*vilower, HYPRE\_Int \*viupper, HYPRE\_Complex \*values)

Add to matrix coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_SStructMatrixSetBoxValues2*.

HYPRE\_Int **HYPRE\_SStructMatrixAddFEMBoxValues**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Complex \*values)

Add finite element stiffness matrix coefficients a box at a time.

The data in *values* is organized as an array of element matrices ordered as in *HYPRE\_SStructMatrixSetBoxValues*. The layout of the data entries of each element matrix is determined by the routines *HYPRE\_SStructGridSetFEMOrdering* and *HYPRE\_SStructGraphSetFEMSparsity*.

HYPRE\_Int **HYPRE\_SStructMatrixAssemble**(*HYPRE\_SStructMatrix* matrix)

Finalize the construction of the matrix before using.

HYPRE\_Int **HYPRE\_SStructMatrixGetBoxValues**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Complex \*values)

Get matrix coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_SStructMatrixSetBoxValues*.

NOTE: Users may get values on any process that owns the associated variables.

NOTE: The entries in this routine must all be of stencil type. Also, they must all represent couplings to the same variable type.

HYPRE\_Int **HYPRE\_SStructMatrixGetBoxValues2**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Int nentries, HYPRE\_Int \*entries, HYPRE\_Int \*vilower, HYPRE\_Int \*viupper, HYPRE\_Complex \*values)

Get matrix coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_SStructMatrixSetBoxValues2*.

HYPRE\_Int **HYPRE\_SStructMatrixGetFEMBoxValues**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Complex \*values)

Does this even make sense to implement?

HYPRE\_Int **HYPRE\_SStructMatrixSetSymmetric**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int part, HYPRE\_Int var, HYPRE\_Int to\_var, HYPRE\_Int symmetric)

Define symmetry properties for the stencil entries in the matrix.

The boolean argument *symmetric* is applied to stencil entries on part *part* that couple variable *var* to variable *to\_var*. A value of -1 may be used for *part*, *var*, or *to\_var* to specify “all”. For example, if *part* and *to\_var*

are set to -1, then the boolean is applied to stencil entries on all parts that couple variable *var* to all other variables.

By default, matrices are assumed to be nonsymmetric. Significant storage savings can be made if the matrix is symmetric.

HYPRE\_Int **HYPRE\_SStructMatrixSetNSSymmetric**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int symmetric)

Define symmetry properties for all non-stencil matrix entries.

HYPRE\_Int **HYPRE\_SStructMatrixSetObjectType**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int type)

Set the storage type of the matrix object to be constructed.

Currently, *type* can be either HYPRE\_SSTRUCT (the default), HYPRE\_STRUCT, or HYPRE\_PARCSR.

➔ See also

*HYPRE\_SStructMatrixGetObject*

HYPRE\_Int **HYPRE\_SStructMatrixGetObject**(*HYPRE\_SStructMatrix* matrix, void \*\*object)

Get a reference to the constructed matrix object.

➔ See also

*HYPRE\_SStructMatrixSetObjectType*

HYPRE\_Int **HYPRE\_SStructMatrixGetGrid**(*HYPRE\_SStructMatrix* matrix, *HYPRE\_SStructGrid* \*grid)

Returns the grid object of a HYPRE\_SStructMatrix.

HYPRE\_Int **HYPRE\_SStructMatrixPrint**(const char \*filename, *HYPRE\_SStructMatrix* matrix, HYPRE\_Int all)

Print the matrix to file.

This is mainly for debugging purposes.

HYPRE\_Int **HYPRE\_SStructMatrixToIJMatrix**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Int fill\_diagonal, *HYPRE\_IJMatrix* \*ijmatrix)

Converts a SStructMatrix to an IJMatrix.

This will sum the structured and unstructured components of the input SStructMatrix and construct the resulting sum as an IJMatrix. When the flag *fill\_diagonal* is turned on, the diagonal coefficient of ghost rows is set to 1.

HYPRE\_Int **HYPRE\_SStructMatrixRead**(MPI\_Comm comm, const char \*filename, *HYPRE\_SStructMatrix* \*matrix\_ptr)

Read the matrix from file.

This is mainly for debugging purposes.

## SStruct Vectors

```
typedef struct hypr_SStructVector_struct *HYPRE_SStructVector
```

The vector object.

```
HYPRE_Int HYPRE_SStructVectorCreate(MPI_Comm comm, HYPRE_SStructGrid grid,
                                     HYPRE_SStructVector *vector)
```

Create a vector object.

```
HYPRE_Int HYPRE_SStructVectorDestroy(HYPRE_SStructVector vector)
```

Destroy a vector object.

```
HYPRE_Int HYPRE_SStructVectorInitialize(HYPRE_SStructVector vector)
```

Prepare a vector object for setting coefficient values.

```
HYPRE_Int HYPRE_SStructVectorSetValues(HYPRE_SStructVector vector, HYPRE_Int part,
                                       HYPRE_Int *index, HYPRE_Int var, HYPRE_Complex
                                       *value)
```

Set vector coefficients index by index.

NOTE: For better efficiency, use *HYPRE\_SStructVectorSetBoxValues* to set coefficients a box at a time.

NOTE: Users are required to set values on all processes that own the associated variables. This means that some data will be multiply defined.

```
HYPRE_Int HYPRE_SStructVectorSetConstantValues(HYPRE_SStructVector vector, HYPRE_Complex
                                                value)
```

Set vector coefficients to a constant value over the grid.

```
HYPRE_Int HYPRE_SStructVectorSetRandomValues(HYPRE_SStructVector vector, HYPRE_Int seed)
```

Set vector coefficients to random values between -1.0 and 1.0 over the grid.

The parameter *seed* controls the generation of random numbers.

```
HYPRE_Int HYPRE_SStructVectorAddToValues(HYPRE_SStructVector vector, HYPRE_Int part,
                                          HYPRE_Int *index, HYPRE_Int var, HYPRE_Complex
                                          *value)
```

Add to vector coefficients index by index.

NOTE: For better efficiency, use *HYPRE\_SStructVectorAddToBoxValues* to set coefficients a box at a time.

NOTE: Users are required to set values on all processes that own the associated variables. This means that some data will be multiply defined.

```
HYPRE_Int HYPRE_SStructVectorAddFEMValues(HYPRE_SStructVector vector, HYPRE_Int part,
                                           HYPRE_Int *index, HYPRE_Complex *values)
```

Add finite element vector coefficients index by index.

The layout of the data in *values* is determined by the routine *HYPRE\_SStructGridSetFEMOrdering*.

NOTE: For better efficiency, use *HYPRE\_SStructVectorAddFEMBoxValues* to set coefficients a box at a time.

```
HYPRE_Int HYPRE_SStructVectorGetValues(HYPRE_SStructVector vector, HYPRE_Int part,
                                       HYPRE_Int *index, HYPRE_Int var, HYPRE_Complex
                                       *value)
```

Get vector coefficients index by index.

Users must first call the routine *HYPRE\_SStructVectorGather* to ensure that data owned by multiple processes is correct.

NOTE: For better efficiency, use *HYPRE\_SStructVectorGetBoxValues* to get coefficients a box at a time.

NOTE: Users may only get values on processes that own the associated variables.

HYPRE\_Int **HYPRE\_SStructVectorGetFEMValues**(*HYPRE\_SStructVector* vector, HYPRE\_Int part, HYPRE\_Int \*index, HYPRE\_Complex \*values)

Get finite element vector coefficients index by index.

The layout of the data in *values* is determined by the routine *HYPRE\_SStructGridSetFEMOrdering*. Users must first call the routine *HYPRE\_SStructVectorGather* to ensure that data owned by multiple processes is correct.

HYPRE\_Int **HYPRE\_SStructVectorSetBoxValues**(*HYPRE\_SStructVector* vector, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Complex \*values)

Set vector coefficients a box at a time.

The data in *values* is ordered as follows:

```
m = 0;
for (k = ilower[2]; k <= iupper[2]; k++)
  for (j = ilower[1]; j <= iupper[1]; j++)
    for (i = ilower[0]; i <= iupper[0]; i++)
    {
      values[m] = ...;
      m++;
    }
```

NOTE: Users are required to set values on all processes that own the associated variables. This means that some data will be multiply defined.

HYPRE\_Int **HYPRE\_SStructVectorAddToBoxValues**(*HYPRE\_SStructVector* vector, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Complex \*values)

Add to vector coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_SStructVectorSetBoxValues*.

NOTE: Users are required to set values on all processes that own the associated variables. This means that some data will be multiply defined.

HYPRE\_Int **HYPRE\_SStructVectorSetBoxValues2**(*HYPRE\_SStructVector* vector, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Int \*vilower, HYPRE\_Int \*viupper, HYPRE\_Complex \*values)

Set vector coefficients a box at a time.

The *values* array is logically box shaped with value-box extents *vilower* and *viupper* that must contain the set-box extents *ilower* and *iupper*. The data in the *values* array is ordered as in *HYPRE\_SStructVectorSetBoxValues*, but based on the value-box extents.

HYPRE\_Int **HYPRE\_SStructVectorAddToBoxValues2**(*HYPRE\_SStructVector* vector, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Int \*vilower, HYPRE\_Int \*viupper, HYPRE\_Complex \*values)

Add to vector coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_SStructVectorSetBoxValues2*.

HYPRE\_Int **HYPRE\_SStructVectorAddFEMBoxValues**(*HYPRE\_SStructVector* vector, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Complex \*values)

Add finite element vector coefficients a box at a time.

The data in *values* is organized as an array of element vectors ordered as in *HYPRE\_SStructVectorSetBoxValues*. The layout of the data entries of each element vector is determined by the routine *HYPRE\_SStructGridSetFEMOrdering*.

HYPRE\_Int **HYPRE\_SStructVectorAssemble**(*HYPRE\_SStructVector* vector)

Finalize the construction of the vector before using.

HYPRE\_Int **HYPRE\_SStructVectorGetBoxValues**(*HYPRE\_SStructVector* vector, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Complex \*values)

Get vector coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_SStructVectorSetBoxValues*. Users must first call the routine *HYPRE\_SStructVectorGather* to ensure that data owned by multiple processes is correct.

NOTE: Users may only get values on processes that own the associated variables.

HYPRE\_Int **HYPRE\_SStructVectorGetBoxValues2**(*HYPRE\_SStructVector* vector, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Int var, HYPRE\_Int \*vilower, HYPRE\_Int \*viupper, HYPRE\_Complex \*values)

Get vector coefficients a box at a time.

The data in *values* is ordered as in *HYPRE\_SStructVectorSetBoxValues2*.

HYPRE\_Int **HYPRE\_SStructVectorGetFEMBoxValues**(*HYPRE\_SStructVector* vector, HYPRE\_Int part, HYPRE\_Int \*ilower, HYPRE\_Int \*iupper, HYPRE\_Complex \*values)

Does this even make sense to implement?

HYPRE\_Int **HYPRE\_SStructVectorGather**(*HYPRE\_SStructVector* vector)

Gather vector data so that efficient GetValues can be done.

This routine must be called prior to calling GetValues to ensure that correct and consistent values are returned, especially for non cell-centered data that is shared between more than one processor.

HYPRE\_Int **HYPRE\_SStructVectorSetObjectType**(*HYPRE\_SStructVector* vector, HYPRE\_Int type)

Set the storage type of the vector object to be constructed.

Currently, *type* can be either HYPRE\_SSTRUCT (the default), HYPRE\_STRUCT, or HYPRE\_PARCSR.

 See also

*HYPRE\_SStructVectorGetObject*

HYPRE\_Int **HYPRE\_SStructVectorGetObject**(*HYPRE\_SStructVector* vector, void \*\*object)

Get a reference to the constructed vector object.

 See also

*HYPRE\_SStructVectorSetObjectType*

`HYPRE_Int HYPRE_SStructVectorPrint`(const char \*filename, *HYPRE\_SStructVector* vector, `HYPRE_Int` all)

Print the vector to file.

This is mainly for debugging purposes.

`HYPRE_Int HYPRE_SStructVectorRead`(`MPI_Comm` comm, const char \*filename, *HYPRE\_SStructVector* \*vector\_ptr)

Read the vector from file.

This is mainly for debugging purposes.

`HYPRE_Int HYPRE_SStructVectorPrintGLVis`(*HYPRE\_SStructVector* vector, const char \*fileprefix)

### SStruct Other

`HYPRE_Int HYPRE_SStructGetAMRObjects`(*HYPRE\_SStructMatrix* matrix, *HYPRE\_SStructVector* rhs, void \*\*matrix\_object, void \*\*rhs\_object)

AMRNEW.

Get a reference to the constructed matrix and right-hand-side (rhs) objects for an AMR system. This routine is similar to the routines *HYPRE\_SStructMatrixGetObject* and *HYPRE\_SStructVectorGetObject*, but ensures that trivial equations such as Dirichlet conditions are also satisfied exactly in the AMR system.

### Basic Matrix/vector routines

`HYPRE_Int HYPRE_SStructVectorCopy`(*HYPRE\_SStructVector* x, *HYPRE\_SStructVector* y)

Copy vector x into y ( $y \leftarrow x$ ).

`HYPRE_Int HYPRE_SStructVectorScale`(`HYPRE_Complex` alpha, *HYPRE\_SStructVector* y)

Scale a vector by  $\alpha$  ( $y \leftarrow \alpha y$ ).

`HYPRE_Int HYPRE_SStructVectorAxpv`(`HYPRE_Complex` alpha, *HYPRE\_SStructVector* x, *HYPRE\_SStructVector* y)

Compute  $y = y + \alpha x$ .

`HYPRE_Int HYPRE_SStructVectorInnerProd`(*HYPRE\_SStructVector* x, *HYPRE\_SStructVector* y, `HYPRE_Real` \*result)

Compute *result*, the inner product of vectors x and y.

`HYPRE_Int HYPRE_SStructMatrixMatvec`(`HYPRE_Complex` alpha, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* x, `HYPRE_Complex` beta, *HYPRE\_SStructVector* y)

Matvec operator.

This operation is  $y = \alpha Ax + \beta y$ . Note that you can do a simple matrix-vector multiply by setting  $\alpha = 1$  and  $\beta = 0$ .

`HYPRE_Int HYPRE_SStructMatrixMatmat`(*HYPRE\_SStructMatrix* A, *HYPRE\_SStructMatrix* B, *HYPRE\_SStructMatrix* \*C)

Matrix-matrix multiply.

HYPRE\_Int **HYPRE\_SStructMatrixScale**(*HYPRE\_SStructMatrix* matrix, HYPRE\_Complex scalar)

Scale a matrix by *scalar* ( $A \leftarrow \text{scalar} \cdot A$ ).

## 8.3 IJ System Interface

### group IJ System Interface

A linear-algebraic conceptual interface.

This interface represents a linear-algebraic conceptual view of a linear system. The ‘I’ and ‘J’ in the name are meant to be mnemonic for the traditional matrix notation A(I,J).

### IJ Matrices

typedef struct hyre\_IJMatrix\_struct \***HYPRE\_IJMatrix**

The matrix object.

HYPRE\_Int **HYPRE\_IJMatrixCreate**(MPI\_Comm comm, HYPRE\_BigInt ilower, HYPRE\_BigInt iupper, HYPRE\_BigInt jlower, HYPRE\_BigInt jupper, *HYPRE\_IJMatrix* \*matrix)

Create a matrix object.

Each process owns some unique consecutive range of rows, indicated by the global row indices *ilower* and *iupper*. The row data is required to be such that the value of *ilower* on any process  $p$  be exactly one more than the value of *iupper* on process  $p - 1$ . Note that the first row of the global matrix may start with any integer value. In particular, one may use zero- or one-based indexing.

For square matrices, *jlower* and *jupper* typically should match *ilower* and *iupper*, respectively. For rectangular matrices, *jlower* and *jupper* should define a partitioning of the columns. This partitioning must be used for any vector  $v$  that will be used in matrix-vector products with the rectangular matrix. The matrix data structure may use *jlower* and *jupper* to store the diagonal blocks (rectangular in general) of the matrix separately from the rest of the matrix.

Collective.

HYPRE\_Int **HYPRE\_IJMatrixDestroy**(*HYPRE\_IJMatrix* matrix)

Destroy a matrix object.

An object should be explicitly destroyed using this destructor when the user’s code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

HYPRE\_Int **HYPRE\_IJMatrixInitialize**(*HYPRE\_IJMatrix* matrix)

Prepare a matrix object for setting coefficient values.

This routine will also re-initialize an already assembled matrix, allowing users to modify coefficient values.

HYPRE\_Int **HYPRE\_IJMatrixInitialize\_v2**(*HYPRE\_IJMatrix* matrix, *HYPRE\_MemoryLocation* memory\_location)

Prepare a matrix object for setting coefficient values.

This routine will also re-initialize an already assembled matrix, allowing users to modify coefficient values. This routine also specifies the memory location, i.e. host or device.

`HYPRE_Int HYPRE_IJMatrixSetValues`(*HYPRE\_IJMatrix* matrix, `HYPRE_Int` nrows, `HYPRE_Int` \*ncols, `const HYPRE_BigInt` \*rows, `const HYPRE_BigInt` \*cols, `const HYPRE_Complex` \*values)

Sets values for *nrows* rows or partial rows of the matrix.

The arrays *ncols* and *rows* are of dimension *nrows* and contain the number of columns in each row and the row indices, respectively. The array *cols* contains the column indices for each of the *rows*, and is ordered by rows. The data in the *values* array corresponds directly to the column entries in *cols*. Erases any previous values at the specified locations and replaces them with new ones, or, if there was no value there before, inserts a new one if set locally. Note that it is not possible to set values on other processors. If one tries to set a value from proc *i* on proc *j*, proc *i* will erase all previous occurrences of this value in its stack (including values generated with `AddToValues`), and treat it like a zero value. The actual value needs to be set on proc *j*.

Note that a threaded version (threaded over the number of rows) will be called if `HYPRE_IJMatrixSetOMPFlag` is set to a value  $\neq 0$ . This requires that  $\text{rows}[i] \neq \text{rows}[j]$  for  $i \neq j$  and is only efficient if a large number of rows is set in one call to `HYPRE_IJMatrixSetValues`.

Not collective.

`HYPRE_Int HYPRE_IJMatrixSetConstantValues`(*HYPRE\_IJMatrix* matrix, `HYPRE_Complex` value)

Sets all matrix coefficients of an already assembled matrix to *value*.

`HYPRE_Int HYPRE_IJMatrixAddToValues`(*HYPRE\_IJMatrix* matrix, `HYPRE_Int` nrows, `HYPRE_Int` \*ncols, `const HYPRE_BigInt` \*rows, `const HYPRE_BigInt` \*cols, `const HYPRE_Complex` \*values)

Adds to values for *nrows* rows or partial rows of the matrix.

Usage details are analogous to `HYPRE_IJMatrixSetValues`. Adds to any previous values at the specified locations, or, if there was no value there before, inserts a new one. `AddToValues` can be used to add to values on other processors.

Note that a threaded version (threaded over the number of rows) will be called if `HYPRE_IJMatrixSetOMPFlag` is set to a value  $\neq 0$ . This requires that  $\text{rows}[i] \neq \text{rows}[j]$  for  $i \neq j$  and is only efficient if a large number of rows is added in one call to `HYPRE_IJMatrixAddToValues`.

Not collective.

`HYPRE_Int HYPRE_IJMatrixSetValues2`(*HYPRE\_IJMatrix* matrix, `HYPRE_Int` nrows, `HYPRE_Int` \*ncols, `const HYPRE_BigInt` \*rows, `const HYPRE_Int` \*row\_indexes, `const HYPRE_BigInt` \*cols, `const HYPRE_Complex` \*values)

Sets values for *nrows* rows or partial rows of the matrix.

Same as `IJMatrixSetValues`, but with an additional *row\_indexes* array that provides indexes into the *cols* and *values* arrays. Because of this, there can be gaps between the row data in these latter two arrays.

`HYPRE_Int HYPRE_IJMatrixAddToValues2`(*HYPRE\_IJMatrix* matrix, `HYPRE_Int` nrows, `HYPRE_Int` \*ncols, `const HYPRE_BigInt` \*rows, `const HYPRE_Int` \*row\_indexes, `const HYPRE_BigInt` \*cols, `const HYPRE_Complex` \*values)

Adds to values for *nrows* rows or partial rows of the matrix.

Same as `IJMatrixAddToValues`, but with an additional *row\_indexes* array that provides indexes into the *cols* and *values* arrays. Because of this, there can be gaps between the row data in these latter two arrays.

`HYPRE_Int HYPRE_IJMatrixAssemble`(*HYPRE\_IJMatrix* matrix)

Finalize the construction of the matrix before using.

HYPRE\_Int **HYPRE\_IJMatrixGetRowCounts**(*HYPRE\_IJMatrix* matrix, HYPRE\_Int nrows, HYPRE\_BigInt \*rows, HYPRE\_Int \*ncols)

Gets number of nonzeros elements for *nrows* rows specified in *rows* and returns them in *ncols*, which needs to be allocated by the user.

HYPRE\_Int **HYPRE\_IJMatrixGetValues**(*HYPRE\_IJMatrix* matrix, HYPRE\_Int nrows, HYPRE\_Int \*ncols, HYPRE\_BigInt \*rows, HYPRE\_BigInt \*cols, HYPRE\_Complex \*values)

Gets values for *nrows* rows or partial rows of the matrix.

Usage details are mostly analogous to *HYPRE\_IJMatrixSetValues*. Note that if *nrows* is negative, the routine will return the column\_indices and matrix coefficients of the (-*nrows*) rows contained in *rows*.

HYPRE\_Int **HYPRE\_IJMatrixGetValues2**(*HYPRE\_IJMatrix* matrix, HYPRE\_Int nrows, HYPRE\_Int \*ncols, HYPRE\_BigInt \*rows, HYPRE\_Int \*row\_indexes, HYPRE\_BigInt \*cols, HYPRE\_Complex \*values)

Gets values for *nrows* rows or partial rows of the matrix.

Same as *IJMatrixGetValues*, but with an additional *row\_indexes* array that provides indexes into the *cols* and *values* arrays. Because of this, there can be gaps between the row data in these latter two arrays.

HYPRE\_Int **HYPRE\_IJMatrixGetValuesAndZeroOut**(*HYPRE\_IJMatrix* matrix, HYPRE\_Int nrows, HYPRE\_Int \*ncols, HYPRE\_BigInt \*rows, HYPRE\_Int \*row\_indexes, HYPRE\_BigInt \*cols, HYPRE\_Complex \*values)

Gets values for *nrows* rows or partial rows of the matrix and zeros out those entries in the matrix.

Same as *IJMatrixGetValues2*, but zeros out the entries after getting them.

HYPRE\_Int **HYPRE\_IJMatrixSetObjectType**(*HYPRE\_IJMatrix* matrix, HYPRE\_Int type)

Set the storage type of the matrix object to be constructed.

Currently, *type* can only be *HYPRE\_PARCSR*.

Not collective, but must be the same on all processes.

 See also

*HYPRE\_IJMatrixGetObject*

HYPRE\_Int **HYPRE\_IJMatrixGetObjectType**(*HYPRE\_IJMatrix* matrix, HYPRE\_Int \*type)

Get the storage type of the constructed matrix object.

HYPRE\_Int **HYPRE\_IJMatrixGetLocalRange**(*HYPRE\_IJMatrix* matrix, HYPRE\_BigInt \*ilower, HYPRE\_BigInt \*iupper, HYPRE\_BigInt \*jlower, HYPRE\_BigInt \*jupper)

Gets range of rows owned by this processor and range of column partitioning for this processor.

HYPRE\_Int **HYPRE\_IJMatrixGetGlobalInfo**(*HYPRE\_IJMatrix* matrix, HYPRE\_BigInt \*global\_num\_rows, HYPRE\_BigInt \*global\_num\_cols, HYPRE\_BigInt \*global\_num\_nonzeros)

Gets global information about the matrix, including the total number of rows, columns, and nonzero elements across all processes.

Collective (must be called by all processes).

#### Parameters

- **matrix** – The IJMatrix object to query.
- **global\_num\_rows** – Pointer to store the total number of rows in the matrix.
- **global\_num\_cols** – Pointer to store the total number of columns in the matrix.
- **global\_num\_nonzeros** – Pointer to store the total number of nonzero elements in the matrix.

#### Returns

HYPRE\_Int Error code.

HYPRE\_Int **HYPRE\_IJMatrixGetObject**(*HYPRE\_IJMatrix* matrix, void \*\*object)

Get a reference to the constructed matrix object.

#### ➔ See also

*HYPRE\_IJMatrix.SetObjectType*

HYPRE\_Int **HYPRE\_IJMatrixSetRowSizes**(*HYPRE\_IJMatrix* matrix, const HYPRE\_Int \*sizes)

(Optional) Set the max number of nonzeros to expect in each row.

The array *sizes* contains estimated sizes for each row on this process. This call can significantly improve the efficiency of matrix construction, and should always be utilized if possible.

Not collective.

HYPRE\_Int **HYPRE\_IJMatrixSetDiagOffdSizes**(*HYPRE\_IJMatrix* matrix, const HYPRE\_Int \*diag\_sizes, const HYPRE\_Int \*offdiag\_sizes)

(Optional) Sets the exact number of nonzeros in each row of the diagonal and off-diagonal blocks.

The diagonal block is the submatrix whose column numbers correspond to rows owned by this process, and the off-diagonal block is everything else. The arrays *diag\_sizes* and *offdiag\_sizes* contain estimated sizes for each row of the diagonal and off-diagonal blocks, respectively. This routine can significantly improve the efficiency of matrix construction, and should always be utilized if possible.

Not collective.

HYPRE\_Int **HYPRE\_IJMatrixSetMaxOffProcElmts**(*HYPRE\_IJMatrix* matrix, HYPRE\_Int max\_off\_proc\_elmts)

(Optional) Sets the maximum number of elements that are expected to be set (or added) on other processors from this processor. This routine can significantly improve the efficiency of matrix construction, and should always be utilized if possible.

Not collective.

HYPRE\_Int **HYPRE\_IJMatrixSetInitAllocation**(*HYPRE\_IJMatrix* matrix, HYPRE\_Int factor)

(Optional, GPU only) Sets the initial memory allocation for matrix assemble, which factor \* local number of rows. Not collective.

HYPRE\_Int **HYPRE\_IJMatrixSetEarlyAssemble**(*HYPRE\_IJMatrix* matrix, HYPRE\_Int early\_assemble)

(Optional, GPU only) Sets if matrix assemble routine does reductions during the accumulation of the entries before calling *HYPRE\_IJMatrixAssemble*.

This early assemble feature may save the peak memory usage but requires extra work. Not collective.

HYPRE\_Int **HYPRE\_IJMatrixSetGrowFactor**(*HYPRE\_IJMatrix* matrix, HYPRE\_Real factor)  
 (Optional, GPU only) Sets the grow factor of memory in matrix assemble when running out of memory.

Not collective.

HYPRE\_Int **HYPRE\_IJMatrixSetPrintLevel**(*HYPRE\_IJMatrix* matrix, HYPRE\_Int print\_level)  
 (Optional) Sets the print level, if the user wants to print error messages.

The default is 0, i.e. no error messages are printed.

HYPRE\_Int **HYPRE\_IJMatrixSetOMPflag**(*HYPRE\_IJMatrix* matrix, HYPRE\_Int omp\_flag)  
 (Optional) if set, will use a threaded version of HYPRE\_IJMatrixSetValues and HYPRE\_IJMatrixAddToValues.

This is only useful if a large number of rows is set or added to at once.

NOTE that the values in the rows array of HYPRE\_IJMatrixSetValues or HYPRE\_IJMatrixAddToValues must be different from each other !!!

This option is VERY inefficient if only a small number of rows is set or added at once and/or if reallocation of storage is required and/or if values are added to off processor values.

HYPRE\_Int **HYPRE\_IJMatrixRead**(const char \*filename, MPI\_Comm comm, HYPRE\_Int type,  
*HYPRE\_IJMatrix* \*matrix)

Read the matrix from file.

This is mainly for debugging purposes.

HYPRE\_Int **HYPRE\_IJMatrixReadMM**(const char \*filename, MPI\_Comm comm, HYPRE\_Int type,  
*HYPRE\_IJMatrix* \*matrix)

Read the matrix from MM file.

This is mainly for debugging purposes.

HYPRE\_Int **HYPRE\_IJMatrixPrint**(*HYPRE\_IJMatrix* matrix, const char \*filename)

Print the matrix to file.

This is mainly for debugging purposes.

HYPRE\_Int **HYPRE\_IJMatrixTranspose**(*HYPRE\_IJMatrix* matrix\_A, *HYPRE\_IJMatrix* \*matrix\_AT)

Transpose an IJMatrix.

HYPRE\_Int **HYPRE\_IJMatrixNorm**(*HYPRE\_IJMatrix* matrix, HYPRE\_Real \*norm)

Computes the infinity norm of an IJMatrix.

HYPRE\_Int **HYPRE\_IJMatrixAdd**(HYPRE\_Complex alpha, *HYPRE\_IJMatrix* matrix\_A, HYPRE\_Complex  
 beta, *HYPRE\_IJMatrix* matrix\_B, *HYPRE\_IJMatrix* \*matrix\_C)

Performs  $C = \alpha * A + \beta * B$ .

HYPRE\_Int **HYPRE\_IJMatrixPrintBinary**(*HYPRE\_IJMatrix* matrix, const char \*filename)

Print the matrix to file in binary format.

This is mainly for debugging purposes.

HYPRE\_Int **HYPRE\_IJMatrixReadBinary**(const char \*filename, MPI\_Comm comm, HYPRE\_Int type,  
*HYPRE\_IJMatrix* \*matrix\_ptr)

Read the matrix from file stored in binary format.

This is mainly for debugging purposes.

HYPRE\_Int **HYPRE\_IJMatrixMigrate**(*HYPRE\_IJMatrix* matrix, *HYPRE\_MemoryLocation* memory\_location)

Migrate the matrix to a given memory location.

HYPRE\_Int **HYPRE\_IJMatrixPartialClone**(*HYPRE\_IJMatrix* matrix\_in, *HYPRE\_IJMatrix* \*matrix\_out)

## IJ Vectors

typedef struct hypre\_IJVector\_struct \***HYPRE\_IJVector**

The vector object.

HYPRE\_Int **HYPRE\_IJVectorCreate**(MPI\_Comm comm, HYPRE\_BigInt jlower, HYPRE\_BigInt jupper, *HYPRE\_IJVector* \*vector)

Create a vector object.

Each process owns some unique consecutive range of vector unknowns, indicated by the global indices *jlower* and *jupper*. The data is required to be such that the value of *jlower* on any process *p* be exactly one more than the value of *jupper* on process *p* - 1. Note that the first index of the global vector may start with any integer value. In particular, one may use zero- or one-based indexing.

Collective.

HYPRE\_Int **HYPRE\_IJVectorDestroy**(*HYPRE\_IJVector* vector)

Destroy a vector object.

An object should be explicitly destroyed using this destructor when the user's code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

HYPRE\_Int **HYPRE\_IJVectorInitializeShell**(*HYPRE\_IJVector* vector)

This function should be called before `HYPRE_IJVectorSetData` if users intend to reuse an existing data pointer, thereby avoiding unnecessary memory copies.

It configures the vector to accept external data without allocating new storage.

HYPRE\_Int **HYPRE\_IJVectorSetData**(*HYPRE\_IJVector* vector, HYPRE\_Complex \*data)

This function sets the internal data pointer of the vector to an external array, allowing direct control over the vector's data storage without transferring ownership.

Users are responsible for managing the memory of the data array, which must remain valid for the vector's lifetime.

Users should call `HYPRE_IJVectorInitializeShell` before this function to prepare the vector for external data. The memory location of the data array is expected to be on the host when hypr is configured without GPU support. If hypr is configured with GPU support, it is assumed that data resides in device memory.

HYPRE\_Int **HYPRE\_IJVectorSetTags**(*HYPRE\_IJVector* vector, HYPRE\_Int owns\_tags, HYPRE\_Int num\_tags, HYPRE\_Int \*tags)

(Optional) Set an array of tags for the vector.

Not collective.

### Parameters

- **vector** – The vector object.

- **owns\_tags** – Whether the vector owns the tags. If true, vector will allocate and copy tags. If false, vector will just point to the input tags array.
- **num\_tags** – The number of tags.
- **tags** – The tags array. Must reside in the same memory location as the vector data (e.g., if vector is on GPU, tags must also be on GPU).

HYPRE\_Int **HYPRE\_IJVectorInitialize**(*HYPRE\_IJVector* vector)

Prepare a vector object for setting coefficient values.

This routine will also re-initialize an already assembled vector, allowing users to modify coefficient values.

HYPRE\_Int **HYPRE\_IJVectorInitialize\_v2**(*HYPRE\_IJVector* vector, *HYPRE\_MemoryLocation* memory\_location)

Prepare a vector object for setting coefficient values.

This routine will also re-initialize an already assembled vector, allowing users to modify coefficient values. This routine also specifies the memory location, i.e. host or device.

HYPRE\_Int **HYPRE\_IJVectorSetMaxOffProcElmts**(*HYPRE\_IJVector* vector, HYPRE\_Int max\_off\_proc\_elmts)

(Optional) Sets the maximum number of elements that are expected to be set (or added) on other processors from this processor. This routine can significantly improve the efficiency of matrix construction, and should always be utilized if possible.

Not collective.

HYPRE\_Int **HYPRE\_IJVectorSetNumComponents**(*HYPRE\_IJVector* vector, HYPRE\_Int num\_components)

(Optional) Sets the number of components (vectors) of a multivector.

A vector is assumed to have a single component when this function is not called. This function must be called prior to `HYPRE_IJVectorInitialize`.

HYPRE\_Int **HYPRE\_IJVectorSetComponent**(*HYPRE\_IJVector* vector, HYPRE\_Int component)

(Optional) Sets the component identifier of a vector with multiple components (multivector).

This can be used for Set/AddTo/Get purposes.

HYPRE\_Int **HYPRE\_IJVectorSetValues**(*HYPRE\_IJVector* vector, HYPRE\_Int nvalues, const HYPRE\_BigInt \*indices, const HYPRE\_Complex \*values)

Sets values in vector.

The arrays *values* and *indices* are of dimension *nvalues* and contain the vector values to be set and the corresponding global vector indices, respectively. Erases any previous values at the specified locations and replaces them with new ones. Note that it is not possible to set values on other processors. If one tries to set a value from proc *i* on proc *j*, proc *i* will erase all previous occurrences of this value in its stack (including values generated with `AddToValues`), and treat it like a zero value. The actual value needs to be set on proc *j*.

Not collective.

HYPRE\_Int **HYPRE\_IJVectorSetConstantValues**(*HYPRE\_IJVector* vector, HYPRE\_Complex value)

Sets all vector coefficients to *value*.

HYPRE\_Int **HYPRE\_IJVectorAddToValues**(*HYPRE\_IJVector* vector, HYPRE\_Int nvalues, const HYPRE\_BigInt \*indices, const HYPRE\_Complex \*values)

Adds to values in vector.

Usage details are analogous to *HYPRE\_IJVectorSetValues*. Adds to any previous values at the specified locations, or, if there was no value there before, inserts a new one. *AddToValues* can be used to add to values on other processors.

Not collective.

`HYPRE_Int HYPRE_IJVectorAssemble(HYPRE_IJVector vector)`

Finalize the construction of the vector before using.

`HYPRE_Int HYPRE_IJVectorUpdateValues(HYPRE_IJVector vector, HYPRE_Int nvalues, const HYPRE_BigInt *indices, const HYPRE_Complex *values, HYPRE_Int action)`

Update vectors by setting (action 1) or adding to (action 0) values in ‘vector’.

Note that this function cannot update values owned by other processes and does not allow repeated index values in ‘indices’.

Not collective.

`HYPRE_Int HYPRE_IJVectorGetValues(HYPRE_IJVector vector, HYPRE_Int nvalues, const HYPRE_BigInt *indices, HYPRE_Complex *values)`

Gets values in vector.

Usage details are analogous to *HYPRE\_IJVectorSetValues*.

Not collective.

`HYPRE_Int HYPRE_IJVectorSetObjectType(HYPRE_IJVector vector, HYPRE_Int type)`

Set the storage type of the vector object to be constructed.

Currently, *type* can only be `HYPRE_PARCSR`.

Not collective, but must be the same on all processes.

 See also

*HYPRE\_IJVectorGetObject*

`HYPRE_Int HYPRE_IJVectorGetObjectType(HYPRE_IJVector vector, HYPRE_Int *type)`

Get the storage type of the constructed vector object.

`HYPRE_Int HYPRE_IJVectorGetLocalRange(HYPRE_IJVector vector, HYPRE_BigInt *jlower, HYPRE_BigInt *jupper)`

Returns range of the part of the vector owned by this processor.

`HYPRE_Int HYPRE_IJVectorGetObject(HYPRE_IJVector vector, void **object)`

Get a reference to the constructed vector object.

 See also

*HYPRE\_IJVectorSetObjectType*

HYPRE\_Int **HYPRE\_IJVectorSetPrintLevel**(*HYPRE\_IJVector* vector, HYPRE\_Int print\_level)  
 (Optional) Sets the print level, if the user wants to print error messages.

The default is 0, i.e. no error messages are printed.

HYPRE\_Int **HYPRE\_IJVectorRead**(const char \*filename, MPI\_Comm comm, HYPRE\_Int type,  
*HYPRE\_IJVector* \*vector)

Read the vector from file.

This is mainly for debugging purposes.

HYPRE\_Int **HYPRE\_IJVectorReadBinary**(const char \*filename, MPI\_Comm comm, HYPRE\_Int type,  
*HYPRE\_IJVector* \*vector)

Read the vector from binary file.

This is mainly for debugging purposes.

HYPRE\_Int **HYPRE\_IJVectorPrint**(*HYPRE\_IJVector* vector, const char \*filename)

Print the vector to file.

This is mainly for debugging purposes.

HYPRE\_Int **HYPRE\_IJVectorPrintBinary**(*HYPRE\_IJVector* vector, const char \*filename)

Print the vector to binary file.

This is mainly for debugging purposes.

HYPRE\_Int **HYPRE\_IJVectorInnerProd**(*HYPRE\_IJVector* x, *HYPRE\_IJVector* y, HYPRE\_Real \*prod)

Computes the inner product between two vectors.

HYPRE\_Int **HYPRE\_IJVectorMigrate**(*HYPRE\_IJVector* vector, *HYPRE\_MemoryLocation*  
 memory\_location)

Migrate the vector to a given memory location.

## 8.4 ParCSR Object Interface

### *group* ParCSR Object Interface

Interface for ParCSR matrices and vectors.

This is an interface for matrices and vectors with object type HYPRE\_PARCSR.

### ParCSR Matrices

typedef struct hypre\_ParCSRMatrix\_struct **\*HYPRE\_ParCSRMatrix**

The matrix object.

ParCSR matrices use a parallel compressed-sparse-row (CSR) storage format that consists of a diagonal CSR matrix for intra-processor couplings and an off-diagonal CSR matrix for inter-processor couplings.

HYPRE\_Int **HYPRE\_ParCSRMatrixCreate**(MPI\_Comm comm, HYPRE\_BigInt global\_num\_rows,  
 HYPRE\_BigInt global\_num\_cols, HYPRE\_BigInt \*row\_starts,  
 HYPRE\_BigInt \*col\_starts, HYPRE\_Int num\_cols\_offd,  
 HYPRE\_Int num\_nonzeros\_diag, HYPRE\_Int  
 num\_nonzeros\_offd, *HYPRE\_ParCSRMatrix* \*matrix)

Create a matrix object.

More info here about arguments...

HYPRE\_Int **HYPRE\_ParCSRMatrixDestroy**(*HYPRE\_ParCSRMatrix* matrix)

Destroy a matrix object.

HYPRE\_Int **HYPRE\_ParCSRMatrixInitialize**(*HYPRE\_ParCSRMatrix* matrix)

Prepare a matrix object for setting coefficient values.

HYPRE\_Int **HYPRE\_ParCSRMatrixRead**(MPI\_Comm comm, const char \*file\_name, *HYPRE\_ParCSRMatrix* \*matrix)

Read a matrix from file.

HYPRE\_Int **HYPRE\_ParCSRMatrixPrint**(*HYPRE\_ParCSRMatrix* matrix, const char \*file\_name)

Print a matrix to file.

## ParCSR Vectors

typedef struct hypr\_ParVector\_struct **HYPRE\_ParVector**

The vector object.

A Par vector is an array storage format compatible with ParCSR matrices.

HYPRE\_Int **HYPRE\_ParVectorCreate**(MPI\_Comm comm, HYPRE\_BigInt global\_size, HYPRE\_BigInt \*partitioning, *HYPRE\_ParVector* \*vector)

Create a vector object.

HYPRE\_Int **HYPRE\_ParVectorDestroy**(*HYPRE\_ParVector* vector)

Destroy a vector object.

HYPRE\_Int **HYPRE\_ParVectorInitialize**(*HYPRE\_ParVector* vector)

Prepare a vector object for setting coefficient values.

HYPRE\_Int **HYPRE\_ParVectorRead**(MPI\_Comm comm, const char \*file\_name, *HYPRE\_ParVector* \*vector)

Read a vector from file.

HYPRE\_Int **HYPRE\_ParVectorPrint**(*HYPRE\_ParVector* vector, const char \*file\_name)

Print a vector to file.

## Basic Matrix/vector routines

HYPRE\_Int **HYPRE\_ParVectorCopy**(*HYPRE\_ParVector* x, *HYPRE\_ParVector* y)

Copy vector x into y ( $y \leftarrow x$ ).

HYPRE\_Int **HYPRE\_ParVectorScale**(HYPRE\_Complex value, *HYPRE\_ParVector* x)

Scale a vector by  $\alpha$  ( $y \leftarrow \alpha y$ ).

HYPRE\_Int **HYPRE\_ParVectorAxy**(HYPRE\_Complex alpha, *HYPRE\_ParVector* x, *HYPRE\_ParVector* y)

Compute  $y = y + \alpha x$ .

HYPRE\_Int **HYPRE\_ParVectorInnerProd**(*HYPRE\_ParVector* x, *HYPRE\_ParVector* y, HYPRE\_Real \*result)

Compute *result*, the inner product of vectors x and y.

HYPRE\_Int **HYPRE\_ParCSRMatrixMatvec**(HYPRE\_Complex alpha, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* x, HYPRE\_Complex beta, *HYPRE\_ParVector* y)

Compute a matrix-vector product  $y = \alpha Ax + \beta y$ .

HYPRE\_Int **HYPRE\_ParCSRMatrixMatvecT**(HYPRE\_Complex alpha, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* x, HYPRE\_Complex beta, *HYPRE\_ParVector* y)

Compute a transpose matrix-vector product  $y = \alpha A^T x + \beta y$ .

HYPRE\_Int **HYPRE\_ParCSRMatrixMatmat**(*HYPRE\_ParCSRMatrix* A, *HYPRE\_ParCSRMatrix* B, *HYPRE\_ParCSRMatrix* \*C)

Matrix-matrix multiply.

## 8.5 Struct Solvers

### group Struct Solvers

Linear solvers for structured grids.

These solvers use matrix/vector storage schemes that are tailored to structured grid problems.

### Struct Solvers

typedef struct hypre\_StructSolver\_struct \***HYPRE\_StructSolver**

The solver object.

typedef HYPRE\_Int (\***HYPRE\_PtrToStructSolverFcn**)(*HYPRE\_StructSolver*, *HYPRE\_StructMatrix*, *HYPRE\_StructVector*, *HYPRE\_StructVector*)

### Struct Jacobi Solver

HYPRE\_Int **HYPRE\_StructJacobiCreate**(MPI\_Comm comm, *HYPRE\_StructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_StructJacobiDestroy**(*HYPRE\_StructSolver* solver)

Destroy a solver object.

An object should be explicitly destroyed using this destructor when the user's code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

HYPRE\_Int **HYPRE\_StructJacobiSetup**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, *HYPRE\_StructVector* b, *HYPRE\_StructVector* x)

Prepare to solve the system.

The coefficient data in  $b$  and  $x$  is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_StructJacobiSolve**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, *HYPRE\_StructVector* b, *HYPRE\_StructVector* x)

Solve the system.

HYPRE\_Int **HYPRE\_StructJacobiSetTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance.

HYPRE\_Int **HYPRE\_StructJacobiGetTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real \*tol)

HYPRE\_Int **HYPRE\_StructJacobiSetMaxIter**(*HYPRE\_StructSolver* solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_StructJacobiGetMaxIter**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*max\_iter)

HYPRE\_Int **HYPRE\_StructJacobiSetZeroGuess**(*HYPRE\_StructSolver* solver)

(Optional) Use a zero initial guess.

This allows the solver to cut corners in the case where a zero initial guess is needed (e.g., for preconditioning) to reduce computational cost.

HYPRE\_Int **HYPRE\_StructJacobiGetZeroGuess**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*zeroguess)

HYPRE\_Int **HYPRE\_StructJacobiSetNonZeroGuess**(*HYPRE\_StructSolver* solver)

(Optional) Use a nonzero initial guess.

This is the default behavior, but this routine allows the user to switch back after using *SetZeroGuess*.

HYPRE\_Int **HYPRE\_StructJacobiGetNumIterations**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_StructJacobiGetFinalRelativeResidualNorm**(*HYPRE\_StructSolver* solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

### Struct PFMG Solver

PFMG is a semicoarsening multigrid solver that uses pointwise relaxation.

For periodic problems, users should try to set the grid size in periodic dimensions to be as close to a power-of-two as possible. That is, if the grid size in a periodic dimension is given by  $N = 2^m * M$  where  $M$  is not a power-of-two, then  $M$  should be as small as possible. Large values of  $M$  will generally result in slower convergence rates.

HYPRE\_Int **HYPRE\_StructPFMGCreate**(MPI\_Comm comm, *HYPRE\_StructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_StructPFMGDestroy**(*HYPRE\_StructSolver* solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_StructPFMGSetup**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, HYPRE\_StructVector b, HYPRE\_StructVector x)

Prepare to solve the system.

The coefficient data in  $b$  and  $x$  is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_StructPFMGsolve**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, HYPRE\_StructVector b, HYPRE\_StructVector x)

Solve the system.

HYPRE\_Int **HYPRE\_StructPFMGSetTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance.

HYPRE\_Int **HYPRE\_StructPFMGGetTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real \*tol)

HYPRE\_Int **HYPRE\_StructPFMGSetMaxIter**(*HYPRE\_StructSolver* solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_StructPFMGGetMaxIter**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*max\_iter)

HYPRE\_Int **HYPRE\_StructPFMGSetMaxLevels**(*HYPRE\_StructSolver* solver, HYPRE\_Int max\_levels)  
 (Optional) Set maximum number of multigrid grid levels.

HYPRE\_Int **HYPRE\_StructPFMGGetMaxLevels**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*max\_levels)

HYPRE\_Int **HYPRE\_StructPFMGSetRelChange**(*HYPRE\_StructSolver* solver, HYPRE\_Int rel\_change)  
 (Optional) Additionally require that the relative difference in successive iterates be small.

HYPRE\_Int **HYPRE\_StructPFMGGetRelChange**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*rel\_change)

HYPRE\_Int **HYPRE\_StructPFMGSetZeroGuess**(*HYPRE\_StructSolver* solver)

(Optional) Use a zero initial guess.

This allows the solver to cut corners in the case where a zero initial guess is needed (e.g., for preconditioning) to reduce computational cost.

HYPRE\_Int **HYPRE\_StructPFMGGetZeroGuess**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*zeroguess)

HYPRE\_Int **HYPRE\_StructPFMGSetNonZeroGuess**(*HYPRE\_StructSolver* solver)

(Optional) Use a nonzero initial guess.

This is the default behavior, but this routine allows the user to switch back after using *SetZeroGuess*.

HYPRE\_Int **HYPRE\_StructPFMGSetRelaxType**(*HYPRE\_StructSolver* solver, HYPRE\_Int relax\_type)

(Optional) Set relaxation type.

Current relaxation methods set by *relax\_type* are:

- 0 : Jacobi
- 1 : Weighted Jacobi (default)
- 2 : Red/Black Gauss-Seidel (symmetric: RB pre-relaxation, BR post-relaxation)
- 3 : Red/Black Gauss-Seidel (nonsymmetric: RB pre- and post-relaxation)

HYPRE\_Int **HYPRE\_StructPFMGGetRelaxType**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*relax\_type)

HYPRE\_Int **HYPRE\_StructPFMGSetJacobiWeight**(*HYPRE\_StructSolver* solver, HYPRE\_Real weight)

HYPRE\_Int **HYPRE\_StructPFMGGetJacobiWeight**(*HYPRE\_StructSolver* solver, HYPRE\_Real \*weight)

HYPRE\_Int **HYPRE\_StructPFMGSetRAPType**(*HYPRE\_StructSolver* solver, HYPRE\_Int rap\_type)

(Optional) Set type of coarse-grid operator to use.

Current operators set by *rap\_type* are:

- 0 : Galerkin (default)
- 1 : non-Galerkin 5-pt or 7-pt stencils

Both operators are constructed algebraically. The non-Galerkin option maintains a 5-pt stencil in 2D and a 7-pt stencil in 3D on all grid levels. The stencil coefficients are computed by averaging techniques.

HYPRE\_Int **HYPRE\_StructPFMGGetRAPType**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*rap\_type)

HYPRE\_Int **HYPRE\_StructPFMGSetMatmultType**(*HYPRE\_StructSolver* solver, HYPRE\_Int matmult\_type)  
 (Optional) Set the kernel type used for computing struct matrix-matrix multiplication.

Current values set by *matmult\_type* are:

- -1 : proxy to the default depending on hypr's build type (CPU or GPU)
- 0 : standard (core) algorithm (default for CPUs)
- 1 : fused algorithm with less, but more computationally intensive BoxLoops (default for GPUs)

HYPRE\_Int **HYPRE\_StructPFMGGetMatmultType**(*HYPRE\_StructSolver* solver, HYPRE\_Int  
 \*matmult\_type)

HYPRE\_Int **HYPRE\_StructPFMGSetNumPreRelax**(*HYPRE\_StructSolver* solver, HYPRE\_Int  
 num\_pre\_relax)

(Optional) Set number of relaxation sweeps before coarse-grid correction.

HYPRE\_Int **HYPRE\_StructPFMGGetNumPreRelax**(*HYPRE\_StructSolver* solver, HYPRE\_Int  
 \*num\_pre\_relax)

HYPRE\_Int **HYPRE\_StructPFMGSetNumPostRelax**(*HYPRE\_StructSolver* solver, HYPRE\_Int  
 num\_post\_relax)

(Optional) Set number of relaxation sweeps after coarse-grid correction.

HYPRE\_Int **HYPRE\_StructPFMGGetNumPostRelax**(*HYPRE\_StructSolver* solver, HYPRE\_Int  
 \*num\_post\_relax)

HYPRE\_Int **HYPRE\_StructPFMGSetSkipRelax**(*HYPRE\_StructSolver* solver, HYPRE\_Int skip\_relax)

(Optional) Skip relaxation on certain grids for isotropic problems.

This can greatly improve efficiency by eliminating unnecessary relaxations when the underlying problem is isotropic.

HYPRE\_Int **HYPRE\_StructPFMGGetSkipRelax**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*skip\_relax)

HYPRE\_Int **HYPRE\_StructPFMGSetDxyz**(*HYPRE\_StructSolver* solver, HYPRE\_Real \*dxyz)

HYPRE\_Int **HYPRE\_StructPFMGSetLogging**(*HYPRE\_StructSolver* solver, HYPRE\_Int logging)

(Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_StructPFMGGetLogging**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*logging)

HYPRE\_Int **HYPRE\_StructPFMGSetPrintLevel**(*HYPRE\_StructSolver* solver, HYPRE\_Int print\_level)

(Optional) Control how much information is printed.

- 0 : no printout (default)
- 1 : print convergence history

HYPRE\_Int **HYPRE\_StructPFMGGetPrintLevel**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*print\_level)

HYPRE\_Int **HYPRE\_StructPFMGGetNumIterations**(*HYPRE\_StructSolver* solver, HYPRE\_Int  
 \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_StructPFMGGetFinalRelativeResidualNorm**(*HYPRE\_StructSolver* solver,  
HYPRE\_Real \*norm)

Return the norm of the final relative residual.

### Struct SMG Solver

SMG is a semicoarsening multigrid solver that uses plane smoothing (in 3D).

The plane smoother calls a 2D SMG algorithm with line smoothing, and the line smoother is cyclic reduction (1D SMG). For periodic problems, the grid size in periodic dimensions currently must be a power-of-two.

HYPRE\_Int **HYPRE\_StructSMGCreate**(MPI\_Comm comm, *HYPRE\_StructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_StructSMGDestroy**(*HYPRE\_StructSolver* solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_StructSMGSetup**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A,  
HYPRE\_StructVector b, HYPRE\_StructVector x)

Prepare to solve the system.

The coefficient data in *b* and *x* is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_StructSMGSolve**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A,  
HYPRE\_StructVector b, HYPRE\_StructVector x)

Solve the system.

HYPRE\_Int **HYPRE\_StructSMGSetMemoryUse**(*HYPRE\_StructSolver* solver, HYPRE\_Int memory\_use)

HYPRE\_Int **HYPRE\_StructSMGGetMemoryUse**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*memory\_use)

HYPRE\_Int **HYPRE\_StructSMGSetTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance.

HYPRE\_Int **HYPRE\_StructSMGGetTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real \*tol)

HYPRE\_Int **HYPRE\_StructSMGSetMaxIter**(*HYPRE\_StructSolver* solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_StructSMGGetMaxIter**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*max\_iter)

HYPRE\_Int **HYPRE\_StructSMGSetRelChange**(*HYPRE\_StructSolver* solver, HYPRE\_Int rel\_change)

(Optional) Additionally require that the relative difference in successive iterates be small.

HYPRE\_Int **HYPRE\_StructSMGGetRelChange**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*rel\_change)

HYPRE\_Int **HYPRE\_StructSMGSetZeroGuess**(*HYPRE\_StructSolver* solver)

(Optional) Use a zero initial guess.

This allows the solver to cut corners in the case where a zero initial guess is needed (e.g., for preconditioning) to reduce computational cost.

HYPRE\_Int **HYPRE\_StructSMGGetZeroGuess**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*zeroguess)

HYPRE\_Int **HYPRE\_StructSMGSetNonZeroGuess**(*HYPRE\_StructSolver* solver)

(Optional) Use a nonzero initial guess.

This is the default behavior, but this routine allows the user to switch back after using *SetZeroGuess*.

HYPRE\_Int **HYPRE\_StructSMGSetNumPreRelax**(*HYPRE\_StructSolver* solver, HYPRE\_Int num\_pre\_relax)  
 (Optional) Set number of relaxation sweeps before coarse-grid correction.

HYPRE\_Int **HYPRE\_StructSMGGetNumPreRelax**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*num\_pre\_relax)

HYPRE\_Int **HYPRE\_StructSMGSetNumPostRelax**(*HYPRE\_StructSolver* solver, HYPRE\_Int num\_post\_relax)  
 (Optional) Set number of relaxation sweeps after coarse-grid correction.

HYPRE\_Int **HYPRE\_StructSMGGetNumPostRelax**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*num\_post\_relax)

HYPRE\_Int **HYPRE\_StructSMGSetLogging**(*HYPRE\_StructSolver* solver, HYPRE\_Int logging)  
 (Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_StructSMGGetLogging**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*logging)

HYPRE\_Int **HYPRE\_StructSMGSetPrintLevel**(*HYPRE\_StructSolver* solver, HYPRE\_Int print\_level)  
 (Optional) Control how much information is printed.

- 0 : no printout (default)
- 1 : print convergence history

HYPRE\_Int **HYPRE\_StructSMGGetPrintLevel**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*print\_level)

HYPRE\_Int **HYPRE\_StructSMGGetNumIterations**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_StructSMGGetFinalRelativeResidualNorm**(*HYPRE\_StructSolver* solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

### Struct CycRed Solver

CycRed is a cyclic reduction solver that simultaneously solves a collection of 1D tridiagonal systems embedded in a d-dimensional grid.

HYPRE\_Int **HYPRE\_StructCycRedCreate**(MPI\_Comm comm, *HYPRE\_StructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_StructCycRedDestroy**(*HYPRE\_StructSolver* solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_StructCycRedSetup**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, HYPRE\_StructVector b, HYPRE\_StructVector x)

Prepare to solve the system.

The coefficient data in *b* and *x* is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_StructCycRedSolve**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, HYPRE\_StructVector b, HYPRE\_StructVector x)

Solve the system.

HYPRE\_Int **HYPRE\_StructCycRedSetTDim**(*HYPRE\_StructSolver* solver, HYPRE\_Int tdim)

(Optional) Set the dimension number for the embedded 1D tridiagonal systems.

The default is *tdim* = 0.

HYPRE\_Int **HYPRE\_StructCycRedSetBase**(*HYPRE\_StructSolver* solver, HYPRE\_Int ndim, HYPRE\_Int \*base\_index, HYPRE\_Int \*base\_stride)

(Optional) Set the base index and stride for the embedded 1D systems.

The stride must be equal one in the dimension corresponding to the 1D systems (see *HYPRE\_StructCycRedSetTDim*).

## Struct PCG Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_StructPCGCreate**(MPI\_Comm comm, *HYPRE\_StructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_StructPCGDestroy**(*HYPRE\_StructSolver* solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_StructPCGSetup**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, HYPRE\_StructVector b, HYPRE\_StructVector x)

HYPRE\_Int **HYPRE\_StructPCGSolve**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, HYPRE\_StructVector b, HYPRE\_StructVector x)

HYPRE\_Int **HYPRE\_StructPCGSetTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_StructPCGSetAbsoluteTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_StructPCGSetMaxIter**(*HYPRE\_StructSolver* solver, HYPRE\_Int max\_iter)

HYPRE\_Int **HYPRE\_StructPCGSetTwoNorm**(*HYPRE\_StructSolver* solver, HYPRE\_Int two\_norm)

HYPRE\_Int **HYPRE\_StructPCGSetRelChange**(*HYPRE\_StructSolver* solver, HYPRE\_Int rel\_change)

HYPRE\_Int **HYPRE\_StructPCGSetPrecond**(*HYPRE\_StructSolver* solver, *HYPRE\_PtrToStructSolverFcn* precondition, *HYPRE\_PtrToStructSolverFcn* precondition\_setup, *HYPRE\_StructSolver* precondition\_solver)

HYPRE\_Int **HYPRE\_StructPCGSetLogging**(*HYPRE\_StructSolver* solver, HYPRE\_Int logging)

HYPRE\_Int **HYPRE\_StructPCGSetPrintLevel**(*HYPRE\_StructSolver* solver, HYPRE\_Int level)

HYPRE\_Int **HYPRE\_StructPCGGetNumIterations**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int **HYPRE\_StructPCGGetFinalRelativeResidualNorm**(*HYPRE\_StructSolver* solver, HYPRE\_Real \*norm)

HYPRE\_Int **HYPRE\_StructPCGGetResidual**(*HYPRE\_StructSolver* solver, void \*\*residual)

HYPRE\_Int **HYPRE\_StructDiagScaleSetup**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, HYPRE\_StructVector y, HYPRE\_StructVector x)

Setup routine for diagonal preconditioning.

HYPRE\_Int **HYPRE\_StructDiagScale**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* HA, *HYPRE\_StructVector* Hy, *HYPRE\_StructVector* Hx)

Solve routine for diagonal preconditioning.

### Struct GMRES Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_StructGMRESCreate**(MPI\_Comm comm, *HYPRE\_StructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_StructGMRESDestroy**(*HYPRE\_StructSolver* solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_StructGMRESSetup**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, *HYPRE\_StructVector* b, *HYPRE\_StructVector* x)

HYPRE\_Int **HYPRE\_StructGMRESSolve**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, *HYPRE\_StructVector* b, *HYPRE\_StructVector* x)

HYPRE\_Int **HYPRE\_StructGMRESSetTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_StructGMRESSetAbsoluteTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_StructGMRESSetMaxIter**(*HYPRE\_StructSolver* solver, HYPRE\_Int max\_iter)

HYPRE\_Int **HYPRE\_StructGMRESSetKDim**(*HYPRE\_StructSolver* solver, HYPRE\_Int k\_dim)

HYPRE\_Int **HYPRE\_StructGMRESSetPrecond**(*HYPRE\_StructSolver* solver, *HYPRE\_PtrToStructSolverFcn* precondition, *HYPRE\_PtrToStructSolverFcn* precondition\_setup, *HYPRE\_StructSolver* precondition\_solver)

HYPRE\_Int **HYPRE\_StructGMRESSetLogging**(*HYPRE\_StructSolver* solver, HYPRE\_Int logging)

HYPRE\_Int **HYPRE\_StructGMRESSetPrintLevel**(*HYPRE\_StructSolver* solver, HYPRE\_Int level)

HYPRE\_Int **HYPRE\_StructGMRESGetNumIterations**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int **HYPRE\_StructGMRESGetFinalRelativeResidualNorm**(*HYPRE\_StructSolver* solver, HYPRE\_Real \*norm)

HYPRE\_Int **HYPRE\_StructGMRESGetResidual**(*HYPRE\_StructSolver* solver, void \*\*residual)

### Struct FlexGMRES Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_StructFlexGMRESCreate**(MPI\_Comm comm, *HYPRE\_StructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_StructFlexGMRESDestroy**(*HYPRE\_StructSolver* solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_StructFlexGMRESSetup**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, *HYPRE\_StructVector* b, *HYPRE\_StructVector* x)

**HYPRE\_Int HYPRE\_StructFlexGMRESSolve**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A,  
*HYPRE\_StructVector* b, *HYPRE\_StructVector* x)

**HYPRE\_Int HYPRE\_StructFlexGMRESSetTol**(*HYPRE\_StructSolver* solver, *HYPRE\_Real* tol)

**HYPRE\_Int HYPRE\_StructFlexGMRESSetAbsoluteTol**(*HYPRE\_StructSolver* solver, *HYPRE\_Real* tol)

**HYPRE\_Int HYPRE\_StructFlexGMRESSetMaxIter**(*HYPRE\_StructSolver* solver, *HYPRE\_Int* max\_iter)

**HYPRE\_Int HYPRE\_StructFlexGMRESSetKDim**(*HYPRE\_StructSolver* solver, *HYPRE\_Int* k\_dim)

**HYPRE\_Int HYPRE\_StructFlexGMRESSetPrecond**(*HYPRE\_StructSolver* solver,  
*HYPRE\_PtrToStructSolverFcn* precond,  
*HYPRE\_PtrToStructSolverFcn* precond\_setup,  
*HYPRE\_StructSolver* precond\_solver)

**HYPRE\_Int HYPRE\_StructFlexGMRESSetLogging**(*HYPRE\_StructSolver* solver, *HYPRE\_Int* logging)

**HYPRE\_Int HYPRE\_StructFlexGMRESSetPrintLevel**(*HYPRE\_StructSolver* solver, *HYPRE\_Int* level)

**HYPRE\_Int HYPRE\_StructFlexGMRESSetNumIterations**(*HYPRE\_StructSolver* solver, *HYPRE\_Int*  
\**num\_iterations*)

**HYPRE\_Int HYPRE\_StructFlexGMRESSetFinalRelativeResidualNorm**(*HYPRE\_StructSolver* solver,  
*HYPRE\_Real* \**norm*)

**HYPRE\_Int HYPRE\_StructFlexGMRESSetResidual**(*HYPRE\_StructSolver* solver, void \*\**residual*)

**HYPRE\_Int HYPRE\_StructFlexGMRESSetModifyPC**(*HYPRE\_StructSolver* solver,  
*HYPRE\_PtrToModifyPCFcn* modify\_pc)

### Struct LGMRES Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

**HYPRE\_Int HYPRE\_StructLGMRESCreate**(*MPI\_Comm* comm, *HYPRE\_StructSolver* \**solver*)  
Create a solver object.

**HYPRE\_Int HYPRE\_StructLGMRESDestroy**(*HYPRE\_StructSolver* solver)  
Destroy a solver object.

**HYPRE\_Int HYPRE\_StructLGMRESSetup**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A,  
*HYPRE\_StructVector* b, *HYPRE\_StructVector* x)

**HYPRE\_Int HYPRE\_StructLGMRESSolve**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A,  
*HYPRE\_StructVector* b, *HYPRE\_StructVector* x)

**HYPRE\_Int HYPRE\_StructLGMRESSetTol**(*HYPRE\_StructSolver* solver, *HYPRE\_Real* tol)

**HYPRE\_Int HYPRE\_StructLGMRESSetAbsoluteTol**(*HYPRE\_StructSolver* solver, *HYPRE\_Real* tol)

**HYPRE\_Int HYPRE\_StructLGMRESSetMaxIter**(*HYPRE\_StructSolver* solver, *HYPRE\_Int* max\_iter)

**HYPRE\_Int HYPRE\_StructLGMRESSetKDim**(*HYPRE\_StructSolver* solver, *HYPRE\_Int* k\_dim)

**HYPRE\_Int HYPRE\_StructLGMRESSetAugDim**(*HYPRE\_StructSolver* solver, *HYPRE\_Int* aug\_dim)

HYPRE\_Int HYPRE\_StructLGMRESSetPrecond(*HYPRE\_StructSolver* solver, *HYPRE\_PtrToStructSolverFcn* precondition, *HYPRE\_PtrToStructSolverFcn* precondition\_setup, *HYPRE\_StructSolver* precondition\_solver)

HYPRE\_Int HYPRE\_StructLGMRESSetLogging(*HYPRE\_StructSolver* solver, HYPRE\_Int logging)

HYPRE\_Int HYPRE\_StructLGMRESSetPrintLevel(*HYPRE\_StructSolver* solver, HYPRE\_Int level)

HYPRE\_Int HYPRE\_StructLGMRESGetNumIterations(*HYPRE\_StructSolver* solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int HYPRE\_StructLGMRESGetFinalRelativeResidualNorm(*HYPRE\_StructSolver* solver, HYPRE\_Real \*norm)

HYPRE\_Int HYPRE\_StructLGMRESGetResidual(*HYPRE\_StructSolver* solver, void \*\*residual)

### Struct BiCGSTAB Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int HYPRE\_StructBiCGSTABCreate(MPI\_Comm comm, *HYPRE\_StructSolver* \*solver)  
Create a solver object.

HYPRE\_Int HYPRE\_StructBiCGSTABDestroy(*HYPRE\_StructSolver* solver)  
Destroy a solver object.

HYPRE\_Int HYPRE\_StructBiCGSTABSetup(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, *HYPRE\_StructVector* b, *HYPRE\_StructVector* x)

HYPRE\_Int HYPRE\_StructBiCGSTABsolve(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A, *HYPRE\_StructVector* b, *HYPRE\_StructVector* x)

HYPRE\_Int HYPRE\_StructBiCGSTABSetTol(*HYPRE\_StructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int HYPRE\_StructBiCGSTABSetAbsoluteTol(*HYPRE\_StructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int HYPRE\_StructBiCGSTABSetMaxIter(*HYPRE\_StructSolver* solver, HYPRE\_Int max\_iter)

HYPRE\_Int HYPRE\_StructBiCGSTABSetPrecond(*HYPRE\_StructSolver* solver, *HYPRE\_PtrToStructSolverFcn* precondition, *HYPRE\_PtrToStructSolverFcn* precondition\_setup, *HYPRE\_StructSolver* precondition\_solver)

HYPRE\_Int HYPRE\_StructBiCGSTABSetLogging(*HYPRE\_StructSolver* solver, HYPRE\_Int logging)

HYPRE\_Int HYPRE\_StructBiCGSTABSetPrintLevel(*HYPRE\_StructSolver* solver, HYPRE\_Int level)

HYPRE\_Int HYPRE\_StructBiCGSTABGetNumIterations(*HYPRE\_StructSolver* solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int HYPRE\_StructBiCGSTABGetFinalRelativeResidualNorm(*HYPRE\_StructSolver* solver, HYPRE\_Real \*norm)

HYPRE\_Int HYPRE\_StructBiCGSTABGetResidual(*HYPRE\_StructSolver* solver, void \*\*residual)

## Struct Hybrid Solver

HYPRE\_Int **HYPRE\_StructHybridCreate**(MPI\_Comm comm, *HYPRE\_StructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_StructHybridDestroy**(*HYPRE\_StructSolver* solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_StructHybridSetup**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A,  
HYPRE\_StructVector b, HYPRE\_StructVector x)

Prepare to solve the system.

The coefficient data in  $b$  and  $x$  is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_StructHybridSolve**(*HYPRE\_StructSolver* solver, *HYPRE\_StructMatrix* A,  
HYPRE\_StructVector b, HYPRE\_StructVector x)

Solve the system.

HYPRE\_Int **HYPRE\_StructHybridSetTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance.

HYPRE\_Int **HYPRE\_StructHybridSetConvergenceTol**(*HYPRE\_StructSolver* solver, HYPRE\_Real cf\_tol)

(Optional) Set an accepted convergence tolerance for diagonal scaling (DS).

The solver will switch preconditioners if the convergence of DS is slower than  $cf\_tol$ .

HYPRE\_Int **HYPRE\_StructHybridSetDSCGMaxIter**(*HYPRE\_StructSolver* solver, HYPRE\_Int ds\_max\_its)

(Optional) Set maximum number of iterations for diagonal scaling (DS).

The solver will switch preconditioners if DS reaches  $ds\_max\_its$ .

HYPRE\_Int **HYPRE\_StructHybridSetPCGMaxIter**(*HYPRE\_StructSolver* solver, HYPRE\_Int pre\_max\_its)

(Optional) Set maximum number of iterations for general preconditioner (PRE).

The solver will stop if PRE reaches  $pre\_max\_its$ .

HYPRE\_Int **HYPRE\_StructHybridSetTwoNorm**(*HYPRE\_StructSolver* solver, HYPRE\_Int two\_norm)

(Optional) Use the two-norm in stopping criteria.

HYPRE\_Int **HYPRE\_StructHybridSetStopCrit**(*HYPRE\_StructSolver* solver, HYPRE\_Int stop\_crit)

HYPRE\_Int **HYPRE\_StructHybridSetRelChange**(*HYPRE\_StructSolver* solver, HYPRE\_Int rel\_change)

(Optional) Additionally require that the relative difference in successive iterates be small.

HYPRE\_Int **HYPRE\_StructHybridSetSolverType**(*HYPRE\_StructSolver* solver, HYPRE\_Int solver\_type)

(Optional) Set the type of Krylov solver to use.

Current krylov methods set by  $solver\_type$  are:

- 0 : PCG (default)
- 1 : GMRES
- 2 : BiCGSTAB

HYPRE\_Int **HYPRE\_StructHybridSetRecomputeResidual**(*HYPRE\_StructSolver* solver, HYPRE\_Int  
recompute\_residual)

(Optional) Set recompute residual (don't rely on 3-term recurrence).

HYPRE\_Int **HYPRE\_StructHybridGetRecomputeResidual**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*recompute\_residual)

(Optional) Get recompute residual option.

HYPRE\_Int **HYPRE\_StructHybridSetRecomputeResidualP**(*HYPRE\_StructSolver* solver, HYPRE\_Int recompute\_residual\_p)

(Optional) Set recompute residual period (don't rely on 3-term recurrence).

Recomputes residual after every specified number of iterations.

HYPRE\_Int **HYPRE\_StructHybridGetRecomputeResidualP**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*recompute\_residual\_p)

(Optional) Get recompute residual period option.

HYPRE\_Int **HYPRE\_StructHybridSetKDim**(*HYPRE\_StructSolver* solver, HYPRE\_Int k\_dim)

(Optional) Set the maximum size of the Krylov space when using GMRES.

HYPRE\_Int **HYPRE\_StructHybridSetPrecond**(*HYPRE\_StructSolver* solver, *HYPRE\_PtrToStructSolverFcn* precond, *HYPRE\_PtrToStructSolverFcn* precond\_setup, *HYPRE\_StructSolver* precond\_solver)

(Optional) Set the preconditioner to use.

HYPRE\_Int **HYPRE\_StructHybridSetLogging**(*HYPRE\_StructSolver* solver, HYPRE\_Int logging)

(Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_StructHybridSetPrintLevel**(*HYPRE\_StructSolver* solver, HYPRE\_Int print\_level)

(Optional) Set the amount of printing to do to the screen.

HYPRE\_Int **HYPRE\_StructHybridGetNumIterations**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*num\_its)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_StructHybridGetDSCGNumIterations**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*ds\_num\_its)

Return the number of diagonal scaling iterations taken.

HYPRE\_Int **HYPRE\_StructHybridGetPCGNumIterations**(*HYPRE\_StructSolver* solver, HYPRE\_Int \*pre\_num\_its)

Return the number of general preconditioning iterations taken.

HYPRE\_Int **HYPRE\_StructHybridGetFinalRelativeResidualNorm**(*HYPRE\_StructSolver* solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

HYPRE\_Int **HYPRE\_StructHybridSetPCGAbsoluteTolFactor**(*HYPRE\_StructSolver* solver, HYPRE\_Real pcg\_atolf)

## Struct LOBPCG Eigensolver

These routines should be used in conjunction with the generic interface in *Eigensolvers*.

HYPRE\_Int **HYPRE\_StructSetupInterpreter**(mv\_InterfaceInterpreter \*i)

Load interface interpreter.

Vector part loaded with *hypr\_StructKrylov* functions and multivector part loaded with *mv\_TempMultiVector* functions.

HYPRE\_Int **HYPRE\_StructSetupMatvec**(HYPRE\_MatvecFunctions \*mv)

Load Matvec interpreter with hyre\_StructKrylov functions.

## 8.6 SStruct Solvers

### group SStruct Solvers

Linear solvers for semi-structured grids.

These solvers use matrix/vector storage schemes that are tailored to semi-structured grid problems.

### SStruct Solvers

typedef struct hyre\_SStructSolver\_struct **\*HYPRE\_SStructSolver**

The solver object.

typedef HYPRE\_Int (**\*HYPRE\_PtrToSStructSolverFcn**)(HYPRE\_SStructSolver, HYPRE\_SStructMatrix, HYPRE\_SStructVector, HYPRE\_SStructVector)

### SStruct SysPFMG Solver

SysPFMG is a semicoarsening multigrid solver similar to PFMG, but for systems of PDEs.

For periodic problems, users should try to set the grid size in periodic dimensions to be as close to a power-of-two as possible (for more details, see Struct PFMG Solver).

HYPRE\_Int **HYPRE\_SStructSysPFMGCreate**(MPI\_Comm comm, HYPRE\_SStructSolver \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_SStructSysPFMGDestroy**(HYPRE\_SStructSolver solver)

Destroy a solver object.

An object should be explicitly destroyed using this destructor when the user's code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetup**(HYPRE\_SStructSolver solver, HYPRE\_SStructMatrix A, HYPRE\_SStructVector b, HYPRE\_SStructVector x)

Prepare to solve the system.

The coefficient data in  $b$  and  $x$  is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_SStructSysPFMGsolve**(HYPRE\_SStructSolver solver, HYPRE\_SStructMatrix A, HYPRE\_SStructVector b, HYPRE\_SStructVector x)

Solve the system.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetTol**(HYPRE\_SStructSolver solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetMaxIter**(HYPRE\_SStructSolver solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetRelChange**(HYPRE\_SStructSolver solver, HYPRE\_Int rel\_change)

(Optional) Additionally require that the relative difference in successive iterates be small.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetZeroGuess**(*HYPRE\_SStructSolver* solver)

(Optional) Use a zero initial guess.

This allows the solver to cut corners in the case where a zero initial guess is needed (e.g., for preconditioning) to reduce computational cost.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetNonZeroGuess**(*HYPRE\_SStructSolver* solver)

(Optional) Use a nonzero initial guess.

This is the default behavior, but this routine allows the user to switch back after using *SetZeroGuess*.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetRelaxType**(*HYPRE\_SStructSolver* solver, HYPRE\_Int relax\_type)

(Optional) Set relaxation type.

Current relaxation methods set by *relax\_type* are:

- 0 : Jacobi
- 1 : Weighted Jacobi (default)
- 2 : Red/Black Gauss-Seidel (symmetric: RB pre-relaxation, BR post-relaxation)

HYPRE\_Int **HYPRE\_SStructSysPFMGSetJacobiWeight**(*HYPRE\_SStructSolver* solver, HYPRE\_Real weight)

(Optional) Set Jacobi Weight.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetNumPreRelax**(*HYPRE\_SStructSolver* solver, HYPRE\_Int num\_pre\_relax)

(Optional) Set number of relaxation sweeps before coarse-grid correction.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetNumPostRelax**(*HYPRE\_SStructSolver* solver, HYPRE\_Int num\_post\_relax)

(Optional) Set number of relaxation sweeps after coarse-grid correction.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetSkipRelax**(*HYPRE\_SStructSolver* solver, HYPRE\_Int skip\_relax)

(Optional) Skip relaxation on certain grids for isotropic problems.

This can greatly improve efficiency by eliminating unnecessary relaxations when the underlying problem is isotropic.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetDxyz**(*HYPRE\_SStructSolver* solver, HYPRE\_Real \*dxyz)

HYPRE\_Int **HYPRE\_SStructSysPFMGSetLogging**(*HYPRE\_SStructSolver* solver, HYPRE\_Int logging)

(Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_SStructSysPFMGSetPrintLevel**(*HYPRE\_SStructSolver* solver, HYPRE\_Int print\_level)

(Optional) Control how much information is printed.

- 0 : no printout (default)
- 1 : print convergence history

HYPRE\_Int **HYPRE\_SStructSysPFMGGetNumIterations**(*HYPRE\_SStructSolver* solver, HYPRE\_Int \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_SStructSysPFMGGetFinalRelativeResidualNorm**(*HYPRE\_SStructSolver* solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

### SStruct SSAMG Solver

The semi-structured algebraic multigrid (SSAMG) method is an iterative solver that can handle problems with multiple parts such as block-structured and structured adaptive mesh refinement grids (SAMR).

HYPRE\_Int **HYPRE\_SStructSSAMGCreate**(MPI\_Comm comm, *HYPRE\_SStructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_SStructSSAMGDestroy**(*HYPRE\_SStructSolver* solver)

Destroy a solver object.

An object should be explicitly destroyed using this destructor when the user's code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

HYPRE\_Int **HYPRE\_SStructSSAMGSetup**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

Prepare to solve the system.

The coefficient data in `{\tt b}` and `{\tt x}` is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_SStructSSAMGSolve**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

Solve the system.

HYPRE\_Int **HYPRE\_SStructSSAMGSetTol**(*HYPRE\_SStructSolver* solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance.

HYPRE\_Int **HYPRE\_SStructSSAMGSetMaxIter**(*HYPRE\_SStructSolver* solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_SStructSSAMGSetMaxLevels**(*HYPRE\_SStructSolver* solver, HYPRE\_Int max\_levels)

(Optional) Set maximum number of levels of the multigrid hierarchy.

HYPRE\_Int **HYPRE\_SStructSSAMGSetRelChange**(*HYPRE\_SStructSolver* solver, HYPRE\_Int rel\_change)

(Optional) Additionally require that the relative difference in successive iterates be small.

HYPRE\_Int **HYPRE\_SStructSSAMGSetNonGalerkinRAP**(*HYPRE\_SStructSolver* solver, HYPRE\_Int non\_galerkin)

(Optional) Set type of coarse-grid operator to use.

Current operators set by `{\tt rap_type}` are:

- 0 : Galerkin (default)
- 1 : non-Galerkin 5-pt or 7-pt stencils

Both operators are constructed algebraically. The non-Galerkin option maintains a 5-pt stencil in 2D and a 7-pt stencil in 3D on all grid levels. The stencil coefficients are computed by averaging techniques.

HYPRE\_Int **HYPRE\_SStructSSAMGSetZeroGuess**(*HYPRE\_SStructSolver* solver)

(Optional) Use a zero initial guess.

This allows the solver to cut corners in the case where a zero initial guess is needed (e.g., for preconditioning) to reduce computational cost.

HYPRE\_Int **HYPRE\_SStructSSAMGSetNonZeroGuess**(*HYPRE\_SStructSolver* solver)

(Optional) Use a nonzero initial guess.

This is the default behavior, but this routine allows the user to switch back after using `{\tt SetZeroGuess}`.

HYPRE\_Int **HYPRE\_SStructSSAMGSetInterpType**(*HYPRE\_SStructSolver* solver, HYPRE\_Int interp\_type)

(Optional) Set interpolation type.

Current interpolation methods set by `{\tt interp_type}` are:

- -1 : Structured interpolation only (default)
- 0 : Structured and classical modified unstructured interpolation

HYPRE\_Int **HYPRE\_SStructSSAMGSetRelaxType**(*HYPRE\_SStructSolver* solver, HYPRE\_Int relax\_type)

(Optional) Set relaxation type.

Current relaxation methods set by `{\tt relax_type}` are:

- 0 : Jacobi
- 1 : Weighted Jacobi (default)
- 2 : L1-Jacobi
- 10 : Red/Black Gauss-Seidel (symmetric: RB pre-relaxation, BR post-relaxation)

HYPRE\_Int **HYPRE\_SStructSSAMGSetSkipRelax**(*HYPRE\_SStructSolver* solver, HYPRE\_Int skip\_relax)

(Optional) Skip relaxation on certain grids for isotropic problems.

This can greatly improve efficiency by eliminating unnecessary relaxations when the underlying problem is isotropic.

HYPRE\_Int **HYPRE\_SStructSSAMGSetRelaxWeight**(*HYPRE\_SStructSolver* solver, HYPRE\_Real weight)

(Optional) Set Jacobi Weight.

HYPRE\_Int **HYPRE\_SStructSSAMGSetNumPreRelax**(*HYPRE\_SStructSolver* solver, HYPRE\_Int num\_pre\_relax)

(Optional) Set number of relaxation sweeps before coarse-grid correction.

HYPRE\_Int **HYPRE\_SStructSSAMGSetNumPostRelax**(*HYPRE\_SStructSolver* solver, HYPRE\_Int num\_post\_relax)

(Optional) Set number of relaxation sweeps after coarse-grid correction.

HYPRE\_Int **HYPRE\_SStructSSAMGSetNumCoarseRelax**(*HYPRE\_SStructSolver* solver, HYPRE\_Int num\_coarse\_relax)

(Optional) Set number of relaxation sweeps in the coarse grid.

HYPRE\_Int **HYPRE\_SStructSSAMGSetMaxCoarseSize**(*HYPRE\_SStructSolver* solver, HYPRE\_Int max\_coarse\_size)

(Optional) Set maximum size of coarse grid.

This option can be disabled by setting `max_coarse_size` to zero.

HYPRE\_Int **HYPRE\_SStructSSAMGSetCoarseSolverType**(*HYPRE\_SStructSolver* solver, HYPRE\_Int csolver\_type)

(Optional) Set coarse solver type for SSAMG.

Current options are

- 0 : Weighted Jacobi (default)
- 1 : BoomerAMG

HYPRE\_Int **HYPRE\_SStructSSAMGSetDxyz**(*HYPRE\_SStructSolver* solver, HYPRE\_Int nparts, HYPRE\_Real \*\*dxyz)

(Optional) Set the grid spacing metric used for coarsening purposes for each part.

HYPRE\_Int **HYPRE\_SStructSSAMGSetLogging**(*HYPRE\_SStructSolver* solver, HYPRE\_Int logging)

(Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_SStructSSAMGSetPrintLevel**(*HYPRE\_SStructSolver* solver, HYPRE\_Int print\_level)

(Optional) Control how much information is printed.

- 0 : no printout (default)
- 1 : print setup info
- 2 : print convergence history

HYPRE\_Int **HYPRE\_SStructSSAMGSetPrintFreq**(*HYPRE\_SStructSolver* solver, HYPRE\_Int print\_freq)

(Optional) Set printing frequency.

HYPRE\_Int **HYPRE\_SStructSSAMGGetNumIterations**(*HYPRE\_SStructSolver* solver, HYPRE\_Int \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_SStructSSAMGGetFinalRelativeResidualNorm**(*HYPRE\_SStructSolver* solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

## SStruct Split Solver

HYPRE\_Int **HYPRE\_SStructSplitCreate**(MPI\_Comm comm, *HYPRE\_SStructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_SStructSplitDestroy**(*HYPRE\_SStructSolver* solver)

Destroy a solver object.

An object should be explicitly destroyed using this destructor when the user's code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

HYPRE\_Int **HYPRE\_SStructSplitSetup**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

Prepare to solve the system.

The coefficient data in *b* and *x* is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_SStructSplitSolve**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

Solve the system.

HYPRE\_Int **HYPRE\_SStructSplitSetTol**(*HYPRE\_SStructSolver* solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance.

HYPRE\_Int **HYPRE\_SStructSplitSetMaxIter**(*HYPRE\_SStructSolver* solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_SStructSplitSetPrintLevel**(*HYPRE\_SStructSolver* solver, HYPRE\_Int print\_level)

(Optional) Control how much information is printed.

- 0 : no printout (default)
- 1 : print convergence history

HYPRE\_Int **HYPRE\_SStructSplitSetLogging**(*HYPRE\_SStructSolver* solver, HYPRE\_Int logging)

(Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_SStructSplitSetZeroGuess**(*HYPRE\_SStructSolver* solver)

(Optional) Use a zero initial guess.

This allows the solver to cut corners in the case where a zero initial guess is needed (e.g., for preconditioning) to reduce computational cost.

HYPRE\_Int **HYPRE\_SStructSplitSetNonZeroGuess**(*HYPRE\_SStructSolver* solver)

(Optional) Use a nonzero initial guess.

This is the default behavior, but this routine allows the user to switch back after using *SetZeroGuess*.

HYPRE\_Int **HYPRE\_SStructSplitSetStructSolver**(*HYPRE\_SStructSolver* solver, HYPRE\_Int ssolver)

(Optional) Set up the type of diagonal struct solver.

Either *ssolver* is set to *HYPRE\_SMG* or *HYPRE\_PFMG*.

HYPRE\_Int **HYPRE\_SStructSplitGetNumIterations**(*HYPRE\_SStructSolver* solver, HYPRE\_Int \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_SStructSplitGetFinalRelativeResidualNorm**(*HYPRE\_SStructSolver* solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

HYPRE\_Int **HYPRE\_SStructSplitPrintLogging**(*HYPRE\_SStructSolver* solver)

**HYPRE\_PFMG**

**HYPRE\_SMG**

**HYPRE\_Jacobi**

## SStruct PCG Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_SStructPCGCreate**(MPI\_Comm comm, *HYPRE\_SStructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_SStructPCGDestroy**(*HYPRE\_SStructSolver* solver)

Destroy a solver object.

An object should be explicitly destroyed using this destructor when the user's code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

HYPRE\_Int **HYPRE\_SStructPCGSetup**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

HYPRE\_Int **HYPRE\_SStructPCGSolve**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

HYPRE\_Int **HYPRE\_SStructPCGSetTol**(*HYPRE\_SStructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_SStructPCGSetAbsoluteTol**(*HYPRE\_SStructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_SStructPCGSetMaxIter**(*HYPRE\_SStructSolver* solver, HYPRE\_Int max\_iter)

HYPRE\_Int **HYPRE\_SStructPCGSetTwoNorm**(*HYPRE\_SStructSolver* solver, HYPRE\_Int two\_norm)

HYPRE\_Int **HYPRE\_SStructPCGSetRelChange**(*HYPRE\_SStructSolver* solver, HYPRE\_Int rel\_change)

HYPRE\_Int **HYPRE\_SStructPCGSetPrecond**(*HYPRE\_SStructSolver* solver, *HYPRE\_PtrToSStructSolverFcn* precondition, *HYPRE\_PtrToSStructSolverFcn* precondition\_setup, void \*precond\_solver)

HYPRE\_Int **HYPRE\_SStructPCGSetLogging**(*HYPRE\_SStructSolver* solver, HYPRE\_Int logging)

HYPRE\_Int **HYPRE\_SStructPCGSetPrintLevel**(*HYPRE\_SStructSolver* solver, HYPRE\_Int level)

HYPRE\_Int **HYPRE\_SStructPCGGetNumIterations**(*HYPRE\_SStructSolver* solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int **HYPRE\_SStructPCGGetFinalRelativeResidualNorm**(*HYPRE\_SStructSolver* solver, HYPRE\_Real \*norm)

HYPRE\_Int **HYPRE\_SStructPCGGetResidual**(*HYPRE\_SStructSolver* solver, void \*\*residual)

HYPRE\_Int **HYPRE\_SStructDiagScaleSetup**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* y, *HYPRE\_SStructVector* x)

Setup routine for diagonal preconditioning.

HYPRE\_Int **HYPRE\_SStructDiagScale**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* y, *HYPRE\_SStructVector* x)

Solve routine for diagonal preconditioning.

### SStruct GMRES Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_SStructGMRESCreate**(MPI\_Comm comm, *HYPRE\_SStructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_SStructGMRESDestroy**(*HYPRE\_SStructSolver* solver)

Destroy a solver object.

An object should be explicitly destroyed using this destructor when the user's code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

HYPRE\_Int **HYPRE\_SStructGMRESSetup**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

HYPRE\_Int **HYPRE\_SStructGMRESSolve**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

HYPRE\_Int **HYPRE\_SStructGMRESSetTol**(*HYPRE\_SStructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_SStructGMRESSetAbsoluteTol**(*HYPRE\_SStructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_SStructGMRESSetMinIter**(*HYPRE\_SStructSolver* solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_SStructGMRESSetMaxIter**(*HYPRE\_SStructSolver* solver, HYPRE\_Int max\_iter)

HYPRE\_Int **HYPRE\_SStructGMRESSetKDim**(*HYPRE\_SStructSolver* solver, HYPRE\_Int k\_dim)

HYPRE\_Int **HYPRE\_SStructGMRESSetStopCrit**(*HYPRE\_SStructSolver* solver, HYPRE\_Int stop\_crit)

HYPRE\_Int **HYPRE\_SStructGMRESSetPrecond**(*HYPRE\_SStructSolver* solver, *HYPRE\_PtrToSStructSolverFcn* precondition, *HYPRE\_PtrToSStructSolverFcn* precondition\_setup, void \*precond\_solver)

HYPRE\_Int **HYPRE\_SStructGMRESSetLogging**(*HYPRE\_SStructSolver* solver, HYPRE\_Int logging)

HYPRE\_Int **HYPRE\_SStructGMRESSetPrintLevel**(*HYPRE\_SStructSolver* solver, HYPRE\_Int print\_level)

HYPRE\_Int **HYPRE\_SStructGMRESGetNumIterations**(*HYPRE\_SStructSolver* solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int **HYPRE\_SStructGMRESGetFinalRelativeResidualNorm**(*HYPRE\_SStructSolver* solver, HYPRE\_Real \*norm)

HYPRE\_Int **HYPRE\_SStructGMRESGetResidual**(*HYPRE\_SStructSolver* solver, void \*\*residual)

### SStruct FlexGMRES Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_SStructFlexGMRESCreate**(MPI\_Comm comm, *HYPRE\_SStructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_SStructFlexGMRESDestroy**(*HYPRE\_SStructSolver* solver)

Destroy a solver object.

An object should be explicitly destroyed using this destructor when the user's code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

HYPRE\_Int **HYPRE\_SStructFlexGMRESSetup**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

HYPRE\_Int **HYPRE\_SStructFlexGMRESSolve**(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

HYPRE\_Int **HYPRE\_SStructFlexGMRESSetTol**(*HYPRE\_SStructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_SStructFlexGMRESSetAbsoluteTol**(*HYPRE\_SStructSolver* solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_SStructFlexGMRESSetMinIter**(*HYPRE\_SStructSolver* solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_SStructFlexGMRESSetMaxIter**(*HYPRE\_SStructSolver* solver, HYPRE\_Int max\_iter)

HYPRE\_Int **HYPRE\_SStructFlexGMRESSetKDim**(*HYPRE\_SStructSolver* solver, HYPRE\_Int k\_dim)

HYPRE\_Int **HYPRE\_SStructFlexGMRESSetPrecond**(*HYPRE\_SStructSolver* solver, *HYPRE\_PtrToSStructSolverFcn* precond, *HYPRE\_PtrToSStructSolverFcn* precond\_setup, void \*precond\_solver)

HYPRE\_Int **HYPRE\_SStructFlexGMRESSetLogging**(*HYPRE\_SStructSolver* solver, HYPRE\_Int logging)

HYPRE\_Int **HYPRE\_SStructFlexGMRESSetPrintLevel**(*HYPRE\_SStructSolver* solver, HYPRE\_Int print\_level)

HYPRE\_Int **HYPRE\_SStructFlexGMRESGetNumIterations**(*HYPRE\_SStructSolver* solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int **HYPRE\_SStructFlexGMRESGetFinalRelativeResidualNorm**(*HYPRE\_SStructSolver* solver, HYPRE\_Real \*norm)

HYPRE\_Int **HYPRE\_SStructFlexGMRESGetResidual**(*HYPRE\_SStructSolver* solver, void \*\*residual)

HYPRE\_Int **HYPRE\_SStructFlexGMRESSetModifyPC**(*HYPRE\_SStructSolver* solver, *HYPRE\_PtrToModifyPCFcn* modify\_pc)

### SStruct LGMRES Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_SStructLGMRESCreate**(MPI\_Comm comm, *HYPRE\_SStructSolver* \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_SStructLGMRESDestroy**(*HYPRE\_SStructSolver* solver)

Destroy a solver object.

An object should be explicitly destroyed using this destructor when the user's code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

`HYPRE_Int HYPRE_SStructLGMRESSetup`(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

`HYPRE_Int HYPRE_SStructLGMRESSolve`(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

`HYPRE_Int HYPRE_SStructLGMRESSetTol`(*HYPRE\_SStructSolver* solver, *HYPRE\_Real* tol)

`HYPRE_Int HYPRE_SStructLGMRESSetAbsoluteTol`(*HYPRE\_SStructSolver* solver, *HYPRE\_Real* tol)

`HYPRE_Int HYPRE_SStructLGMRESSetMinIter`(*HYPRE\_SStructSolver* solver, *HYPRE\_Int* min\_iter)

`HYPRE_Int HYPRE_SStructLGMRESSetMaxIter`(*HYPRE\_SStructSolver* solver, *HYPRE\_Int* max\_iter)

`HYPRE_Int HYPRE_SStructLGMRESSetKDim`(*HYPRE\_SStructSolver* solver, *HYPRE\_Int* k\_dim)

`HYPRE_Int HYPRE_SStructLGMRESSetAugDim`(*HYPRE\_SStructSolver* solver, *HYPRE\_Int* aug\_dim)

`HYPRE_Int HYPRE_SStructLGMRESSetPrecond`(*HYPRE\_SStructSolver* solver, *HYPRE\_PtrToSStructSolverFcn* precondition, *HYPRE\_PtrToSStructSolverFcn* precondition\_setup, void \*precond\_solver)

`HYPRE_Int HYPRE_SStructLGMRESSetLogging`(*HYPRE\_SStructSolver* solver, *HYPRE\_Int* logging)

`HYPRE_Int HYPRE_SStructLGMRESSetPrintLevel`(*HYPRE\_SStructSolver* solver, *HYPRE\_Int* print\_level)

`HYPRE_Int HYPRE_SStructLGMRESGetNumIterations`(*HYPRE\_SStructSolver* solver, *HYPRE\_Int* \*num\_iterations)

`HYPRE_Int HYPRE_SStructLGMRESGetFinalRelativeResidualNorm`(*HYPRE\_SStructSolver* solver, *HYPRE\_Real* \*norm)

`HYPRE_Int HYPRE_SStructLGMRESGetResidual`(*HYPRE\_SStructSolver* solver, void \*\*residual)

### **SStruct BiCGSTAB Solver**

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

`HYPRE_Int HYPRE_SStructBiCGSTABCreate`(*MPI\_Comm* comm, *HYPRE\_SStructSolver* \*solver)

Create a solver object.

`HYPRE_Int HYPRE_SStructBiCGSTABDestroy`(*HYPRE\_SStructSolver* solver)

Destroy a solver object.

An object should be explicitly destroyed using this destructor when the user's code no longer needs direct access to it. Once destroyed, the object must not be referenced again. Note that the object may not be deallocated at the completion of this call, since there may be internal package references to the object. The object will then be destroyed when all internal reference counts go to zero.

`HYPRE_Int HYPRE_SStructBiCGSTABSetup`(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

`HYPRE_Int HYPRE_SStructBiCGSTABSolve`(*HYPRE\_SStructSolver* solver, *HYPRE\_SStructMatrix* A, *HYPRE\_SStructVector* b, *HYPRE\_SStructVector* x)

`HYPRE_Int HYPRE_SStructBiCGSTABSetTol`(*HYPRE\_SStructSolver* solver, *HYPRE\_Real* tol)

HYPRE\_Int **HYPRE\_SStructBiCGSTABSetAbsoluteTol**(*HYPRE\_SStructSolver* solver, HYPRE\_Real tol)  
 HYPRE\_Int **HYPRE\_SStructBiCGSTABSetMinIter**(*HYPRE\_SStructSolver* solver, HYPRE\_Int min\_iter)  
 HYPRE\_Int **HYPRE\_SStructBiCGSTABSetMaxIter**(*HYPRE\_SStructSolver* solver, HYPRE\_Int max\_iter)  
 HYPRE\_Int **HYPRE\_SStructBiCGSTABSetStopCrit**(*HYPRE\_SStructSolver* solver, HYPRE\_Int stop\_crit)  
 HYPRE\_Int **HYPRE\_SStructBiCGSTABSetPrecond**(*HYPRE\_SStructSolver* solver,  
   *HYPRE\_PtrToSStructSolverFcn* precondition,  
   *HYPRE\_PtrToSStructSolverFcn* precondition\_setup, void  
   \*precond\_solver)  
  
 HYPRE\_Int **HYPRE\_SStructBiCGSTABSetLogging**(*HYPRE\_SStructSolver* solver, HYPRE\_Int logging)  
 HYPRE\_Int **HYPRE\_SStructBiCGSTABSetPrintLevel**(*HYPRE\_SStructSolver* solver, HYPRE\_Int level)  
 HYPRE\_Int **HYPRE\_SStructBiCGSTABGetNumIterations**(*HYPRE\_SStructSolver* solver, HYPRE\_Int  
   \*num\_iterations)  
  
 HYPRE\_Int **HYPRE\_SStructBiCGSTABGetFinalRelativeResidualNorm**(*HYPRE\_SStructSolver* solver,  
   HYPRE\_Real \*norm)  
  
 HYPRE\_Int **HYPRE\_SStructBiCGSTABGetResidual**(*HYPRE\_SStructSolver* solver, void \*\*residual)

### SStruct LOBPCG Eigensolver

These routines should be used in conjunction with the generic interface in *Eigensolvers*.

HYPRE\_Int **HYPRE\_SStructSetupInterpreter**(mv\_InterfaceInterpreter \*i)

Load interface interpreter.

Vector part loaded with *hyre\_SStructKrylov* functions and multivector part loaded with *mv\_TempMultiVector* functions.

HYPRE\_Int **HYPRE\_SStructSetupMatvec**(HYPRE\_MatvecFunctions \*mv)

Load Matvec interpreter with *hyre\_SStructKrylov* functions.

## 8.7 ParCSR Solvers

### group ParCSR Solvers

Linear solvers for sparse matrix systems.

These solvers use matrix/vector storage schemes that are tailored for general sparse matrix systems.

### ParCSR Solvers

typedef HYPRE\_Int (\***HYPRE\_PtrToParSolverFcn**)(HYPRE\_Solver, *HYPRE\_ParCSRMatrix*,  
*HYPRE\_ParVector*, *HYPRE\_ParVector*)

The solver object.

### ParCSR BoomerAMG Solver and Preconditioner

Parallel unstructured algebraic multigrid solver and preconditioner

HYPRE\_Int **HYPRE\_BoomerAMGCreate**(HYPRE\_Solver \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_BoomerAMGDestroy**(HYPRE\_Solver solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_BoomerAMGSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Set up the BoomerAMG solver or preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] object to be set up.
- **A** – [IN] ParCSR matrix used to construct the solver/preconditioner.
- **b** – Ignored by this function.
- **x** – Ignored by this function.

HYPRE\_Int **HYPRE\_BoomerAMGSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Solve the system or apply AMG as a preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] solver or preconditioner object to be applied.
- **A** – [IN] ParCSR matrix, matrix of the linear system to be solved
- **b** – [IN] right hand side of the linear system to be solved
- **x** – [OUT] approximated solution of the linear system to be solved

HYPRE\_Int **HYPRE\_BoomerAMGSolveT**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Solve the transpose system  $A^T x = b$  or apply AMG as a preconditioner to the transpose system .

Note that this function should only be used when preconditioning CGNR with BoomerAMG. It can only be used with Jacobi smoothing (relax\_type 0 or 7) and without CF smoothing, i.e relax\_order needs to be set to 0. If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] solver or preconditioner object to be applied.
- **A** – [IN] ParCSR matrix
- **b** – [IN] right hand side of the linear system to be solved
- **x** – [OUT] approximated solution of the linear system to be solved

HYPRE\_Int **HYPRE\_BoomerAMGSetOldDefault**(HYPRE\_Solver solver)

Recovers old default for coarsening and interpolation, i.e Falgout coarsening and untruncated modified classical interpolation.

This option might be preferred for 2 dimensional problems.

HYPRE\_Int **HYPRE\_BoomerAMGGetResidual**(HYPRE\_Solver solver, *HYPRE\_ParVector* \*residual)

Returns the residual.

HYPRE\_Int **HYPRE\_BoomerAMGGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

Returns the number of iterations taken.

HYPRE\_Int **HYPRE\_BoomerAMGGetCumNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*cum\_num\_iterations)

HYPRE\_Int **HYPRE\_BoomerAMGGetCumNnzAP**(HYPRE\_Solver solver, HYPRE\_Real \*cum\_nnz\_AP)

Returns cumulative num of nonzeros for A and P operators.

HYPRE\_Int **HYPRE\_BoomerAMGSetCumNnzAP**(HYPRE\_Solver solver, HYPRE\_Real cum\_nnz\_AP)

Activates cumulative num of nonzeros for A and P operators.

Needs to be set to a positive number for activation.

HYPRE\_Int **HYPRE\_BoomerAMGGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*rel\_resid\_norm)

Returns the norm of the final relative residual.

HYPRE\_Int **HYPRE\_BoomerAMGSetNumFunctions**(HYPRE\_Solver solver, HYPRE\_Int num\_functions)

(Optional) Sets the size of the system of PDEs, if using the systems version.

The default is 1, i.e. a scalar system.

HYPRE\_Int **HYPRE\_BoomerAMGGetNumFunctions**(HYPRE\_Solver solver, HYPRE\_Int \*num\_functions)

HYPRE\_Int **HYPRE\_BoomerAMGSetFilterFunctions**(HYPRE\_Solver solver, HYPRE\_Int filter\_functions)

(Optional) Sets filtering for system of PDEs (*num\_functions* > 1).

**Note**

This option assumes that variables are stored in an interleaved format, where multiple variables are combined in a single vector. Enabling filtering can be beneficial when the problem has multiple coupled variables (functions) that are not strongly coupled.

**Parameters**

**filter\_functions** – An integer flag to enable or disable filtering of inter-variable connections in the input matrix used for preconditioning.

- A value of 0 (default) indicates no filtering, preserving all inter-variable connections.
- A value of 1 enables filtering, removing inter-variable connections to lower operator and memory complexities.

HYPRE\_Int **HYPRE\_BoomerAMGGetFilterFunctions**(HYPRE\_Solver solver, HYPRE\_Int \*filter\_functions)

HYPRE\_Int **HYPRE\_BoomerAMGSetDofFunc**(HYPRE\_Solver solver, HYPRE\_Int \*dof\_func)

(Optional) Sets the mapping that assigns the function to each variable, if using the systems version.

If no assignment is made and the number of functions is  $k > 1$ , the mapping generated is  $(0, 1, \dots, k-1, 0, 1, \dots, k-1, \dots)$ .

HYPRE\_Int **HYPRE\_BoomerAMGSetConvergeType**(HYPRE\_Solver solver, HYPRE\_Int type)  
 (Optional) Set the type convergence checking 0: (default) norm(r)/norm(b), or norm(r) when b == 0 1: nomr(r) / norm(r\_0)

HYPRE\_Int **HYPRE\_BoomerAMGGetConvergeType**(HYPRE\_Solver solver, HYPRE\_Int \*type)

HYPRE\_Int **HYPRE\_BoomerAMGSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)  
 (Optional) Set the convergence tolerance, if BoomerAMG is used as a solver.  
 If it is used as a preconditioner, it should be set to 0. The default is 1.e-6.

HYPRE\_Int **HYPRE\_BoomerAMGGetTol**(HYPRE\_Solver solver, HYPRE\_Real \*tol)

HYPRE\_Int **HYPRE\_BoomerAMGSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)  
 (Optional) Sets maximum number of iterations, if BoomerAMG is used as a solver.  
 If it is used as a preconditioner, it should be set to 1. The default is 20.

HYPRE\_Int **HYPRE\_BoomerAMGGetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int \*max\_iter)

HYPRE\_Int **HYPRE\_BoomerAMGSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)  
 (Optional)

HYPRE\_Int **HYPRE\_BoomerAMGSetMaxCoarseSize**(HYPRE\_Solver solver, HYPRE\_Int max\_coarse\_size)  
 (Optional) Sets maximum size of coarsest grid.  
 The default is 9.

HYPRE\_Int **HYPRE\_BoomerAMGGetMaxCoarseSize**(HYPRE\_Solver solver, HYPRE\_Int \*max\_coarse\_size)

HYPRE\_Int **HYPRE\_BoomerAMGSetMinCoarseSize**(HYPRE\_Solver solver, HYPRE\_Int min\_coarse\_size)  
 (Optional) Sets minimum size of coarsest grid.  
 The default is 1.

HYPRE\_Int **HYPRE\_BoomerAMGGetMinCoarseSize**(HYPRE\_Solver solver, HYPRE\_Int \*min\_coarse\_size)

HYPRE\_Int **HYPRE\_BoomerAMGSetMaxLevels**(HYPRE\_Solver solver, HYPRE\_Int max\_levels)  
 (Optional) Sets maximum number of multigrid levels.  
 The default is 25.

HYPRE\_Int **HYPRE\_BoomerAMGGetMaxLevels**(HYPRE\_Solver solver, HYPRE\_Int \*max\_levels)

HYPRE\_Int **HYPRE\_BoomerAMGSetCoarsenCutFactor**(HYPRE\_Solver solver, HYPRE\_Int coarsen\_cut\_factor)  
 (Optional) Sets cut factor for choosing isolated points during coarsening according to the rows' density.  
 The default is 0. If  $nnzrow > coarsen\_cut\_factor * avg\_nnzrow$ , where  $avg\_nnzrow$  is the average number of nonzeros per row of the global matrix, holds for a given row, it is set as fine, and interpolation weights are not computed.

HYPRE\_Int **HYPRE\_BoomerAMGGetCoarsenCutFactor**(HYPRE\_Solver solver, HYPRE\_Int \*coarsen\_cut\_factor)

HYPRE\_Int **HYPRE\_BoomerAMGSetStrongThreshold**(HYPRE\_Solver solver, HYPRE\_Real strong\_threshold)  
 (Optional) Sets AMG strength threshold.  
 The default is 0.25. For 2D Laplace operators, 0.25 is a good value, for 3D Laplace operators, 0.5 or 0.6 is a better value. For elasticity problems, a large strength threshold, such as 0.9, is often better.

HYPRE\_Int **HYPRE\_BoomerAMGGetStrongThreshold**(HYPRE\_Solver solver, HYPRE\_Real  
\*strong\_threshold)

HYPRE\_Int **HYPRE\_BoomerAMGSetStrongThresholdR**(HYPRE\_Solver solver, HYPRE\_Real  
strong\_threshold)

(Optional) The strong threshold for R is strong connections used in building an approximate ideal restriction.

Default value is 0.25.

HYPRE\_Int **HYPRE\_BoomerAMGGetStrongThresholdR**(HYPRE\_Solver solver, HYPRE\_Real  
\*strong\_threshold)

HYPRE\_Int **HYPRE\_BoomerAMGSetFilterThresholdR**(HYPRE\_Solver solver, HYPRE\_Real  
filter\_threshold)

(Optional) The filter threshold for R is used to eliminate small entries of the approximate ideal restriction after building it.

Default value is 0.0, which disables filtering.

HYPRE\_Int **HYPRE\_BoomerAMGGetFilterThresholdR**(HYPRE\_Solver solver, HYPRE\_Real  
\*filter\_threshold)

HYPRE\_Int **HYPRE\_BoomerAMGSetSCommPkgSwitch**(HYPRE\_Solver solver, HYPRE\_Real  
S\_commpkg\_switch)

(Optional) Deprecated.

This routine now has no effect.

HYPRE\_Int **HYPRE\_BoomerAMGSetMaxRowSum**(HYPRE\_Solver solver, HYPRE\_Real max\_row\_sum)

(Optional) Sets a parameter to modify the definition of strength for diagonal dominant portions of the matrix.

The default is 0.9. If *max\_row\_sum* is 1, no checking for diagonally dominant rows is performed.

HYPRE\_Int **HYPRE\_BoomerAMGSetCoarsenType**(HYPRE\_Solver solver, HYPRE\_Int coarsen\_type)

(Optional) Defines which parallel coarsening algorithm is used.

There are the following options for *coarsen\_type*:

- 0 : CLJP-coarsening (a parallel coarsening algorithm using independent sets.
- 1 : classical Ruge-Stueben coarsening on each processor, no boundary treatment (not recommended!)
- 3 : classical Ruge-Stueben coarsening on each processor, followed by a third pass, which adds coarse points on the boundaries
- 6 : Falgout coarsening (uses 1 first, followed by CLJP using the interior coarse points generated by 1 as its first independent set)
- 7 : CLJP-coarsening (using a fixed random vector, for debugging purposes only)
- 8 : PMIS-coarsening (a parallel coarsening algorithm using independent sets, generating lower complexities than CLJP, might also lead to slower convergence)
- 9 : PMIS-coarsening (using a fixed random vector, for debugging purposes only)
- 10 : HMIS-coarsening (uses one pass Ruge-Stueben on each processor independently, followed by PMIS using the interior C-points generated as its first independent set)
- 11 : one-pass Ruge-Stueben coarsening on each processor, no boundary treatment (not recommended!)

- 21 : CGC coarsening by M. Griebel, B. Metsch and A. Schweitzer
- 22 : CGC-E coarsening by M. Griebel, B. Metsch and A.Schweitzer

The default is 10.

HYPRE\_Int **HYPRE\_BoomerAMGGetCoarsenType**(HYPRE\_Solver solver, HYPRE\_Int \*coarsen\_type)

HYPRE\_Int **HYPRE\_BoomerAMGSetNonGalerkinTol**(HYPRE\_Solver solver, HYPRE\_Real nongalerkin\_tol)

(Optional) Defines the non-Galerkin drop-tolerance for sparsifying coarse grid operators and thus reducing communication.

Value specified here is set on all levels. This routine should be used before HYPRE\_BoomerAMGSetLevelNonGalerkinTol, which then can be used to change individual levels if desired

HYPRE\_Int **HYPRE\_BoomerAMGSetLevelNonGalerkinTol**(HYPRE\_Solver solver, HYPRE\_Real nongalerkin\_tol, HYPRE\_Int level)

(Optional) Defines the level specific non-Galerkin drop-tolerances for sparsifying coarse grid operators and thus reducing communication.

A drop-tolerance of 0.0 means to skip doing non-Galerkin on that level. The maximum drop tolerance for a level is 1.0, although much smaller values such as 0.03 or 0.01 are recommended.

Note that if the user wants to set a specific tolerance on all levels, HYPRE\_BoomerAMGSetNonGalerkinTol should be used. Individual levels can then be changed using this routine.

In general, it is safer to drop more aggressively on coarser levels. For instance, one could use 0.0 on the finest level, 0.01 on the second level and then using 0.05 on all remaining levels. The best way to achieve this is to set 0.05 on all levels with HYPRE\_BoomerAMGSetNonGalerkinTol and then change the tolerance on level 0 to 0.0 and the tolerance on level 1 to 0.01 with HYPRE\_BoomerAMGSetLevelNonGalerkinTol. Like many AMG parameters, these drop tolerances can be tuned. It is also common to delay the start of the non-Galerkin process further to a later level than level 1.

#### Parameters

- **solver** – [IN] solver or preconditioner object to be applied.
- **nongalerkin\_tol** – [IN] level specific drop tolerance
- **level** – [IN] level on which drop tolerance is used

HYPRE\_Int **HYPRE\_BoomerAMGSetNonGalerkTol**(HYPRE\_Solver solver, HYPRE\_Int nongalerk\_num\_tol, HYPRE\_Real \*nongalerk\_tol)

(Optional) Defines the non-Galerkin drop-tolerance (old version)

HYPRE\_Int **HYPRE\_BoomerAMGSetMeasureType**(HYPRE\_Solver solver, HYPRE\_Int measure\_type)

(Optional) Defines whether local or global measures are used.

HYPRE\_Int **HYPRE\_BoomerAMGGetMeasureType**(HYPRE\_Solver solver, HYPRE\_Int \*measure\_type)

HYPRE\_Int **HYPRE\_BoomerAMGSetSetupType**(HYPRE\_Solver solver, HYPRE\_Int setup\_type)

HYPRE\_Int **HYPRE\_BoomerAMGSetAggNumLevels**(HYPRE\_Solver solver, HYPRE\_Int agg\_num\_levels)

(Optional) Defines the number of levels of aggressive coarsening.

The default is 0, i.e. no aggressive coarsening.

HYPRE\_Int **HYPRE\_BoomerAMGSetNumPaths**(HYPRE\_Solver solver, HYPRE\_Int num\_paths)

(Optional) Defines the degree of aggressive coarsening.

The default is 1. Larger numbers lead to less aggressive coarsening.

HYPRE\_Int **HYPRE\_BoomerAMGSetCGCIts**(HYPRE\_Solver solver, HYPRE\_Int its)

(optional) Defines the number of pathes for CGC-coarsening.

HYPRE\_Int **HYPRE\_BoomerAMGSetNodal**(HYPRE\_Solver solver, HYPRE\_Int nodal)

(Optional) Sets whether to use the nodal systems coarsening.

Should be used for linear systems generated from systems of PDEs. The default is 0 (unknown-based coarsening, only coarsens within same function). For the remaining options a nodal matrix is generated by applying a norm to the nodal blocks and applying the coarsening algorithm to this matrix.

- 1 : Frobenius norm
- 2 : sum of absolute values of elements in each block
- 3 : largest element in each block (not absolute value)
- 4 : row-sum norm
- 6 : sum of all values in each block

HYPRE\_Int **HYPRE\_BoomerAMGSetNodalDiag**(HYPRE\_Solver solver, HYPRE\_Int nodal\_diag)

(Optional) Sets whether to give special treatment to diagonal elements in the nodal systems version.

The default is 0. If set to 1, the diagonal entry is set to the negative sum of all off diagonal entries. If set to 2, the signs of all diagonal entries are inverted.

HYPRE\_Int **HYPRE\_BoomerAMGSetNodalLevels**(HYPRE\_Solver solver, HYPRE\_Int nodal\_levels)

(Optional) Sets the number of levels on which nodal coarsening should be performed (nodal coarsening requires that numFunctions be set to something other than 1).

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetKeepSameSign**(HYPRE\_Solver solver, HYPRE\_Int keep\_same\_sign)

(Optional) Sets whether to keep same sign in S for nodal > 0 The default is 0, i.e., discard those elements.

HYPRE\_Int **HYPRE\_BoomerAMGSetInterpType**(HYPRE\_Solver solver, HYPRE\_Int interp\_type)

(Optional) Defines which parallel interpolation operator is used.

There are the following options for *interp\_type*:

- 0 : classical modified interpolation
- 1 : LS interpolation (for use with GSMG)
- 2 : classical modified interpolation for hyperbolic PDEs
- 3 : direct interpolation (with separation of weights) (also for GPU use)
- 4 : multipass interpolation
- 5 : multipass interpolation (with separation of weights)
- 6 : extended+i interpolation (also for GPU use)
- 7 : extended+i (if no common C neighbor) interpolation
- 8 : standard interpolation
- 9 : standard interpolation (with separation of weights)

- 10 : classical block interpolation (for use with nodal systems version only)
- 11 : classical block interpolation (for use with nodal systems version only) with diagonalized diagonal blocks
- 12 : FF interpolation
- 13 : FF1 interpolation
- 14 : extended interpolation (also for GPU use)
- 15 : interpolation with adaptive weights (GPU use only)
- 16 : extended interpolation in matrix-matrix form
- 17 : extended+i interpolation in matrix-matrix form
- 18 : extended+e interpolation in matrix-matrix form

The default is ext+i interpolation (interp\_type 6) truncated to at most 4 elements per row. (see HYPRE\_BoomerAMGSetPMaxElmts).

HYPRE\_Int **HYPRE\_BoomerAMGSetInterpRefine**(HYPRE\_Solver solver, HYPRE\_Int num\_refine)

(Optional) Sets the number of Jacobi interpolation improvement steps.

Each improvement step smooths the interpolation matrix and generally reduces the operator complexity. The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetTruncFactor**(HYPRE\_Solver solver, HYPRE\_Real trunc\_factor)

(Optional) Defines a truncation factor for the interpolation.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGGetTruncFactor**(HYPRE\_Solver solver, HYPRE\_Real \*trunc\_factor)

HYPRE\_Int **HYPRE\_BoomerAMGSetPMaxElmts**(HYPRE\_Solver solver, HYPRE\_Int P\_max\_elmts)

(Optional) Defines the maximal number of elements per row for the interpolation.

The default is 4. To turn off truncation, it needs to be set to 0.

HYPRE\_Int **HYPRE\_BoomerAMGGetPMaxElmts**(HYPRE\_Solver solver, HYPRE\_Int \*P\_max\_elmts)

HYPRE\_Int **HYPRE\_BoomerAMGSetSepWeight**(HYPRE\_Solver solver, HYPRE\_Int sep\_weight)

(Optional) Defines whether separation of weights is used when defining strength for standard interpolation or multipass interpolation.

Default: 0, i.e. no separation of weights used.

HYPRE\_Int **HYPRE\_BoomerAMGSetAggInterpType**(HYPRE\_Solver solver, HYPRE\_Int agg\_interp\_type)

(Optional) Defines the interpolation used on levels of aggressive coarsening. The default is 4, i.e.

multipass interpolation. The following options exist:

- 1 : 2-stage extended+i
- 2 : 2-stage standard
- 3 : 2-stage extended
- 4 : multipass (default)
- 5 : 2-stage extended in matrix-matrix form
- 6 : 2-stage extended+i in matrix-matrix form

- 7 : 2-stage extended+e in matrix-matrix form
- 8 : multipass in matrix-matrix form

HYPRE\_Int **HYPRE\_BoomerAMGSetAggTruncFactor**(HYPRE\_Solver solver, HYPRE\_Real  
agg\_trunc\_factor)

(Optional) Defines the truncation factor for the interpolation used for aggressive coarsening.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetAggP12TruncFactor**(HYPRE\_Solver solver, HYPRE\_Real  
agg\_P12\_trunc\_factor)

(Optional) Defines the truncation factor for the matrices P1 and P2 which are used to build 2-stage interpolation.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetAggPMaxElmts**(HYPRE\_Solver solver, HYPRE\_Int agg\_P\_max\_elmts)

(Optional) Defines the maximal number of elements per row for the interpolation used for aggressive coarsening.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetAggP12MaxElmts**(HYPRE\_Solver solver, HYPRE\_Int  
agg\_P12\_max\_elmts)

(Optional) Defines the maximal number of elements per row for the matrices P1 and P2 which are used to build 2-stage interpolation.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetInterpVectors**(HYPRE\_Solver solver, HYPRE\_Int num\_vectors,  
*HYPRE\_ParVector* \*interp\_vectors)

(Optional) Allows the user to incorporate additional vectors into the interpolation for systems AMG, e.g. rigid body modes for linear elasticity problems. This can only be used in context with nodal coarsening and still requires the user to choose an interpolation.

HYPRE\_Int **HYPRE\_BoomerAMGSetInterpVecVariant**(HYPRE\_Solver solver, HYPRE\_Int var)

(Optional) Defines the interpolation variant used for HYPRE\_BoomerAMGSetInterpVectors:

- 1 : GM approach 1
- 2 : GM approach 2 (to be preferred over 1)
- 3 : LN approach

HYPRE\_Int **HYPRE\_BoomerAMGSetSmoothInterpVectors**(HYPRE\_Solver solver, HYPRE\_Int  
smooth\_interp\_vectors)

(Optional) Controls whether to apply smoothing to the interpolation vectors used in GMG interpolation.

Set to 1 to enable smoothing, 0 to disable. The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetInterpVecQMax**(HYPRE\_Solver solver, HYPRE\_Int q\_max)

(Optional) Defines the maximal elements per row for Q, the additional columns added to the original interpolation matrix P, to reduce complexity.

The default is no truncation.

HYPRE\_Int **HYPRE\_BoomerAMGSetInterpVecAbsQTrunc**(HYPRE\_Solver solver, HYPRE\_Real q\_trunc)  
 (Optional) Defines a truncation factor for Q, the additional columns added to the original interpolation matrix P, to reduce complexity.

The default is no truncation.

HYPRE\_Int **HYPRE\_BoomerAMGSetGSMG**(HYPRE\_Solver solver, HYPRE\_Int gsmg)  
 (Optional) Specifies the use of GSMG - geometrically smooth coarsening and interpolation.

Currently any nonzero value for gsmg will lead to the use of GSMG. The default is 0, i.e. (GSMG is not used)

HYPRE\_Int **HYPRE\_BoomerAMGSetNumSamples**(HYPRE\_Solver solver, HYPRE\_Int num\_samples)  
 (Optional) Defines the number of sample vectors used in GSMG or LS interpolation.

HYPRE\_Int **HYPRE\_BoomerAMGSetCycleType**(HYPRE\_Solver solver, HYPRE\_Int cycle\_type)  
 (Optional) Defines the type of cycle.

For a V-cycle, set *cycle\_type* to 1, for a W-cycle set *cycle\_type* to 2. The default is 1.

HYPRE\_Int **HYPRE\_BoomerAMGGetCycleType**(HYPRE\_Solver solver, HYPRE\_Int \*cycle\_type)

HYPRE\_Int **HYPRE\_BoomerAMGSetFCycle**(HYPRE\_Solver solver, HYPRE\_Int fcycle)  
 (Optional) Specifies the use of Full multigrid cycle.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGGetFCycle**(HYPRE\_Solver solver, HYPRE\_Int \*fcycle)

HYPRE\_Int **HYPRE\_BoomerAMGSetAdditive**(HYPRE\_Solver solver, HYPRE\_Int addlvl)  
 (Optional) Defines use of an additive V(1,1)-cycle using the classical additive method starting at level 'addlvl'.

The multiplicative approach is used on levels 0, ... 'addlvl+1'. 'addlvl' needs to be > -1 for this to have an effect. Can only be used with weighted Jacobi and I1-Jacobi(default).

Can only be used when AMG is used as a preconditioner !!!

HYPRE\_Int **HYPRE\_BoomerAMGGetAdditive**(HYPRE\_Solver solver, HYPRE\_Int \*additive)

HYPRE\_Int **HYPRE\_BoomerAMGSetMultAdditive**(HYPRE\_Solver solver, HYPRE\_Int addlvl)  
 (Optional) Defines use of an additive V(1,1)-cycle using the mult-additive method starting at level 'addlvl'.

The multiplicative approach is used on levels 0, ... 'addlvl+1'. 'addlvl' needs to be > -1 for this to have an effect. Can only be used with weighted Jacobi and I1-Jacobi(default).

Can only be used when AMG is used as a preconditioner !!!

HYPRE\_Int **HYPRE\_BoomerAMGGetMultAdditive**(HYPRE\_Solver solver, HYPRE\_Int \*mult\_additive)

HYPRE\_Int **HYPRE\_BoomerAMGSetSimple**(HYPRE\_Solver solver, HYPRE\_Int addlvl)  
 (Optional) Defines use of an additive V(1,1)-cycle using the simplified mult-additive method starting at level 'addlvl'.

The multiplicative approach is used on levels 0, ... 'addlvl+1'. 'addlvl' needs to be > -1 for this to have an effect. Can only be used with weighted Jacobi and I1-Jacobi(default).

Can only be used when AMG is used as a preconditioner !!!

HYPRE\_Int **HYPRE\_BoomerAMGGetSimple**(HYPRE\_Solver solver, HYPRE\_Int \*simple)

HYPRE\_Int **HYPRE\_BoomerAMGSetAddLastLvl**(HYPRE\_Solver solver, HYPRE\_Int add\_last\_lvl)

(Optional) Defines last level where additive, mult-additive or simple cycle is used.

The multiplicative approach is used on levels > add\_last\_lvl.

Can only be used when AMG is used as a preconditioner !!!

HYPRE\_Int **HYPRE\_BoomerAMGSetMultAddTruncFactor**(HYPRE\_Solver solver, HYPRE\_Real add\_trunc\_factor)

(Optional) Defines the truncation factor for the smoothed interpolation used for mult-additive or simple method.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetMultAddPMaxElmts**(HYPRE\_Solver solver, HYPRE\_Int add\_P\_max\_elmts)

(Optional) Defines the maximal number of elements per row for the smoothed interpolation used for mult-additive or simple method.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetAddRelaxType**(HYPRE\_Solver solver, HYPRE\_Int add\_rlx\_type)

(Optional) Defines the relaxation type used in the (mult)additive cycle portion (also affects simple method.)  
The default is 18 (L1-Jacobi).

Currently the only other option allowed is 0 (Jacobi) which should be used in combination with HYPRE\_BoomerAMGSetAddRelaxWt.

HYPRE\_Int **HYPRE\_BoomerAMGSetAddPMaxElmts**(HYPRE\_Solver solver, HYPRE\_Int add\_P\_max\_elmts)

(Optional) Sets the maximum number of elements per row for additive interpolation matrices.

This controls sparsity of P in additive cycles. The default is 0 (no limit).

HYPRE\_Int **HYPRE\_BoomerAMGSetAddTruncFactor**(HYPRE\_Solver solver, HYPRE\_Real add\_trunc\_factor)

(Optional) Sets the truncation factor for additive interpolation.

Elements in P smaller than (truncation factor \* max element in row) are dropped. The default is 0.0 (no truncation).

HYPRE\_Int **HYPRE\_BoomerAMGSetAddRelaxWt**(HYPRE\_Solver solver, HYPRE\_Real add\_rlx\_wt)

(Optional) Defines the relaxation weight used for Jacobi within the (mult)additive or simple cycle portion.

The default is 1. The weight only affects the Jacobi method, and has no effect on L1-Jacobi

HYPRE\_Int **HYPRE\_BoomerAMGSetSeqThreshold**(HYPRE\_Solver solver, HYPRE\_Int seq\_threshold)

(Optional) Sets maximal size for agglomeration or redundant coarse grid solve.

When the system is smaller than this threshold, sequential AMG is used on process 0 or on all remaining active processes (if redundant = 1).

HYPRE\_Int **HYPRE\_BoomerAMGGetSeqThreshold**(HYPRE\_Solver solver, HYPRE\_Int \*seq\_threshold)

HYPRE\_Int **HYPRE\_BoomerAMGSetRedundant**(HYPRE\_Solver solver, HYPRE\_Int redundant)

(Optional) operates switch for redundancy.

Needs to be used with HYPRE\_BoomerAMGSetSeqThreshold. Default is 0, i.e. no redundancy.

HYPRE\_Int **HYPRE\_BoomerAMGGetRedundant**(HYPRE\_Solver solver, HYPRE\_Int \*redundant)

HYPRE\_Int **HYPRE\_BoomerAMGSetNumGridSweeps**(HYPRE\_Solver solver, HYPRE\_Int \*num\_grid\_sweeps)

(Optional) Defines the number of sweeps for the fine and coarse grid, the up and down cycle.

Note: This routine will be phased out!!!! Use HYPRE\_BoomerAMGSetNumSweeps or HYPRE\_BoomerAMGSetCycleNumSweeps instead.

HYPRE\_Int **HYPRE\_BoomerAMGSetNumSweeps**(HYPRE\_Solver solver, HYPRE\_Int num\_sweeps)

(Optional) Sets the number of sweeps.

On the finest level, the up and the down cycle the number of sweeps are set to *num\_sweeps* and on the coarsest level to 1. The default is 1.

HYPRE\_Int **HYPRE\_BoomerAMGSetCycleNumSweeps**(HYPRE\_Solver solver, HYPRE\_Int num\_sweeps, HYPRE\_Int k)

(Optional) Sets the number of sweeps at a specified cycle.

There are the following options for *k*:

- 1 : the down cycle
- 2 : the up cycle
- 3 : the coarsest level

HYPRE\_Int **HYPRE\_BoomerAMGGetCycleNumSweeps**(HYPRE\_Solver solver, HYPRE\_Int \*num\_sweeps, HYPRE\_Int k)

HYPRE\_Int **HYPRE\_BoomerAMGSetGridRelaxType**(HYPRE\_Solver solver, HYPRE\_Int \*grid\_relax\_type)

(Optional) Defines which smoother is used on the fine and coarse grid, the up and down cycle.

Note: This routine will be phased out!!!! Use HYPRE\_BoomerAMGSetRelaxType or HYPRE\_BoomerAMGSetCycleRelaxType instead.

HYPRE\_Int **HYPRE\_BoomerAMGSetRelaxType**(HYPRE\_Solver solver, HYPRE\_Int relax\_type)

(Optional) Defines the smoother to be used.

It uses the given smoother on the fine grid, the up and the down cycle and sets the solver on the coarsest level to Gaussian elimination (9). The default is  $\ell_1$ -Gauss-Seidel, forward solve (13) on the down cycle and backward solve (14) on the up cycle.

There are the following options for *relax\_type*:

- 0 : Jacobi
- 1 : Gauss-Seidel, sequential (very slow!)
- 2 : Gauss-Seidel, interior points in parallel, boundary sequential (slow!)
- 3 : hybrid Gauss-Seidel or SOR, forward solve
- 4 : hybrid Gauss-Seidel or SOR, backward solve
- 5 : hybrid chaotic Gauss-Seidel (works only with OpenMP)
- 6 : hybrid symmetric Gauss-Seidel or SSOR
- 7 : Jacobi (uses Matvec)
- 8 :  $\ell_1$ -scaled hybrid symmetric Gauss-Seidel

- 9 : Gaussian elimination (only on coarsest level)
- 10 : On-processor direct forward solve for matrices with triangular structure
- 11 : Two Stage approximation to GS. Uses the strict lower part of the diagonal matrix
- 12 : Two Stage approximation to GS. Uses the strict lower part of the diagonal matrix and a second iteration for additional error approximation
- 13 :  $\ell_1$  Gauss-Seidel, forward solve
- 14 :  $\ell_1$  Gauss-Seidel, backward solve
- 15 : CG (warning - not a fixed smoother - may require FGMRES)
- 16 : Chebyshev
- 17 : FCF-Jacobi
- 18 :  $\ell_1$ -scaled jacobi
- 19 : Gaussian elimination (old version)
- 21 : The same as 8 except forcing serialization on CPU (#OMP-thread = 1)
- 29 : Direct solve: use Gaussian elimination & BLAS (with pivoting) (old version)
- 30 : Kaczmarz
- 88: The same methods as 8 with a convergent 11-term
- 89: Symmetric 11-hybrid Gauss-Seidel (i.e., 13 followed by 14)
- 98 : LU with pivoting
- 99 : LU with pivoting -199 : Matvec with the inverse

HYPRE\_Int HYPRE\_BoomerAMGSetCycleRelaxType(HYPRE\_Solver solver, HYPRE\_Int relax\_type, HYPRE\_Int k)

(Optional) Defines the smoother at a given cycle.

For options of *relax\_type* see description of HYPRE\_BoomerAMGSetRelaxType. In addition, the following options for *relax\_type* are available when choosing the coarsest level solver (k = 3):

For coarsest level systems formed via a sub-communicator defined with active ranks:

- 9 : hypr's internal Gaussian elimination (host only).
- 99 : LU factorization with pivoting.
- 199 : explicit (dense) inverse.

For coarsest level systems formed via hypr\_DataExchangeList:

- 19 : hypr's internal Gaussian elimination (host only).
- 98 : LU factorization with pivoting.
- 198 : explicit (dense) inverse.

Options for *k* are

- 1 : the down cycle
- 2 : the up cycle
- 3 : the coarsest level

HYPRE\_Int **HYPRE\_BoomerAMGGetCycleRelaxType**(HYPRE\_Solver solver, HYPRE\_Int \*relax\_type, HYPRE\_Int k)

HYPRE\_Int **HYPRE\_BoomerAMGSetRelaxOrder**(HYPRE\_Solver solver, HYPRE\_Int relax\_order)

(Optional) Defines in which order the points are relaxed.

There are the following options for *relax\_order*:

- 0 : the points are relaxed in natural or lexicographic order on each processor
- 1 : CF-relaxation is used, i.e on the fine grid and the down cycle the coarse points are relaxed first, followed by the fine points; on the up cycle the F-points are relaxed first, followed by the C-points. On the coarsest level, if an iterative scheme is used, the points are relaxed in lexicographic order.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetGridRelaxPoints**(HYPRE\_Solver solver, HYPRE\_Int \*\*grid\_relax\_points)

(Optional) Defines in which order the points are relaxed.

See also HYPRE\_BoomerAMGSetRelaxOrder.

HYPRE\_Int **HYPRE\_BoomerAMGSetRelaxWeight**(HYPRE\_Solver solver, HYPRE\_Real \*relax\_weight)

(Optional) Defines the relaxation weight for smoothed Jacobi and hybrid SOR.

Note: This routine will be phased out!!!! Use HYPRE\_BoomerAMGSetRelaxWt or HYPRE\_BoomerAMGSetLevelRelaxWt instead.

HYPRE\_Int **HYPRE\_BoomerAMGSetRelaxWt**(HYPRE\_Solver solver, HYPRE\_Real relax\_weight)

(Optional) Defines the relaxation weight for smoothed Jacobi and hybrid SOR on all levels.

Values for *relax\_weight* are

- > 0 : this assigns the given relaxation weight on all levels
- = 0 : the weight is determined on each level with the estimate  $\frac{3}{4\|D^{-1/2}AD^{-1/2}\|}$ , where  $D$  is the diagonal of  $A$  (this should only be used with Jacobi)
- = -k : the relaxation weight is determined with at most k CG steps on each level (this should only be used for symmetric positive definite problems)

The default is 1.

HYPRE\_Int **HYPRE\_BoomerAMGSetLevelRelaxWt**(HYPRE\_Solver solver, HYPRE\_Real relax\_weight, HYPRE\_Int level)

(Optional) Defines the relaxation weight for smoothed Jacobi and hybrid SOR on the user defined level.

Note that the finest level is denoted 0, the next coarser level 1, etc. For nonpositive *relax\_weight*, the parameter is determined on the given level as described for HYPRE\_BoomerAMGSetRelaxWt. The default is 1.

HYPRE\_Int **HYPRE\_BoomerAMGSetOmega**(HYPRE\_Solver solver, HYPRE\_Real \*omega)

(Optional) Defines the outer relaxation weight for hybrid SOR.

Note: This routine will be phased out!!!! Use HYPRE\_BoomerAMGSetOuterWt or HYPRE\_BoomerAMGSetLevelOuterWt instead.

HYPRE\_Int **HYPRE\_BoomerAMGSetOuterWt**(HYPRE\_Solver solver, HYPRE\_Real omega)

(Optional) Defines the outer relaxation weight for hybrid SOR and SSOR on all levels.

Values for *omega* are

- $> 0$  : this assigns the same outer relaxation weight  $\omega$  on each level
- $= -k$  : an outer relaxation weight is determined with at most  $k$  CG steps on each level (this only makes sense for symmetric positive definite problems and smoothers such as SSOR)

The default is 1.

HYPRE\_Int **HYPRE\_BoomerAMGSetLevelOuterWt**(HYPRE\_Solver solver, HYPRE\_Real  $\omega$ , HYPRE\_Int level)

(Optional) Defines the outer relaxation weight for hybrid SOR or SSOR on the user defined level.

Note that the finest level is denoted 0, the next coarser level 1, etc. For nonpositive  $\omega$ , the parameter is determined on the given level as described for HYPRE\_BoomerAMGSetOuterWt. The default is 1.

HYPRE\_Int **HYPRE\_BoomerAMGSetChebyOrder**(HYPRE\_Solver solver, HYPRE\_Int order)

(Optional) See *HYPRE\_ParChebySetOrder*

HYPRE\_Int **HYPRE\_BoomerAMGSetChebyFraction**(HYPRE\_Solver solver, HYPRE\_Real ratio)

(Optional) See *HYPRE\_ParChebySetEigRatio*

HYPRE\_Int **HYPRE\_BoomerAMGSetChebyScale**(HYPRE\_Solver solver, HYPRE\_Int scale)

(Optional) See *HYPRE\_ParChebySetScale*

HYPRE\_Int **HYPRE\_BoomerAMGSetChebyVariant**(HYPRE\_Solver solver, HYPRE\_Int variant)

(Optional) See *HYPRE\_ParChebySetVariant*

HYPRE\_Int **HYPRE\_BoomerAMGSetChebyEigEst**(HYPRE\_Solver solver, HYPRE\_Int eig\_est)

(Optional) See *HYPRE\_ParChebySetEigEst*

HYPRE\_Int **HYPRE\_BoomerAMGSetSmoothType**(HYPRE\_Solver solver, HYPRE\_Int smooth\_type)

(Optional) Enables the use of more complex smoothers.

The following options exist for *smooth\_type*:

- 4 : FSAI (routines needed to set: HYPRE\_BoomerAMGSetFSAIMaxSteps, HYPRE\_BoomerAMGSetFSAIMaxStepSize, HYPRE\_BoomerAMGSetFSAIEigMaxIters, HYPRE\_BoomerAMGSetFSAIKapTolerance)
- 5 : ParILUK (routines needed to set: HYPRE\_ILUSetLevelOfFill, HYPRE\_ILUSetType)
- 6 : Schwarz (routines needed to set: HYPRE\_BoomerAMGSetDomainType, HYPRE\_BoomerAMGSetOverlap, HYPRE\_BoomerAMGSetVariant, HYPRE\_BoomerAMGSetSchwarzRlxWeight)
- 7 : Pilut (routines needed to set: HYPRE\_BoomerAMGSetDropTol, HYPRE\_BoomerAMGSetMaxNzPerRow)
- 8 : ParaSails (routines needed to set: HYPRE\_BoomerAMGSetSym, HYPRE\_BoomerAMGSetLevel, HYPRE\_BoomerAMGSetFilter, HYPRE\_BoomerAMGSetThreshold)
- 9 : Euclid (routines needed to set: HYPRE\_BoomerAMGSetEuclidFile)

The default is 6. Also, if no smoother parameters are set via the routines mentioned in the table above, default values are used.

HYPRE\_Int **HYPRE\_BoomerAMGGetSmoothType**(HYPRE\_Solver solver, HYPRE\_Int \*smooth\_type)

HYPRE\_Int **HYPRE\_BoomerAMGSetSmoothNumLevels**(HYPRE\_Solver solver, HYPRE\_Int smooth\_num\_levels)

(Optional) Sets the number of levels for more complex smoothers.

The smoothers, as defined by HYPRE\_BoomerAMGSetSmoothType, will be used on level 0 (the finest level) through level *smooth\_num\_levels-1*. The default is 0, i.e. no complex smoothers are used.

HYPRE\_Int **HYPRE\_BoomerAMGGetSmoothNumLevels**(HYPRE\_Solver solver, HYPRE\_Int \*smooth\_num\_levels)

HYPRE\_Int **HYPRE\_BoomerAMGSetSmoothNumSweeps**(HYPRE\_Solver solver, HYPRE\_Int smooth\_num\_sweeps)

(Optional) Sets the number of sweeps for more complex smoothers.

The default is 1.

HYPRE\_Int **HYPRE\_BoomerAMGGetSmoothNumSweeps**(HYPRE\_Solver solver, HYPRE\_Int \*smooth\_num\_sweeps)

HYPRE\_Int **HYPRE\_BoomerAMGSetFlexibleCycleStruct**(HYPRE\_Solver solver, HYPRE\_Int \*cycle\_struct\_flexible)

(Optional) Defines a flexible cycle structure for the BoomerAMG cycle.

Cycle structure is defined by an array of integers with values:

-2 : terminate the cycle -1 : move to the next coarser level 0 : remain on the current level 1 : move to the next finer level

HYPRE\_Int **HYPRE\_BoomerAMGSetFlexibleRelaxTypes**(HYPRE\_Solver solver, HYPRE\_Int \*relax\_types\_flexible)

(Optional) Defines the relax orders (lexicographical or CF) used during a flexible cycle.

Different relax orders may be used at each point in the cycle. This must be used in conjunction with a cycle structure defined by HYPRE\_BoomerAMGSetFlexibleRelaxTypes. The array of relax orders passed must have the same length as the array defining the cycle structure. See also HYPRE\_BoomerAMGSetRelaxOrder.

HYPRE\_Int **HYPRE\_BoomerAMGSetFlexibleRelaxOrders**(HYPRE\_Solver solver, HYPRE\_Int \*relax\_orders\_flexible)

(Optional) Defines the relax orders (lexicographical or CF) used during a flexible cycle.

Different relax orders may be used at each point in the cycle. This must be used in conjunction with a cycle structure defined by HYPRE\_BoomerAMGSetFlexibleRelaxTypes. The array of relax orders passed must have the same length as the array defining the cycle structure. See also HYPRE\_BoomerAMGSetRelaxOrder.

HYPRE\_Int **HYPRE\_BoomerAMGSetFlexibleRelaxWeights**(HYPRE\_Solver solver, HYPRE\_Real \*relax\_weights\_flexible)

(Optional) Defines the relaxation weights used during a flexible cycle.

Different relaxation weights may be used at each point in the cycle. This must be used in conjunction with a cycle structure defined by HYPRE\_BoomerAMGSetFlexibleRelaxTypes. The array of relax weights passed must have the same length as the array defining the cycle structure. See also HYPRE\_BoomerAMGSetRelaxWeight.

HYPRE\_Int **HYPRE\_BoomerAMGSetFlexibleOuterWeights**(HYPRE\_Solver solver, HYPRE\_Real \*outer\_weights\_flexible)

(Optional) Defines the outer relaxation weights used during a flexible cycle.

Different outer weights may be used at each point in the cycle. This must be used in conjunction with a cycle structure defined by `HYPRE_BoomerAMGSetFlexibleRelaxTypes`. The array of outer weights passed must have the same length as the array defining the cycle structure. See also `HYPRE_BoomerAMGSetOuterWt`.

`HYPRE_Int HYPRE_BoomerAMGSetFlexibleCGCScalingFactors`(`HYPRE_Solver solver`, `HYPRE_Real *cgc_scaling_factors_flexible`)

(Optional) Defines the a scaling factor for the coarse-grid correction during a flexible cycle.

Different scaling factors may be used at each point in the cycle. This must be used in conjunction with a cycle structure defined by `HYPRE_BoomerAMGSetFlexibleRelaxTypes`. The array of coarse- grid correction scaling factors passed must have the same length as the array defining the cycle structure. Note, however, that the scaling factors will only be used when the flexible cycle performs a coarse-grid correction, i.e. when the value of the array passed for the cycle structure is 1.

`HYPRE_Int HYPRE_BoomerAMGSetVariant`(`HYPRE_Solver solver`, `HYPRE_Int variant`)

(Optional) Defines which variant of the Schwarz method is used.

The following options exist for *variant*:

- 0 : hybrid multiplicative Schwarz method (no overlap across processor boundaries)
- 1 : hybrid additive Schwarz method (no overlap across processor boundaries)
- 2 : additive Schwarz method
- 3 : hybrid multiplicative Schwarz method (with overlap across processor boundaries)

The default is 0.

`HYPRE_Int HYPRE_BoomerAMGGetVariant`(`HYPRE_Solver solver`, `HYPRE_Int *variant`)

`HYPRE_Int HYPRE_BoomerAMGSetOverlap`(`HYPRE_Solver solver`, `HYPRE_Int overlap`)

(Optional) Defines the overlap for the Schwarz method.

The following options exist for overlap:

- 0 : no overlap
- 1 : minimal overlap (default)
- 2 : overlap generated by including all neighbors of domain boundaries

`HYPRE_Int HYPRE_BoomerAMGGetOverlap`(`HYPRE_Solver solver`, `HYPRE_Int *overlap`)

`HYPRE_Int HYPRE_BoomerAMGSetDomainType`(`HYPRE_Solver solver`, `HYPRE_Int domain_type`)

(Optional) Defines the type of domain used for the Schwarz method.

The following options exist for *domain\_type*:

- 0 : each point is a domain
- 1 : each node is a domain (only of interest in “systems” AMG)
- 2 : each domain is generated by agglomeration (default)

`HYPRE_Int HYPRE_BoomerAMGGetDomainType`(`HYPRE_Solver solver`, `HYPRE_Int *domain_type`)

HYPRE\_Int **HYPRE\_BoomerAMGSetSchwarzRlxWeight**(HYPRE\_Solver solver, HYPRE\_Real schwarz\_rlx\_weight)

(Optional) Defines a smoothing parameter for the additive Schwarz method.

HYPRE\_Int **HYPRE\_BoomerAMGGetSchwarzRlxWeight**(HYPRE\_Solver solver, HYPRE\_Real \*schwarz\_rlx\_weight)

HYPRE\_Int **HYPRE\_BoomerAMGSetSchwarzUseNonSymm**(HYPRE\_Solver solver, HYPRE\_Int use\_nonsymm)

(Optional) Indicates that the aggregates may not be SPD for the Schwarz method.

The following options exist for *use\_nonsymm*:

- 0 : assume SPD (default)
- 1 : assume non-symmetric

HYPRE\_Int **HYPRE\_BoomerAMGSetSym**(HYPRE\_Solver solver, HYPRE\_Int sym)

(Optional) Defines symmetry for ParaSAILS.

For further explanation see description of ParaSAILS.

HYPRE\_Int **HYPRE\_BoomerAMGSetLevel**(HYPRE\_Solver solver, HYPRE\_Int level)

(Optional) Defines number of levels for ParaSAILS.

For further explanation see description of ParaSAILS.

HYPRE\_Int **HYPRE\_BoomerAMGSetThreshold**(HYPRE\_Solver solver, HYPRE\_Real threshold)

(Optional) Defines threshold for ParaSAILS.

For further explanation see description of ParaSAILS.

HYPRE\_Int **HYPRE\_BoomerAMGSetFilter**(HYPRE\_Solver solver, HYPRE\_Real filter)

(Optional) Defines filter for ParaSAILS.

For further explanation see description of ParaSAILS.

HYPRE\_Int **HYPRE\_BoomerAMGSetDropTol**(HYPRE\_Solver solver, HYPRE\_Real drop\_tol)

(Optional) Defines drop tolerance for PILUT.

For further explanation see description of PILUT.

HYPRE\_Int **HYPRE\_BoomerAMGSetMaxNzPerRow**(HYPRE\_Solver solver, HYPRE\_Int max\_nz\_per\_row)

(Optional) Defines maximal number of nonzeros for PILUT.

For further explanation see description of PILUT.

HYPRE\_Int **HYPRE\_BoomerAMGSetEuclidFile**(HYPRE\_Solver solver, char \*euclidfile)

(Optional) Defines name of an input file for Euclid parameters.

For further explanation see description of Euclid.

HYPRE\_Int **HYPRE\_BoomerAMGSetEuLevel**(HYPRE\_Solver solver, HYPRE\_Int eu\_level)

(Optional) Defines number of levels for ILU(k) in Euclid.

For further explanation see description of Euclid.

HYPRE\_Int **HYPRE\_BoomerAMGSetEuSparseA**(HYPRE\_Solver solver, HYPRE\_Real eu\_sparse\_A)

(Optional) Defines filter for ILU(k) for Euclid.

For further explanation see description of Euclid.

- HYPRE\_Int **HYPRE\_BoomerAMGSetEuBJ**(HYPRE\_Solver solver, HYPRE\_Int eu\_bj)  
 (Optional) Defines use of block jacobi ILUT for Euclid.  
 For further explanation see description of Euclid.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILUType**(HYPRE\_Solver solver, HYPRE\_Int ilu\_type)  
 Defines type of ILU smoother to use For further explanation see description of ILU.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILULevel**(HYPRE\_Solver solver, HYPRE\_Int ilu\_lfil)  
 Defines level k for ILU(k) smoother For further explanation see description of ILU.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILUMaxRowNnz**(HYPRE\_Solver solver, HYPRE\_Int ilu\_max\_row\_nnz)  
 Defines max row nonzeros for ILUT smoother For further explanation see description of ILU.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILUMaxIter**(HYPRE\_Solver solver, HYPRE\_Int ilu\_max\_iter)  
 Defines number of iterations for ILU smoother on each level For further explanation see description of ILU.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILUDroptol**(HYPRE\_Solver solver, HYPRE\_Real ilu\_droptol)  
 Defines drop tolerance for iLUT smoother For further explanation see description of ILU.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILUTriSolve**(HYPRE\_Solver solver, HYPRE\_Int ilu\_tri\_solve)  
 (Optional) Defines triangular solver for ILU(k,T) smoother: 0-iterative, 1-direct (default) For further explanation see description of ILU.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILULowerJacobiIters**(HYPRE\_Solver solver, HYPRE\_Int ilu\_lower\_jacobi\_iters)  
 (Optional) Defines number of lower Jacobi iterations for ILU(k,T) smoother triangular solve.  
 For further explanation see description of ILU.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILUUpperJacobiIters**(HYPRE\_Solver solver, HYPRE\_Int ilu\_upper\_jacobi\_iters)  
 (Optional) Defines number of upper Jacobi iterations for ILU(k,T) smoother triangular solve.  
 For further explanation see description of ILU.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILULocalReordering**(HYPRE\_Solver solver, HYPRE\_Int ilu\_reordering\_type)  
 (Optional) Set Local Reordering paramter (1==RCM, 0==None) For further explanation see description of ILU.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILUIterSetupType**(HYPRE\_Solver solver, HYPRE\_Int ilu\_iter\_setup\_type)  
 (Optional) Set iterative ILU's algorithm type.  
 For further explanation see *HYPRE\_ILUSetIterativeSetupType*.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILUIterSetupOption**(HYPRE\_Solver solver, HYPRE\_Int ilu\_iter\_setup\_option)  
 (Optional) Set iterative ILU's option.  
 For further explanation see *HYPRE\_ILUSetIterativeSetupOption*.
- HYPRE\_Int **HYPRE\_BoomerAMGSetILUIterSetupMaxIter**(HYPRE\_Solver solver, HYPRE\_Int ilu\_iter\_setup\_max\_iter)  
 (Optional) Set iterative ILU's max.  
 number of iterations. For further explanation see *HYPRE\_ILUSetIterativeSetupMaxIter*.

HYPRE\_Int **HYPRE\_BoomerAMGSetILUIterSetupTolerance**(HYPRE\_Solver solver, HYPRE\_Real ilu\_iter\_setup\_tolerance)

(Optional) Set iterative ILU's tolerance.

For further explanation see *HYPRE\_ILUSetIterativeSetupTolerance*.

HYPRE\_Int **HYPRE\_BoomerAMGSetFSAIAlgoType**(HYPRE\_Solver solver, HYPRE\_Int algo\_type)

(Optional) Defines the algorithm type for setting up FSAI For further explanation see *HYPRE\_FSAISetAlgoType*.

HYPRE\_Int **HYPRE\_BoomerAMGSetFSAILocalSolveType**(HYPRE\_Solver solver, HYPRE\_Int local\_solve\_type)

(Optional) Sets the solver type for solving local linear systems in FSAI.

For further explanation see *HYPRE\_FSAISetLocalSolveType*.

HYPRE\_Int **HYPRE\_BoomerAMGSetFSAIMaxSteps**(HYPRE\_Solver solver, HYPRE\_Int max\_steps)

(Optional) Defines maximum number of steps for FSAI.

For further explanation see *HYPRE\_FSAISetMaxSteps*.

HYPRE\_Int **HYPRE\_BoomerAMGSetFSAIMaxStepSize**(HYPRE\_Solver solver, HYPRE\_Int max\_step\_size)

(Optional) Defines maximum step size for FSAI.

For further explanation see *HYPRE\_FSAISetMaxStepSize*.

HYPRE\_Int **HYPRE\_BoomerAMGSetFSAIMaxNnzRow**(HYPRE\_Solver solver, HYPRE\_Int max\_nnz\_row)

(Optional) Defines maximum number of nonzero entries per row for FSAI.

For further explanation see *HYPRE\_FSAISetMaxNnzRow*.

HYPRE\_Int **HYPRE\_BoomerAMGSetFSAINumLevels**(HYPRE\_Solver solver, HYPRE\_Int num\_levels)

(Optional) Defines number of levels for computing the candidate pattern for FSAI For further explanation see *HYPRE\_FSAISetNumLevels*.

HYPRE\_Int **HYPRE\_BoomerAMGSetFSAIThreshold**(HYPRE\_Solver solver, HYPRE\_Real threshold)

(Optional) Defines the threshold for computing the candidate pattern for FSAI For further explanation see *HYPRE\_FSAISetThreshold*.

HYPRE\_Int **HYPRE\_BoomerAMGSetFSAIEigMaxIters**(HYPRE\_Solver solver, HYPRE\_Int eig\_max\_iters)

(Optional) Defines maximum number of iterations for estimating the largest eigenvalue of the FSAI pre-conditioned matrix ( $G^T * G * A$ ).

For further explanation see *HYPRE\_FSAISetEigMaxIters*.

HYPRE\_Int **HYPRE\_BoomerAMGSetFSAIKapTolerance**(HYPRE\_Solver solver, HYPRE\_Real kap\_tolerance)

(Optional) Defines the kaporin dropping tolerance.

For further explanation see *HYPRE\_FSAISetKapTolerance*.

HYPRE\_Int **HYPRE\_BoomerAMGSetRestriction**(HYPRE\_Solver solver, HYPRE\_Int restr\_par)

(Optional) Defines which parallel restriction operator is used.

There are the following options for restr\_type:

- 0 :  $P^T$  - Transpose of the interpolation operator
- 1 : AIR-1 - Approximate Ideal Restriction (distance 1)

- 2 : AIR-2 - Approximate Ideal Restriction (distance 2)

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetIsTriangular**(HYPRE\_Solver solver, HYPRE\_Int is\_triangular)

(Optional) Assumes the matrix is triangular in some ordering to speed up the setup time of approximate ideal restriction.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetGMRESSwitchR**(HYPRE\_Solver solver, HYPRE\_Int gmres\_switch)

(Optional) Set local problem size at which GMRES is used over a direct solve in approximating ideal restriction.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetADropTol**(HYPRE\_Solver solver, HYPRE\_Real A\_drop\_tol)

(Optional) Defines the drop tolerance for the A-matrices from the 2nd level of AMG.

The default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGSetADropType**(HYPRE\_Solver solver, HYPRE\_Int A\_drop\_type)

(Optional) Drop the entries that are not on the diagonal and smaller than its row norm: type 1: 1-norm, 2: 2-norm, -1: infinity norm

HYPRE\_Int **HYPRE\_BoomerAMGSetPrintFileName**(HYPRE\_Solver solver, const char \*print\_file\_name)

(Optional) Name of file to which BoomerAMG will print; cf HYPRE\_BoomerAMGSetPrintLevel.

(Presently this is ignored).

HYPRE\_Int **HYPRE\_BoomerAMGSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

(Optional) Requests automatic printing of setup and solve information.

- 0 : no printout (default)
- 1 : print setup information
- 2 : print solve information
- 3 : print both setup and solve information

Note, that if one desires to print information and uses BoomerAMG as a preconditioner, suggested *print\_level* is 1 to avoid excessive output, and use *print\_level* of solver for solve phase information.

HYPRE\_Int **HYPRE\_BoomerAMGGetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int \*print\_level)

HYPRE\_Int **HYPRE\_BoomerAMGSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

(Optional) Requests additional computations for diagnostic and similar data to be logged by the user.

Default to 0 to do nothing. The latest residual will be available if logging > 1.

HYPRE\_Int **HYPRE\_BoomerAMGGetLogging**(HYPRE\_Solver solver, HYPRE\_Int \*logging)

HYPRE\_Int **HYPRE\_BoomerAMGSetDebugFlag**(HYPRE\_Solver solver, HYPRE\_Int debug\_flag)

(Optional)

HYPRE\_Int **HYPRE\_BoomerAMGGetDebugFlag**(HYPRE\_Solver solver, HYPRE\_Int \*debug\_flag)

HYPRE\_Int **HYPRE\_BoomerAMGInitGridRelaxation**(HYPRE\_Int \*\*num\_grid\_sweeps\_ptr, HYPRE\_Int \*\*grid\_relax\_type\_ptr, HYPRE\_Int \*\*\*grid\_relax\_points\_ptr, HYPRE\_Int coarsen\_type, HYPRE\_Real \*\*relax\_weights\_ptr, HYPRE\_Int max\_levels)

(Optional) This routine will be eliminated in the future.

HYPRE\_Int **HYPRE\_BoomerAMGSetRAP2**(HYPRE\_Solver solver, HYPRE\_Int rap2)

(Optional) If rap2 not equal 0, the triple matrix product RAP is replaced by two matrix products.

(Required for triple matrix product generation on GPUs)

HYPRE\_Int **HYPRE\_BoomerAMGSetModuleRAP2**(HYPRE\_Solver solver, HYPRE\_Int mod\_rap2)

(Optional) If mod\_rap2 not equal 0, the triple matrix product RAP is replaced by two matrix products with modularized kernels (Required for triple matrix product generation on GPUs)

HYPRE\_Int **HYPRE\_BoomerAMGSetKeepTranspose**(HYPRE\_Solver solver, HYPRE\_Int keepTranspose)

(Optional) If set to 1, the local interpolation transposes will be saved to use more efficient matvecs instead of matvecTs (Recommended for efficient use on GPUs)

HYPRE\_Int **HYPRE\_BoomerAMGSetPlotGrids**(HYPRE\_Solver solver, HYPRE\_Int plotgrids)

HYPRE\_BoomerAMGSetPlotGrids.

HYPRE\_Int **HYPRE\_BoomerAMGSetPlotFileName**(HYPRE\_Solver solver, const char \*plotfilename)

HYPRE\_BoomerAMGSetPlotFilename.

HYPRE\_Int **HYPRE\_BoomerAMGSetCoordDim**(HYPRE\_Solver solver, HYPRE\_Int coorddim)

HYPRE\_BoomerAMGSetCoordDim.

HYPRE\_Int **HYPRE\_BoomerAMGSetCoordinates**(HYPRE\_Solver solver, float \*coordinates)

HYPRE\_BoomerAMGSetCoordinates.

HYPRE\_Int **HYPRE\_BoomerAMGGetGridHierarchy**(HYPRE\_Solver solver, HYPRE\_Int \*cgrid)

(Optional) Get the coarse grid hierarchy.

Assumes input/ output array is preallocated to the size of the local matrix. On return, *cgrid*[*i*] returns the last grid level containing node *i*.

#### Parameters

- **solver** – [IN] solver or preconditioner
- **cgrid** – [IN/ OUT] preallocated array. On return, contains grid hierarchy info.

HYPRE\_Int **HYPRE\_BoomerAMGSetCPoints**(HYPRE\_Solver solver, HYPRE\_Int cpt\_coarse\_level, HYPRE\_Int num\_cpt\_coarse, HYPRE\_BigInt \*cpt\_coarse\_index)

(Optional) Fix C points to be kept till a specified coarse level.

#### Parameters

- **solver** – [IN] solver or preconditioner
- **cpt\_coarse\_level** – [IN] coarse level up to which to keep C points
- **num\_cpt\_coarse** – [IN] number of C points to be kept
- **cpt\_coarse\_index** – [IN] indexes of C points to be kept

HYPRE\_Int **HYPRE\_BoomerAMGSetCpointsToKeep**(HYPRE\_Solver solver, HYPRE\_Int cpt\_coarse\_level, HYPRE\_Int num\_cpt\_coarse, HYPRE\_BigInt \*cpt\_coarse\_index)

(Optional) Deprecated function.

Use HYPRE\_BoomerAMGSetCPoints instead.

HYPRE\_Int **HYPRE\_BoomerAMGSetFPoints**(HYPRE\_Solver solver, HYPRE\_Int num\_fpt, HYPRE\_BigInt \*fpt\_index)

(Optional) Set fine points in the first level.

**Parameters**

- **solver** – [IN] solver or preconditioner
- **num\_fpt** – [IN] number of fine points
- **fpt\_index** – [IN] global indices of fine points

HYPRE\_Int **HYPRE\_BoomerAMGSetIsolatedFPoints**(HYPRE\_Solver solver, HYPRE\_Int num\_isolated\_fpt, HYPRE\_BigInt \*isolated\_fpt\_index)

(Optional) Set isolated fine points in the first level.

Interpolation weights are not computed for these points.

**Parameters**

- **solver** – [IN] solver or preconditioner
- **num\_isolated\_fpt** – [IN] number of isolated fine points
- **isolated\_fpt\_index** – [IN] global indices of isolated fine points

HYPRE\_Int **HYPRE\_BoomerAMGSetSabs**(HYPRE\_Solver solver, HYPRE\_Int Sabs)

(Optional) if Sabs equals 1, the strength of connection test is based on the absolute value of the matrix coefficients

HYPRE\_Int **HYPRE\_BoomerAMGGetMaxRowSum**(HYPRE\_Solver solver, HYPRE\_Real \*max\_row\_sum)

HYPRE\_Int **HYPRE\_BoomerAMGSetPostInterpType**(HYPRE\_Solver solver, HYPRE\_Int post\_interp\_type)

(Optional) Switches on use of Jacobi interpolation after computing an original interpolation

HYPRE\_Int **HYPRE\_BoomerAMGGetPostInterpType**(HYPRE\_Solver solver, HYPRE\_Int \*post\_interp\_type)

HYPRE\_Int **HYPRE\_BoomerAMGSetJacobiTruncThreshold**(HYPRE\_Solver solver, HYPRE\_Real jacobi\_trunc\_threshold)

(Optional) Sets a truncation threshold for Jacobi interpolation.

HYPRE\_Int **HYPRE\_BoomerAMGGetJacobiTruncThreshold**(HYPRE\_Solver solver, HYPRE\_Real \*jacobi\_trunc\_threshold)

HYPRE\_Int **HYPRE\_BoomerAMGSetNumCRRelaxSteps**(HYPRE\_Solver solver, HYPRE\_Int num\_CR\_relax\_steps)

(Optional) Defines the number of relaxation steps for CR The default is 2.

HYPRE\_Int **HYPRE\_BoomerAMGSetCRRate**(HYPRE\_Solver solver, HYPRE\_Real CR\_rate)

(Optional) Defines convergence rate for CR The default is 0.7.

HYPRE\_Int **HYPRE\_BoomerAMGSetCRStrongTh**(HYPRE\_Solver solver, HYPRE\_Real CR\_strong\_th)  
 (Optional) Defines strong threshold for CR The default is 0.0.

HYPRE\_Int **HYPRE\_BoomerAMGSetCRUseCG**(HYPRE\_Solver solver, HYPRE\_Int CR\_use\_CG)  
 (Optional) Defines whether to use CG

HYPRE\_Int **HYPRE\_BoomerAMGSetISType**(HYPRE\_Solver solver, HYPRE\_Int IS\_type)  
 (Optional) Defines the Type of independent set algorithm used for CR

### ParCSR BoomerAMGDD Solver and Preconditioner

Communication reducing solver and preconditioner built on top of algebraic multigrid

HYPRE\_Int **HYPRE\_BoomerAMGDDCreate**(HYPRE\_Solver \*solver)  
 Create a solver object.

HYPRE\_Int **HYPRE\_BoomerAMGDDDestroy**(HYPRE\_Solver solver)  
 Destroy a solver object.

HYPRE\_Int **HYPRE\_BoomerAMGDDSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A,  
*HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Set up the BoomerAMGDD solver or preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] object to be set up.
- **A** – [IN] ParCSR matrix used to construct the solver/preconditioner.
- **b** – Ignored by this function.
- **x** – Ignored by this function.

HYPRE\_Int **HYPRE\_BoomerAMGDDsolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A,  
*HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Solve the system or apply AMG-DD as a preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] solver or preconditioner object to be applied.
- **A** – [IN] ParCSR matrix, matrix of the linear system to be solved
- **b** – [IN] right hand side of the linear system to be solved
- **x** – [OUT] approximated solution of the linear system to be solved

HYPRE\_Int **HYPRE\_BoomerAMGDDSetFACNumRelax**(HYPRE\_Solver solver, HYPRE\_Int  
 amgdd\_fac\_num\_relax)

(Optional) Set the number of pre- and post-relaxations per level for AMG-DD inner FAC cycles.

Default is 1.

HYPRE\_Int **HYPRE\_BoomerAMGDDSetFACNumCycles**(HYPRE\_Solver solver, HYPRE\_Int  
 amgdd\_fac\_num\_cycles)

(Optional) Set the number of inner FAC cycles per AMG-DD iteration.

Default is 2.

HYPRE\_Int **HYPRE\_BoomerAMGDDSetFACCycleType**(HYPRE\_Solver solver, HYPRE\_Int  
amgdd\_fac\_cycle\_type)

(Optional) Set the cycle type for the AMG-DD inner FAC cycles.

1 (default) = V-cycle, 2 = W-cycle, 3 = F-cycle

HYPRE\_Int **HYPRE\_BoomerAMGDDSetFACRelaxType**(HYPRE\_Solver solver, HYPRE\_Int  
amgdd\_fac\_relax\_type)

(Optional) Set the relaxation type for the AMG-DD inner FAC cycles.

0 = Jacobi, 1 = Gauss-Seidel, 2 = ordered Gauss-Seidel, 3 (default) = C/F L1-scaled Jacobi

HYPRE\_Int **HYPRE\_BoomerAMGDDSetFACRelaxWeight**(HYPRE\_Solver solver, HYPRE\_Real  
amgdd\_fac\_relax\_weight)

(Optional) Set the relaxation weight for the AMG-DD inner FAC cycles.

Default is 1.0.

HYPRE\_Int **HYPRE\_BoomerAMGDDSetStartLevel**(HYPRE\_Solver solver, HYPRE\_Int start\_level)

(Optional) Set the AMG-DD start level.

Default is 0.

HYPRE\_Int **HYPRE\_BoomerAMGDDSetPadding**(HYPRE\_Solver solver, HYPRE\_Int padding)

(Optional) Set the AMG-DD padding.

Default is 1.

HYPRE\_Int **HYPRE\_BoomerAMGDDSetNumGhostLayers**(HYPRE\_Solver solver, HYPRE\_Int  
num\_ghost\_layers)

(Optional) Set the AMG-DD number of ghost layers.

Default is 1.

HYPRE\_Int **HYPRE\_BoomerAMGDDSetUserFACRelaxation**(HYPRE\_Solver solver, HYPRE\_Int  
(\*userFACRelaxation)(void \*amgdd\_vdata,  
HYPRE\_Int level, HYPRE\_Int cycle\_param))

(Optional) Pass a custom user-defined function as a relaxation method for the AMG-DD FAC cycles.

Function should have the following form, where amgdd\_solver is of type `hypr_ParAMGDDData*` and level is the level on which to relax: `HYPRE_Int userFACRelaxation( HYPRE_Solver amgdd_solver, HYPRE_Int level )`

HYPRE\_Int **HYPRE\_BoomerAMGDDGetAMG**(HYPRE\_Solver solver, HYPRE\_Solver \*amg\_solver)

(Optional) Get the underlying AMG hierarchy as a HYPRE\_Solver object.

HYPRE\_Int **HYPRE\_BoomerAMGDDGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real  
\*rel\_resid\_norm)

Returns the norm of the final relative residual.

HYPRE\_Int **HYPRE\_BoomerAMGDDGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int  
\*num\_iterations)

Returns the number of iterations taken.

### ParCSR FSAI Solver and Preconditioner

An adaptive factorized sparse approximate inverse solver/preconditioner/smoothen that computes a sparse approximation  $G$  to the inverse of the lower cholesky factor of  $A$  such that  $M^{-1} \approx G^T * G$ .

HYPRE\_Int **HYPRE\_FSAICreate**(HYPRE\_Solver \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_FSAIDestroy**(HYPRE\_Solver solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_FSAISetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Set up the FSAI solver or preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

**Parameters**

- **solver** – [IN] object to be set up.
- **A** – [IN] ParCSR matrix used to construct the solver/preconditioner.
- **b** – Ignored by this function.
- **x** – Ignored by this function.

HYPRE\_Int **HYPRE\_FSAISolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Solve the system or apply FSAI as a preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

**Parameters**

- **solver** – [IN] solver or preconditioner object to be applied.
- **A** – [IN] ParCSR matrix, matrix of the linear system to be solved
- **b** – [IN] right hand side of the linear system to be solved
- **x** – [OUT] approximated solution of the linear system to be solved

HYPRE\_Int **HYPRE\_FSAISetAlgoType**(HYPRE\_Solver solver, HYPRE\_Int algo\_type)

(Optional) Sets the algorithm type used to compute the lower triangular factor G

- 1: Adaptive (can use OpenMP **with** static scheduling)
- 2: Adaptive OpenMP **with** dynamic scheduling
- 3: Static - power pattern

HYPRE\_Int **HYPRE\_FSAISetLocalSolveType**(HYPRE\_Solver solver, HYPRE\_Int local\_solve\_type)

(Optional) Sets the solver type for solving local linear systems in FSAI.

This option makes sense only for GPU runs.

- 0: Gauss-Jordan solver
- 1: Vendor solver (cuSOLVER/rocSOLVER)
- 2: MAGMA solver

HYPRE\_Int **HYPRE\_FSAISetMaxSteps**(HYPRE\_Solver solver, HYPRE\_Int max\_steps)

(Optional) Sets the maximum number of steps for computing the sparsity pattern of G.

This input parameter makes sense when using adaptive FSAI, i.e., algorithm type 1 or 2.

`HYPRE_Int HYPRE_FSAISetMaxStepSize(HYPRE_Solver solver, HYPRE_Int max_step_size)`  
 (Optional) Sets the maximum step size for computing the sparsity pattern of G.  
 This input parameter makes sense when using adaptive FSAI, i.e., algorithm type 1 or 2.

`HYPRE_Int HYPRE_FSAISetMaxNnzRow(HYPRE_Solver solver, HYPRE_Int max_nnz_row)`  
 (Optional) Sets the maximum number of off-diagonal entries per row of G.  
 This input parameter makes sense when using static FSAI, i.e., algorithm type 3.

`HYPRE_Int HYPRE_FSAISetNumLevels(HYPRE_Solver solver, HYPRE_Int num_levels)`  
 (Optional) Sets the number of levels for computing the candidate pattern of G.  
 This input parameter must be a positive integer and it makes sense when using static FSAI, i.e., algorithm type 3.

`HYPRE_Int HYPRE_FSAISetThreshold(HYPRE_Solver solver, HYPRE_Real threshold)`  
 (Optional) Sets the threshold for computing the candidate pattern of G This input parameter makes sense when using static FSAI, i.e., algorithm type 3.

`HYPRE_Int HYPRE_FSAISetKapTolerance(HYPRE_Solver solver, HYPRE_Real kap_tolerance)`  
 (Optional) Sets the kaporin gradient reduction factor for computing the sparsity pattern of G.  
 This input parameter makes sense when using adaptive FSAI, i.e., algorithm types 1 or 2.

`HYPRE_Int HYPRE_FSAISetOmega(HYPRE_Solver solver, HYPRE_Real omega)`  
 (Optional) Sets the relaxation factor for FSAI.  
 This input parameter makes sense to all algorithm types for setting up FSAI.

`HYPRE_Int HYPRE_FSAISetMaxIterations(HYPRE_Solver solver, HYPRE_Int max_iterations)`  
 (Optional) Sets the maximum number of iterations (sweeps) for FSAI.  
 This input parameter makes sense to all algorithm types for setting up FSAI.

`HYPRE_Int HYPRE_FSAISetEigMaxIters(HYPRE_Solver solver, HYPRE_Int eig_max_iters)`  
 (Optional) Set number of iterations for computing maximum eigenvalue of the preconditioned operator.  
 This input parameter makes sense to all algorithm types for setting up FSAI.

`HYPRE_Int HYPRE_FSAISetTolerance(HYPRE_Solver solver, HYPRE_Real tolerance)`  
 (Optional) Set the convergence tolerance, if FSAI is used as a solver.  
 This input parameter makes sense to all algorithm types for setting up FSAI. When using FSAI as a preconditioner, set the tolerance to 0.0. The default is  $10^{-6}$ .

`HYPRE_Int HYPRE_FSAISetPrintLevel(HYPRE_Solver solver, HYPRE_Int print_level)`  
 (Optional) Requests automatic printing of setup information.

- 0 : no printout (default)
- 1 : print setup information

`HYPRE_Int HYPRE_FSAISetZeroGuess(HYPRE_Solver solver, HYPRE_Int zero_guess)`  
 (Optional) Use a zero initial guess.  
 This allows the solver to cut corners in the case where a zero initial guess is needed (e.g., for preconditioning) to reduce computational cost.

## ParCSR Chebyshev Solver and Preconditioner

Solver based on Chebyshev polynomials

HYPRE\_Int **HYPRE\_ParChebyCreate**(HYPRE\_Solver \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_ParChebyDestroy**(HYPRE\_Solver solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_ParChebySetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Set up the Chebyshev solver or preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

### Parameters

- **solver** – [IN] object to be set up.
- **A** – [IN] ParCSR matrix used to construct the solver/preconditioner.
- **b** – Ignored by this function.
- **x** – Ignored by this function.

HYPRE\_Int **HYPRE\_ParChebySolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Solve the system or apply Chebyshev as a preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

### Parameters

- **solver** – [IN] solver or preconditioner object to be applied.
- **A** – [IN] ParCSR matrix, matrix of the linear system to be solved
- **b** – [IN] right hand side of the linear system to be solved
- **x** – [OUT] approximated solution of the linear system to be solved

HYPRE\_Int **HYPRE\_ParChebySetMaxIterations**(HYPRE\_Solver solver, HYPRE\_Int max\_iterations)

(Optional) Sets the maximum number of iterations (sweeps) for Chebyshev.

Default is 100.

HYPRE\_Int **HYPRE\_ParChebySetTolerance**(HYPRE\_Solver solver, HYPRE\_Real tolerance)

(Optional) Set the convergence tolerance used by Chebyshev.

When using Chebyshev as a preconditioner, set the tolerance to 0.0. The default is  $10^{-6}$ .

HYPRE\_Int **HYPRE\_ParChebySetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

(Optional) Requests automatic printing of setup information.

- 0 : no printout (default)
- 1 : print setup information

HYPRE\_Int **HYPRE\_ParChebySetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

(Optional) Requests additional computations for diagnostic and similar data to be logged by the user.

Default to 0 to do nothing. The latest residual will be available if logging > 1.

HYPRE\_Int **HYPRE\_ParChebySetOrder**(HYPRE\_Solver solver, HYPRE\_Int order)

(Optional) Defines the Order for Chebyshev smoother.

The default is 2 (valid options are 1-4).

HYPRE\_Int **HYPRE\_ParChebySetVariant**(HYPRE\_Solver solver, HYPRE\_Int variant)

(Optional) Defines which polynomial variant should be used.

The default is 0 (i.e., scaled).

HYPRE\_Int **HYPRE\_ParChebySetScale**(HYPRE\_Solver solver, HYPRE\_Int scale)

(Optional) Defines whether matrix should be scaled.

The default is 1 (i.e., scaled).

HYPRE\_Int **HYPRE\_ParChebySetEigRatio**(HYPRE\_Solver solver, HYPRE\_Real eig\_ratio)

(Optional) Fraction of the spectrum to use for the Chebyshev smoother.

The default is .3 (i.e., damp on upper 30% of the spectrum).

HYPRE\_Int **HYPRE\_ParChebySetEigEst**(HYPRE\_Solver solver, HYPRE\_Int eig\_est)

(Optional) Defines how to estimate eigenvalues.

The default is 10 CG iterations are used to find extreme eigenvalues. If `eig_est` is 0, the largest eigenvalue is estimated using Gershgorin, while the smallest eigenvalue is set to 0. If `eig_est` is a positive number `n`, `n` iterations of CG are used to estimate the smallest and largest eigenvalue.

HYPRE\_Int **HYPRE\_ParChebySetMinMaxEigEst**(HYPRE\_Solver solver, HYPRE\_Real min\_eig\_est,  
HYPRE\_Real max\_eig\_est)

(Optional) Set minimum and maximum eigenvalues

HYPRE\_Int **HYPRE\_ParChebyGetMinMaxEigEst**(HYPRE\_Solver solver, HYPRE\_Real \*min\_eig\_est,  
HYPRE\_Real \*max\_eig\_est)

(Optional) Get minimum and maximum eigenvalues

HYPRE\_Int **HYPRE\_ParChebySetZeroGuess**(HYPRE\_Solver solver, HYPRE\_Int zero\_guess)

(Optional) Use a zero initial guess.

This allows the solver to cut corners in the case where a zero initial guess is needed (e.g., for preconditioning) to reduce computational cost.

### ParCSR ParaSails Preconditioner

Parallel sparse approximate inverse preconditioner for the ParCSR matrix format.

HYPRE\_Int **HYPRE\_ParaSailsCreate**(MPI\_Comm comm, HYPRE\_Solver \*solver)

Create a ParaSails preconditioner.

HYPRE\_Int **HYPRE\_ParaSailsDestroy**(HYPRE\_Solver solver)

Destroy a ParaSails preconditioner.

HYPRE\_Int **HYPRE\_ParaSailsSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Set up the ParaSails preconditioner.

This function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] Preconditioner object to set up.

- **A** – [IN] ParCSR matrix used to construct the preconditioner.
- **b** – Ignored by this function.
- **x** – Ignored by this function.

HYPRE\_Int **HYPRE\_ParaSailsSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Apply the ParaSails preconditioner.

This function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] Preconditioner object to apply.
- **A** – Ignored by this function.
- **b** – [IN] Vector to precondition.
- **x** – [OUT] Preconditioned vector.

HYPRE\_Int **HYPRE\_ParaSailsSetParams**(HYPRE\_Solver solver, HYPRE\_Real thresh, HYPRE\_Int nlevels)

Set the threshold and levels parameter for the ParaSails preconditioner.

The accuracy and cost of ParaSails are parameterized by these two parameters. Lower values of the threshold parameter and higher values of levels parameter lead to more accurate, but more expensive preconditioners.

#### Parameters

- **solver** – [IN] Preconditioner object for which to set parameters.
- **thresh** – [IN] Value of threshold parameter,  $0 \leq \text{thresh} \leq 1$ . The default value is 0.1.
- **nlevels** – [IN] Value of levels parameter,  $0 \leq \text{nlevels}$ . The default value is 1.

HYPRE\_Int **HYPRE\_ParaSailsSetFilter**(HYPRE\_Solver solver, HYPRE\_Real filter)

Set the filter parameter for the ParaSails preconditioner.

#### Parameters

- **solver** – [IN] Preconditioner object for which to set filter parameter.
- **filter** – [IN] Value of filter parameter. The filter parameter is used to drop small nonzeros in the preconditioner, to reduce the cost of applying the preconditioner. Values from 0.05 to 0.1 are recommended. The default value is 0.1.

HYPRE\_Int **HYPRE\_ParaSailsSetSym**(HYPRE\_Solver solver, HYPRE\_Int sym)

Set the symmetry parameter for the ParaSails preconditioner.

Values for *sym*

- 0 : nonsymmetric and/or indefinite problem, and nonsymmetric preconditioner
- 1 : SPD problem, and SPD (factored) preconditioner
- 2 : nonsymmetric, definite problem, and SPD (factored) preconditioner

#### Parameters

- **solver** – [IN] Preconditioner object for which to set symmetry parameter.
- **sym** – [IN] Symmetry parameter.

HYPRE\_Int **HYPRE\_ParaSailsSetLoadbal**(HYPRE\_Solver solver, HYPRE\_Real loadbal)

Set the load balance parameter for the ParaSails preconditioner.

**Parameters**

- **solver** – [IN] Preconditioner object for which to set the load balance parameter.
- **loadbal** – [IN] Value of the load balance parameter,  $0 \leq \text{loadbal} \leq 1$ . A zero value indicates that no load balance is attempted; a value of unity indicates that perfect load balance will be attempted. The recommended value is 0.9 to balance the overhead of data exchanges for load balancing. No load balancing is needed if the preconditioner is very sparse and fast to construct. The default value when this parameter is not set is 0.

HYPRE\_Int **HYPRE\_ParaSailsSetReuse**(HYPRE\_Solver solver, HYPRE\_Int reuse)

Set the pattern reuse parameter for the ParaSails preconditioner.

**Parameters**

- **solver** – [IN] Preconditioner object for which to set the pattern reuse parameter.
- **reuse** – [IN] Value of the pattern reuse parameter. A nonzero value indicates that the pattern of the preconditioner should be reused for subsequent constructions of the preconditioner. A zero value indicates that the preconditioner should be constructed from scratch. The default value when this parameter is not set is 0.

HYPRE\_Int **HYPRE\_ParaSailsSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

Set the logging parameter for the ParaSails preconditioner.

**Parameters**

- **solver** – [IN] Preconditioner object for which to set the logging parameter.
- **logging** – [IN] Value of the logging parameter. A nonzero value sends statistics of the setup procedure to stdout. The default value when this parameter is not set is 0.

HYPRE\_Int **HYPRE\_ParaSailsBuildIJMatrix**(HYPRE\_Solver solver, *HYPRE\_IJMatrix* \*pij\_A)

Build IJ Matrix of the sparse approximate inverse (factor).

This function explicitly creates the IJ Matrix corresponding to the sparse approximate inverse or the inverse factor. Example: *HYPRE\_IJMatrix* ij\_A; **HYPRE\_ParaSailsBuildIJMatrix**(solver, &ij\_A);

**Parameters**

- **solver** – [IN] Preconditioner object.
- **pij\_A** – [OUT] Pointer to the IJ Matrix.

HYPRE\_Int **HYPRE\_ParCSRParaSailsCreate**(MPI\_Comm comm, HYPRE\_Solver \*solver)

HYPRE\_Int **HYPRE\_ParCSRParaSailsDestroy**(HYPRE\_Solver solver)

HYPRE\_Int **HYPRE\_ParCSRParaSailsSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRParaSailsSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRParaSailsSetParams**(HYPRE\_Solver solver, HYPRE\_Real thresh, HYPRE\_Int nlevels)

HYPRE\_Int **HYPRE\_ParaSailsSetThresh**(HYPRE\_Solver solver, HYPRE\_Real thresh)

HYPRE\_Int **HYPRE\_ParCSRSailsGetThresh**(HYPRE\_Solver solver, HYPRE\_Real \*thresh)  
 HYPRE\_Int **HYPRE\_ParCSRSailsSetNlevels**(HYPRE\_Solver solver, HYPRE\_Int nlevels)  
 HYPRE\_Int **HYPRE\_ParCSRSailsGetNlevels**(HYPRE\_Solver solver, HYPRE\_Int \*nlevels)  
 HYPRE\_Int **HYPRE\_ParCSRSailsSetFilter**(HYPRE\_Solver solver, HYPRE\_Real filter)  
 HYPRE\_Int **HYPRE\_ParCSRSailsGetFilter**(HYPRE\_Solver solver, HYPRE\_Real \*filter)  
 HYPRE\_Int **HYPRE\_ParCSRSailsSetSym**(HYPRE\_Solver solver, HYPRE\_Int sym)  
 HYPRE\_Int **HYPRE\_ParCSRSailsGetSym**(HYPRE\_Solver solver, HYPRE\_Int \*sym)  
 HYPRE\_Int **HYPRE\_ParCSRSailsSetLoadbal**(HYPRE\_Solver solver, HYPRE\_Real loadbal)  
 HYPRE\_Int **HYPRE\_ParCSRSailsGetLoadbal**(HYPRE\_Solver solver, HYPRE\_Real \*loadbal)  
 HYPRE\_Int **HYPRE\_ParCSRSailsSetReuse**(HYPRE\_Solver solver, HYPRE\_Int reuse)  
 HYPRE\_Int **HYPRE\_ParCSRSailsGetReuse**(HYPRE\_Solver solver, HYPRE\_Int \*reuse)  
 HYPRE\_Int **HYPRE\_ParCSRSailsSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)  
 HYPRE\_Int **HYPRE\_ParCSRSailsGetLogging**(HYPRE\_Solver solver, HYPRE\_Int \*logging)

### ParCSR Euclid Preconditioner

MPI Parallel ILU preconditioner

Options summary:

Option	Default	Synopsis
-level	1	ILU(k) factorization level
-bj	0 (false)	Use Block Jacobi ILU instead of PILU
-eu_stats	0 (false)	Print internal timing and statistics
-eu_mem	0 (false)	Print internal memory usage

HYPRE\_Int **HYPRE\_EuclidCreate**(MPI\_Comm comm, HYPRE\_Solver \*solver)

Create a Euclid object.

HYPRE\_Int **HYPRE\_EuclidDestroy**(HYPRE\_Solver solver)

Destroy a Euclid object.

HYPRE\_Int **HYPRE\_EuclidSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Set up the Euclid preconditioner.

This function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] Preconditioner object to set up.
- **A** – [IN] ParCSR matrix used to construct the preconditioner.
- **b** – Ignored by this function.

- **x** – Ignored by this function.

HYPRE\_Int **HYPRE\_EuclidSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Apply the Euclid preconditioner.

This function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] Preconditioner object to apply.
- **A** – Ignored by this function.
- **b** – [IN] Vector to precondition.
- **x** – [OUT] Preconditioned vector.

HYPRE\_Int **HYPRE\_EuclidSetParams**(HYPRE\_Solver solver, HYPRE\_Int argc, char \*argv[])

Insert (name, value) pairs in Euclid’s options database by passing Euclid the command line (or an array of strings).

All Euclid options (e.g, level, drop-tolerance) are stored in this database. If a (name, value) pair already exists, this call updates the value. See also: *HYPRE\_EuclidSetParamsFromFile*.

#### Parameters

- **argc** – [IN] Length of argv array
- **argv** – [IN] Array of strings

HYPRE\_Int **HYPRE\_EuclidSetParamsFromFile**(HYPRE\_Solver solver, char \*filename)

Insert (name, value) pairs in Euclid’s options database.

Each line of the file should either begin with a “#”, indicating a comment line, or contain a (name value) pair, e.g:

```
>cat optionsFile
\#sample runtime parameter file
-blockJacobi 3
-matFile /home/hysom/myfile.euclid
-doSomething true
-xx_coeff -1.0
```

See also: *HYPRE\_EuclidSetParams*.

#### Parameters

**filename**[IN] – Pathname/filename to read

HYPRE\_Int **HYPRE\_EuclidSetLevel**(HYPRE\_Solver solver, HYPRE\_Int level)

Set level k for ILU(k) factorization, default: 1.

HYPRE\_Int **HYPRE\_EuclidSetBJ**(HYPRE\_Solver solver, HYPRE\_Int bj)

Use block Jacobi ILU preconditioning instead of PILU.

HYPRE\_Int **HYPRE\_EuclidSetStats**(HYPRE\_Solver solver, HYPRE\_Int eu\_stats)

If *eu\_stats* not equal 0, a summary of runtime settings and timing information is printed to stdout.

HYPRE\_Int **HYPRE\_EuclidSetMem**(HYPRE\_Solver solver, HYPRE\_Int eu\_mem)

If *eu\_mem* not equal 0, a summary of Euclid’s memory usage is printed to stdout.

HYPRE\_Int **HYPRE\_EuclidSetSparseA**(HYPRE\_Solver solver, HYPRE\_Real sparse\_A)

Defines a drop tolerance for ILU(k).

Default: 0 Use with HYPRE\_EuclidSetRowScale. Note that this can destroy symmetry in a matrix.

HYPRE\_Int **HYPRE\_EuclidSetRowScale**(HYPRE\_Solver solver, HYPRE\_Int row\_scale)

If *row\_scale* not equal 0, values are scaled prior to factorization so that largest value in any row is +1 or -1.

Note that this can destroy symmetry in a matrix.

HYPRE\_Int **HYPRE\_EuclidSetILUT**(HYPRE\_Solver solver, HYPRE\_Real drop\_tol)

uses ILUT and defines a drop tolerance relative to the largest absolute value of any entry in the row being factored.

### ParCSR Pilut Preconditioner

HYPRE\_Int **HYPRE\_ParCSRPilutCreate**(MPI\_Comm comm, HYPRE\_Solver \*solver)

Create a preconditioner object.

HYPRE\_Int **HYPRE\_ParCSRPilutDestroy**(HYPRE\_Solver solver)

Destroy a preconditioner object.

HYPRE\_Int **HYPRE\_ParCSRPilutSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A,  
*HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRPilutSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A,  
*HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Precondition the system.

HYPRE\_Int **HYPRE\_ParCSRPilutSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_ParCSRPilutSetDropTolerance**(HYPRE\_Solver solver, HYPRE\_Real tol)

(Optional)

HYPRE\_Int **HYPRE\_ParCSRPilutSetFactorRowSize**(HYPRE\_Solver solver, HYPRE\_Int size)

(Optional)

HYPRE\_Int **HYPRE\_ParCSRPilutSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

### ParCSR AMS Solver and Preconditioner

Parallel auxiliary space Maxwell solver and preconditioner

HYPRE\_Int **HYPRE\_AMSCreate**(HYPRE\_Solver \*solver)

Create an AMS solver object.

HYPRE\_Int **HYPRE\_AMSDestroy**(HYPRE\_Solver solver)

Destroy an AMS solver object.

HYPRE\_Int **HYPRE\_AMSSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b,  
*HYPRE\_ParVector* x)

Set up the AMS solver or preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] object to be set up.

- **A** – [IN] ParCSR matrix used to construct the solver/preconditioner.
- **b** – Ignored by this function.
- **x** – Ignored by this function.

HYPRE\_Int **HYPRE\_AMSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Solve the system or apply AMS as a preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] solver or preconditioner object to be applied.
- **A** – [IN] ParCSR matrix, matrix of the linear system to be solved
- **b** – [IN] right hand side of the linear system to be solved
- **x** – [OUT] approximated solution of the linear system to be solved

HYPRE\_Int **HYPRE\_AMSSetDimension**(HYPRE\_Solver solver, HYPRE\_Int dim)

(Optional) Sets the problem dimension (2 or 3).

The default is 3.

HYPRE\_Int **HYPRE\_AMSSetDiscreteGradient**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* G)

Sets the discrete gradient matrix *G*.

This function should be called before *HYPRE\_AMSSetup*()!

HYPRE\_Int **HYPRE\_AMSSetCoordinateVectors**(HYPRE\_Solver solver, *HYPRE\_ParVector* x, *HYPRE\_ParVector* y, *HYPRE\_ParVector* z)

Sets the *x*, *y* and *z* coordinates of the vertices in the mesh.

Either *HYPRE\_AMSSetCoordinateVectors*() or *HYPRE\_AMSSetEdgeConstantVectors*() should be called before *HYPRE\_AMSSetup*()!

HYPRE\_Int **HYPRE\_AMSSetEdgeConstantVectors**(HYPRE\_Solver solver, *HYPRE\_ParVector* G<sub>x</sub>, *HYPRE\_ParVector* G<sub>y</sub>, *HYPRE\_ParVector* G<sub>z</sub>)

Sets the vectors *G<sub>x</sub>*, *G<sub>y</sub>* and *G<sub>z</sub>* which give the representations of the constant vector fields (1,0,0), (0,1,0) and (0,0,1) in the edge element basis.

Either *HYPRE\_AMSSetCoordinateVectors*() or *HYPRE\_AMSSetEdgeConstantVectors*() should be called before *HYPRE\_AMSSetup*()!

HYPRE\_Int **HYPRE\_AMSSetInterpolations**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* Pi, *HYPRE\_ParCSRMatrix* Pix, *HYPRE\_ParCSRMatrix* Piy, *HYPRE\_ParCSRMatrix* Piz)

(Optional) Set the (components of) the Nedelec interpolation matrix  $\Pi = [\Pi^x, \Pi^y, \Pi^z]$ .

This function is generally intended to be used only for high-order Nedelec discretizations (in the lowest order case,  $\Pi$  is constructed internally in AMS from the discrete gradient matrix and the coordinates of the vertices), though it can also be used in the lowest-order case or for other types of discretizations (e.g. ones based on the second family of Nedelec elements).

By definition,  $\Pi$  is the matrix representation of the linear operator that interpolates (high-order) vector nodal finite elements into the (high-order) Nedelec space. The component matrices are defined as  $\Pi^x \varphi = \Pi(\varphi, 0, 0)$  and similarly for  $\Pi^y$  and  $\Pi^z$ . Note that all these operators depend on the choice of the basis and degrees of freedom in the high-order spaces.

The column numbering of  $\Pi$  should be node-based, i.e. the  $x/y/z$  components of the first node (vertex or high-order dof) should be listed first, followed by the  $x/y/z$  components of the second node and so on (see the documentation of `HYPRE_BoomerAMGSetDofFunc`).

If used, this function should be called before `HYPRE_AMSSetup()` and there is no need to provide the vertex coordinates. Furthermore, only one of the sets  $\{\Pi\}$  and  $\{\Pi^x, \Pi^y, \Pi^z\}$  needs to be specified (though it is OK to provide both). If  $\Pi_x$  is NULL, then scalar  $\Pi$ -based AMS cycles, i.e. those with `cycle_type > 10`, will be unavailable. Similarly, AMS cycles based on monolithic  $\Pi$  (`cycle_type < 10`) require that  $\Pi$  is not NULL.

`HYPRE_Int HYPRE_AMSSetAlphaPoissonMatrix`(`HYPRE_Solver solver`, `HYPRE_ParCSRMatrix A_alpha`)

(Optional) Sets the matrix  $A_\alpha$  corresponding to the Poisson problem with coefficient  $\alpha$  (the curl-curl term coefficient in the Maxwell problem).

If this function is called, the coarse space solver on the range of  $\Pi^T$  is a block-diagonal version of  $A_\Pi$ . If this function is not called, the coarse space solver on the range of  $\Pi^T$  is constructed as  $\Pi^T A \Pi$  in `HYPRE_AMSSetup()`. See the user's manual for more details.

`HYPRE_Int HYPRE_AMSSetBetaPoissonMatrix`(`HYPRE_Solver solver`, `HYPRE_ParCSRMatrix A_beta`)

(Optional) Sets the matrix  $A_\beta$  corresponding to the Poisson problem with coefficient  $\beta$  (the mass term coefficient in the Maxwell problem).

If not given, the Poisson matrix will be computed in `HYPRE_AMSSetup()`. If the given matrix is NULL, we assume that  $\beta$  is identically 0 and use two-level (instead of three-level) methods. See the user's manual for more details.

`HYPRE_Int HYPRE_AMSSetInteriorNodes`(`HYPRE_Solver solver`, `HYPRE_ParVector interior_nodes`)

(Optional) Set the list of nodes which are interior to a zero-conductivity region.

This way, a more robust solver is constructed, that can be iterated to lower tolerance levels. A node is interior if its entry in the array is 1.0. This function should be called before `HYPRE_AMSSetup()`!

`HYPRE_Int HYPRE_AMSSetProjectionFrequency`(`HYPRE_Solver solver`, `HYPRE_Int projection_frequency`)

(Optional) Set the frequency at which a projection onto the compatible subspace for problems with zero-conductivity regions is performed.

The default value is 5.

`HYPRE_Int HYPRE_AMSSetMaxIter`(`HYPRE_Solver solver`, `HYPRE_Int maxit`)

(Optional) Sets maximum number of iterations, if AMS is used as a solver.

To use AMS as a preconditioner, set the maximum number of iterations to 1. The default is 20.

`HYPRE_Int HYPRE_AMSSetTol`(`HYPRE_Solver solver`, `HYPRE_Real tol`)

(Optional) Set the convergence tolerance, if AMS is used as a solver.

When using AMS as a preconditioner, set the tolerance to 0.0. The default is  $10^{-6}$ .

`HYPRE_Int HYPRE_AMSSetCycleType`(`HYPRE_Solver solver`, `HYPRE_Int cycle_type`)

(Optional) Choose which three-level solver to use.

Possible values are:

- 1 : 3-level multiplicative solver (01210)
- 2 : 3-level additive solver (0+1+2)
- 3 : 3-level multiplicative solver (02120)

- 4 : 3-level additive solver (010+2)
- 5 : 3-level multiplicative solver (0102010)
- 6 : 3-level additive solver (1+020)
- 7 : 3-level multiplicative solver (0201020)
- 8 : 3-level additive solver (0(1+2)0)
- 11 : 5-level multiplicative solver (013454310)
- 12 : 5-level additive solver (0+1+3+4+5)
- 13 : 5-level multiplicative solver (034515430)
- 14 : 5-level additive solver (01(3+4+5)10)

The default is 1. See the user's manual for more details.

HYPRE\_Int **HYPRE\_AMSSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

(Optional) Control how much information is printed during the solution iterations.

The default is 1 (print residual norm at each step).

HYPRE\_Int **HYPRE\_AMSSetSmoothingOptions**(HYPRE\_Solver solver, HYPRE\_Int relax\_type,  
HYPRE\_Int relax\_times, HYPRE\_Real relax\_weight,  
HYPRE\_Real omega)

(Optional) Sets relaxation parameters for  $A$ .

The defaults are 2, 1, 1.0, 1.0.

The available options for *relax\_type* are:

- 1 :  $\ell_1$ -scaled Jacobi
- 2 :  $\ell_1$ -scaled block symmetric Gauss-Seidel/SSOR
- 3 : Kaczmarz
- 4 : truncated version of  $\ell_1$ -scaled block symmetric Gauss-Seidel/SSOR
- 16 : Chebyshev

HYPRE\_Int **HYPRE\_AMSSetAlphaAMGOptions**(HYPRE\_Solver solver, HYPRE\_Int alpha\_coarsen\_type,  
HYPRE\_Int alpha\_agg\_levels, HYPRE\_Int  
alpha\_relax\_type, HYPRE\_Real alpha\_strength\_threshold,  
HYPRE\_Int alpha\_interp\_type, HYPRE\_Int alpha\_Pmax)

(Optional) Sets AMG parameters for  $B_{II}$ .

The defaults are 10, 1, 3, 0.25, 0, 0. See the user's manual for more details.

HYPRE\_Int **HYPRE\_AMSSetAlphaAMGCoarseRelaxType**(HYPRE\_Solver solver, HYPRE\_Int  
alpha\_coarse\_relax\_type)

(Optional) Sets the coarsest level relaxation in the AMG solver for  $B_{II}$ .

The default is 8 (11-GS). Use 9, 19, 29 or 99 for a direct solver.

HYPRE\_Int **HYPRE\_AMSSetBetaAMGOptions**(HYPRE\_Solver solver, HYPRE\_Int beta\_coarsen\_type,  
HYPRE\_Int beta\_agg\_levels, HYPRE\_Int beta\_relax\_type,  
HYPRE\_Real beta\_strength\_threshold, HYPRE\_Int  
beta\_interp\_type, HYPRE\_Int beta\_Pmax)

(Optional) Sets AMG parameters for  $B_G$ .

The defaults are 10, 1, 3, 0.25, 0, 0. See the user's manual for more details.

HYPRE\_Int **HYPRE\_AMSSetBetaAMGCoarseRelaxType**(HYPRE\_Solver solver, HYPRE\_Int beta\_coarse\_relax\_type)

(Optional) Sets the coarsest level relaxation in the AMG solver for  $B_G$ .

The default is 8 (I1-GS). Use 9, 19, 29 or 99 for a direct solver.

HYPRE\_Int **HYPRE\_AMSGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)  
Returns the number of iterations taken.

HYPRE\_Int **HYPRE\_AMSGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*rel\_resid\_norm)

Returns the norm of the final relative residual.

HYPRE\_Int **HYPRE\_AMSProjectOutGradients**(HYPRE\_Solver solver, *HYPRE\_ParVector* x)

For problems with zero-conductivity regions, project the vector onto the compatible subspace:  $x = (I - G_0(G_0^t G_0)^{-1} G_0^T)x$ , where  $G_0$  is the discrete gradient restricted to the interior nodes of the regions with zero conductivity.

This ensures that x is orthogonal to the gradients in the range of  $G_0$ .

This function is typically called after the solution iteration is complete, in order to facilitate the visualization of the computed field. Without it the values in the zero-conductivity regions contain kernel components.

HYPRE\_Int **HYPRE\_AMSConstructDiscreteGradient**(*HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* x\_coord, HYPRE\_BigInt \*edge\_vertex, HYPRE\_Int edge\_orientation, *HYPRE\_ParCSRMatrix* \*G)

Construct and return the lowest-order discrete gradient matrix G using some edge and vertex information.

We assume that *edge\_vertex* lists the edge vertices consecutively, and that the orientation of all edges is consistent.

If *edge\_orientation* = 1, the edges are already oriented.

If *edge\_orientation* = 2, the orientation of edge i depends only on the sign of *edge\_vertex*[2\*i+1] - *edge\_vertex*[2\*i].

## ParCSR ADS Solver and Preconditioner

Parallel auxiliary space divergence solver and preconditioner

HYPRE\_Int **HYPRE\_ADSCreate**(HYPRE\_Solver \*solver)

Create an ADS solver object.

HYPRE\_Int **HYPRE\_ADSDestroy**(HYPRE\_Solver solver)

Destroy an ADS solver object.

HYPRE\_Int **HYPRE\_ADSSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Set up the ADS solver or preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

### Parameters

- **solver** – [IN] object to be set up.

- **A** – [IN] ParCSR matrix used to construct the solver/preconditioner.
- **b** – Ignored by this function.
- **x** – Ignored by this function.

HYPRE\_Int **HYPRE\_ADSSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Solve the system or apply ADS as a preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] solver or preconditioner object to be applied.
- **A** – [IN] ParCSR matrix, matrix of the linear system to be solved
- **b** – [IN] right hand side of the linear system to be solved
- **x** – [OUT] approximated solution of the linear system to be solved

HYPRE\_Int **HYPRE\_ADSSetDiscreteCurl**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* C)

Sets the discrete curl matrix *C*.

This function should be called before *HYPRE\_ADSSetup()*!

HYPRE\_Int **HYPRE\_ADSSetDiscreteGradient**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* G)

Sets the discrete gradient matrix *G*.

This function should be called before *HYPRE\_ADSSetup()*!

HYPRE\_Int **HYPRE\_ADSSetCoordinateVectors**(HYPRE\_Solver solver, *HYPRE\_ParVector* x, *HYPRE\_ParVector* y, *HYPRE\_ParVector* z)

Sets the *x*, *y* and *z* coordinates of the vertices in the mesh.

This function should be called before *HYPRE\_ADSSetup()*!

HYPRE\_Int **HYPRE\_ADSSetInterpolations**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* RT\_Pi, *HYPRE\_ParCSRMatrix* RT\_Pix, *HYPRE\_ParCSRMatrix* RT\_Piy, *HYPRE\_ParCSRMatrix* RT\_Piz, *HYPRE\_ParCSRMatrix* ND\_Pi, *HYPRE\_ParCSRMatrix* ND\_Pix, *HYPRE\_ParCSRMatrix* ND\_Piy, *HYPRE\_ParCSRMatrix* ND\_Piz)

(Optional) Set the (components of) the Raviart-Thomas ( $\Pi_{RT}$ ) and the Nedelec ( $\Pi_{ND}$ ) interpolation matrices.

This function is generally intended to be used only for high-order  $H(\text{div})$  discretizations (in the lowest order case, these matrices are constructed internally in ADS from the discrete gradient and curl matrices and the coordinates of the vertices), though it can also be used in the lowest-order case or for other types of discretizations.

By definition,  $RT\_Pi$  and  $ND\_Pi$  are the matrix representations of the linear operators  $\Pi_{RT}$  and  $\Pi_{ND}$  that interpolate (high-order) vector nodal finite elements into the (high-order) Raviart-Thomas and Nedelec spaces. The component matrices are defined in both cases as  $\Pi^x \varphi = \Pi(\varphi, 0, 0)$  and similarly for  $\Pi^y$  and  $\Pi^z$ . Note that all these operators depend on the choice of the basis and degrees of freedom in the high-order spaces.

The column numbering of  $RT\_Pi$  and  $ND\_Pi$  should be node-based, i.e. the *x/ y/ z* components of the first node (vertex or high-order dof) should be listed first, followed by the *x/ y/ z* components of the second node and so on (see the documentation of *HYPRE\_BoomerAMGSetDofFunc*).

If used, this function should be called before `hypr_ADSSetup()` and there is no need to provide the vertex coordinates. Furthermore, only one of the sets  $\{\Pi_{RT}\}$  and  $\{\Pi_{RT}^x, \Pi_{RT}^y, \Pi_{RT}^z\}$  needs to be specified (though it is OK to provide both). If `RT_Pix` is NULL, then scalar  $\Pi$ -based ADS cycles, i.e. those with `cycle_type > 10`, will be unavailable. Similarly, ADS cycles based on monolithic  $\Pi$  (`cycle_type < 10`) require that `RT_Pi` is not NULL. The same restrictions hold for the sets  $\{\Pi_{ND}\}$  and  $\{\Pi_{ND}^x, \Pi_{ND}^y, \Pi_{ND}^z\}$ ; only one of them needs to be specified, and the availability of each enables different AMS cycle type options.

`HYPRE_Int HYPRE_ADSSetMaxIter`(`HYPRE_Solver solver`, `HYPRE_Int maxit`)

(Optional) Sets maximum number of iterations, if ADS is used as a solver.

To use ADS as a preconditioner, set the maximum number of iterations to 1. The default is 20.

`HYPRE_Int HYPRE_ADSSetTol`(`HYPRE_Solver solver`, `HYPRE_Real tol`)

(Optional) Set the convergence tolerance, if ADS is used as a solver.

When using ADS as a preconditioner, set the tolerance to 0.0. The default is  $10^{-6}$ .

`HYPRE_Int HYPRE_ADSSetCycleType`(`HYPRE_Solver solver`, `HYPRE_Int cycle_type`)

(Optional) Choose which auxiliary-space solver to use.

Possible values are:

- 1 : 3-level multiplicative solver (01210)
- 2 : 3-level additive solver (0+1+2)
- 3 : 3-level multiplicative solver (02120)
- 4 : 3-level additive solver (010+2)
- 5 : 3-level multiplicative solver (0102010)
- 6 : 3-level additive solver (1+020)
- 7 : 3-level multiplicative solver (0201020)
- 8 : 3-level additive solver (0(1+2)0)
- 11 : 5-level multiplicative solver (013454310)
- 12 : 5-level additive solver (0+1+3+4+5)
- 13 : 5-level multiplicative solver (034515430)
- 14 : 5-level additive solver (01(3+4+5)10)

The default is 1. See the user's manual for more details.

`HYPRE_Int HYPRE_ADSSetPrintLevel`(`HYPRE_Solver solver`, `HYPRE_Int print_level`)

(Optional) Control how much information is printed during the solution iterations.

The default is 1 (print residual norm at each step).

`HYPRE_Int HYPRE_ADSSetSmoothingOptions`(`HYPRE_Solver solver`, `HYPRE_Int relax_type`,  
`HYPRE_Int relax_times`, `HYPRE_Real relax_weight`,  
`HYPRE_Real omega`)

(Optional) Sets relaxation parameters for  $A$ .

The defaults are 2, 1, 1.0, 1.0.

The available options for `relax_type` are:

- 1 :  $\ell_1$ -scaled Jacobi
- 2 :  $\ell_1$ -scaled block symmetric Gauss-Seidel/SSOR
- 3 : Kaczmarz
- 4 : truncated version of  $\ell_1$ -scaled block symmetric Gauss-Seidel/SSOR
- 16 : Chebyshev

HYPRE\_Int **HYPRE\_ADSSetChebySmoothingOptions**(HYPRE\_Solver solver, HYPRE\_Int cheby\_order, HYPRE\_Real cheby\_fraction)

(Optional) Sets parameters for Chebyshev relaxation.

The defaults are 2, 0.3.

HYPRE\_Int **HYPRE\_ADSSetAMSOptions**(HYPRE\_Solver solver, HYPRE\_Int cycle\_type, HYPRE\_Int coarsen\_type, HYPRE\_Int agg\_levels, HYPRE\_Int relax\_type, HYPRE\_Real strength\_threshold, HYPRE\_Int interp\_type, HYPRE\_Int Pmax)

(Optional) Sets AMS parameters for  $B_C$ .

The defaults are 11, 10, 1, 3, 0.25, 0, 0. Note that *cycle\_type* should be greater than 10, unless the high-order interface of HYPRE\_ADSSetInterpolations is being used! See the user's manual for more details.

HYPRE\_Int **HYPRE\_ADSSetAMGOptions**(HYPRE\_Solver solver, HYPRE\_Int coarsen\_type, HYPRE\_Int agg\_levels, HYPRE\_Int relax\_type, HYPRE\_Real strength\_threshold, HYPRE\_Int interp\_type, HYPRE\_Int Pmax)

(Optional) Sets AMG parameters for  $B_{II}$ .

The defaults are 10, 1, 3, 0.25, 0, 0. See the user's manual for more details.

HYPRE\_Int **HYPRE\_ADSSetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

Returns the number of iterations taken.

HYPRE\_Int **HYPRE\_ADSSetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*rel\_resid\_norm)

Returns the norm of the final relative residual.

## ParCSR PCG Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_ParCSRPCGCreate**(MPI\_Comm comm, HYPRE\_Solver \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_ParCSRPCGDestroy**(HYPRE\_Solver solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_ParCSRPCGSetup**(HYPRE\_Solver solver, HYPRE\_ParCSRMatrix A, HYPRE\_ParVector b, HYPRE\_ParVector x)

HYPRE\_Int **HYPRE\_ParCSRPCGSolve**(HYPRE\_Solver solver, HYPRE\_ParCSRMatrix A, HYPRE\_ParVector b, HYPRE\_ParVector x)

HYPRE\_Int **HYPRE\_ParCSRPCGSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_ParCSRPCGSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_ParCSRPCGSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

HYPRE\_Int **HYPRE\_ParCSRPCGSetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int stop\_crit)

HYPRE\_Int **HYPRE\_ParCSRPCGSetTwoNorm**(HYPRE\_Solver solver, HYPRE\_Int two\_norm)

HYPRE\_Int **HYPRE\_ParCSRPCGSetRelChange**(HYPRE\_Solver solver, HYPRE\_Int rel\_change)

HYPRE\_Int **HYPRE\_ParCSRPCGSetPrecond**(HYPRE\_Solver solver, *HYPRE\_PtrToParSolverFcn* precond, *HYPRE\_PtrToParSolverFcn* precond\_setup, HYPRE\_Solver precond\_solver)

HYPRE\_Int **HYPRE\_ParCSRPCGSetPreconditioner**(HYPRE\_Solver solver, HYPRE\_Solver precond)

HYPRE\_Int **HYPRE\_ParCSRPCGGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data)

HYPRE\_Int **HYPRE\_ParCSRPCGSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

HYPRE\_Int **HYPRE\_ParCSRPCGSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

HYPRE\_Int **HYPRE\_ParCSRPCGGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int **HYPRE\_ParCSRPCGGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

HYPRE\_Int **HYPRE\_ParCSRPCGGetResidual**(HYPRE\_Solver solver, *HYPRE\_ParVector* \*residual)  
Returns the residual.

HYPRE\_Int **HYPRE\_ParCSRDiagScaleSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* y, *HYPRE\_ParVector* x)  
Setup routine for diagonal preconditioning.

HYPRE\_Int **HYPRE\_ParCSRDiagScale**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* HA, *HYPRE\_ParVector* Hy, *HYPRE\_ParVector* Hx)  
Solve routine for diagonal preconditioning.

HYPRE\_Int **HYPRE\_ParCSROnProcTriSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* HA, *HYPRE\_ParVector* Hy, *HYPRE\_ParVector* Hx)  
Setup routine for on-processor triangular solve as preconditioning.

HYPRE\_Int **HYPRE\_ParCSROnProcTriSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* HA, *HYPRE\_ParVector* Hy, *HYPRE\_ParVector* Hx)  
Solve routine for on-processor triangular solve as preconditioning.

### ParCSR GMRES Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_ParCSRGMRESCreate**(MPI\_Comm comm, HYPRE\_Solver \*solver)  
Create a solver object.

HYPRE\_Int **HYPRE\_ParCSRGMRESDestroy**(HYPRE\_Solver solver)  
Destroy a solver object.

HYPRE\_Int **HYPRE\_ParCSRGMRESSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRGMRESSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRGMRESSetRefSolution**(HYPRE\_Solver solver, *HYPRE\_ParVector* ref\_solution)

HYPRE\_Int **HYPRE\_ParCSRGMRESGetRefSolution**(HYPRE\_Solver solver, *HYPRE\_ParVector* \*ref\_solution)

HYPRE\_Int **HYPRE\_ParCSRGMRESSetKDim**(HYPRE\_Solver solver, HYPRE\_Int k\_dim)

HYPRE\_Int **HYPRE\_ParCSRGMRESSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_ParCSRGMRESSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real a\_tol)

HYPRE\_Int **HYPRE\_ParCSRGMRESSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_ParCSRGMRESSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

HYPRE\_Int **HYPRE\_ParCSRGMRESSetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int stop\_crit)

HYPRE\_Int **HYPRE\_ParCSRGMRESSetPrecond**(HYPRE\_Solver solver, *HYPRE\_PtrToParSolverFcn* precond, *HYPRE\_PtrToParSolverFcn* precond\_setup, HYPRE\_Solver precond\_solver)

HYPRE\_Int **HYPRE\_ParCSRGMRESGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data)

HYPRE\_Int **HYPRE\_ParCSRGMRESSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

HYPRE\_Int **HYPRE\_ParCSRGMRESSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

HYPRE\_Int **HYPRE\_ParCSRGMRESGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int **HYPRE\_ParCSRGMRESGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

HYPRE\_Int **HYPRE\_ParCSRGMRESGetResidual**(HYPRE\_Solver solver, *HYPRE\_ParVector* \*residual)  
Returns the residual.

HYPRE\_Int **HYPRE\_ParCSRCOGMRESCreate**(MPI\_Comm comm, HYPRE\_Solver \*solver)  
Create a solver object.

HYPRE\_Int **HYPRE\_ParCSRCOGMRESDestroy**(HYPRE\_Solver solver)  
Destroy a solver object.

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSetKDim**(HYPRE\_Solver solver, HYPRE\_Int k\_dim)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSetUnroll**(HYPRE\_Solver solver, HYPRE\_Int unroll)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSetCGS**(HYPRE\_Solver solver, HYPRE\_Int cgs)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real a\_tol)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSetPrecond**(HYPRE\_Solver solver, *HYPRE\_PtrToParSolverFcn* precondition, *HYPRE\_PtrToParSolverFcn* precondition\_setup, HYPRE\_Solver precondition\_solver)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

HYPRE\_Int **HYPRE\_ParCSRCOGMRESGetResidual**(HYPRE\_Solver solver, *HYPRE\_ParVector* \*residual)

Returns the residual.

### ParCSR FlexGMRES Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESCreate**(MPI\_Comm comm, HYPRE\_Solver \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESDestroy**(HYPRE\_Solver solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESSetKDim**(HYPRE\_Solver solver, HYPRE\_Int k\_dim)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real a\_tol)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESSetPrecond**(HYPRE\_Solver solver, *HYPRE\_PtrToParSolverFcn* precondition, *HYPRE\_PtrToParSolverFcn* precondition\_setup, HYPRE\_Solver precondition\_solver)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESGetResidual**(HYPRE\_Solver solver, *HYPRE\_ParVector* \*residual)

HYPRE\_Int **HYPRE\_ParCSRFlexGMRESSetModifyPC**(HYPRE\_Solver solver, *HYPRE\_PtrToModifyPCFcn* modify\_pc)

### ParCSR LGMRES Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_ParCSRLGMRESCreate**(MPI\_Comm comm, HYPRE\_Solver \*solver)  
Create a solver object.

HYPRE\_Int **HYPRE\_ParCSRLGMRESDestroy**(HYPRE\_Solver solver)  
Destroy a solver object.

HYPRE\_Int **HYPRE\_ParCSRLGMRESSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRLGMRESSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRLGMRESSetKDim**(HYPRE\_Solver solver, HYPRE\_Int k\_dim)

HYPRE\_Int **HYPRE\_ParCSRLGMRESSetAugDim**(HYPRE\_Solver solver, HYPRE\_Int aug\_dim)

HYPRE\_Int **HYPRE\_ParCSRLGMRESSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_ParCSRLGMRESSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real a\_tol)

HYPRE\_Int **HYPRE\_ParCSRLGMRESSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_ParCSRLGMRESSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

HYPRE\_Int **HYPRE\_ParCSRLGMRESSetPrecond**(HYPRE\_Solver solver, *HYPRE\_PtrToParSolverFcn* precondition, *HYPRE\_PtrToParSolverFcn* precondition\_setup, HYPRE\_Solver precondition\_solver)

HYPRE\_Int **HYPRE\_ParCSRLGMRESGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data)

HYPRE\_Int **HYPRE\_ParCSRLGMRESSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

HYPRE\_Int **HYPRE\_ParCSRLGMRESSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

HYPRE\_Int **HYPRE\_ParCSRLGMRESGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

HYPRE\_Int **HYPRE\_ParCSRLGMRESGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

HYPRE\_Int **HYPRE\_ParCSRLGMRESGetResidual**(HYPRE\_Solver solver, *HYPRE\_ParVector* \*residual)

### ParCSR BiCGSTAB Solver

These routines should be used in conjunction with the generic interface in *Krylov Solvers*.

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABCreate**(MPI\_Comm comm, HYPRE\_Solver \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABDestroy**(HYPRE\_Solver solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A,  
*HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A,  
*HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real a\_tol)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABSetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int stop\_crit)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABSetPrecond**(HYPRE\_Solver solver, *HYPRE\_PtrToParSolverFcn*  
precond, *HYPRE\_PtrToParSolverFcn* precondition\_setup,  
HYPRE\_Solver precondition\_solver)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int  
\*num\_iterations)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABGetFinalRelativeResidualNorm**(HYPRE\_Solver solver,  
HYPRE\_Real \*norm)

HYPRE\_Int **HYPRE\_ParCSRBiCGSTABGetResidual**(HYPRE\_Solver solver, *HYPRE\_ParVector* \*residual)

### ParCSR Hybrid Solver

HYPRE\_Int **HYPRE\_ParCSRHybridCreate**(HYPRE\_Solver \*solver)

Create solver object.

HYPRE\_Int **HYPRE\_ParCSRHybridDestroy**(HYPRE\_Solver solver)

Destroy solver object.

HYPRE\_Int **HYPRE\_ParCSRHybridSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A,  
*HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Setup the hybrid solver.

#### Parameters

- **solver** – [IN] object to be set up.

- **A** – [IN] ParCSR matrix used to construct the solver/preconditioner.
- **b** – Ignored by this function.
- **x** – Ignored by this function.

HYPRE\_Int **HYPRE\_ParCSRHybridSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Solve linear system.

**Parameters**

- **solver** – [IN] solver or preconditioner object to be applied.
- **A** – [IN] ParCSR matrix, matrix of the linear system to be solved
- **b** – [IN] right hand side of the linear system to be solved
- **x** – [OUT] approximated solution of the linear system to be solved

HYPRE\_Int **HYPRE\_ParCSRHybridSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

Set the convergence tolerance for the Krylov solver.

The default is 1.e-6.

HYPRE\_Int **HYPRE\_ParCSRHybridSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

Set the absolute convergence tolerance for the Krylov solver.

The default is 0.

HYPRE\_Int **HYPRE\_ParCSRHybridSetConvergenceTol**(HYPRE\_Solver solver, HYPRE\_Real cf\_tol)

Set the desired convergence factor.

HYPRE\_Int **HYPRE\_ParCSRHybridSetDSCGMaxIter**(HYPRE\_Solver solver, HYPRE\_Int dscg\_max\_its)

Set the maximal number of iterations for the diagonally preconditioned solver.

HYPRE\_Int **HYPRE\_ParCSRHybridSetPCGMaxIter**(HYPRE\_Solver solver, HYPRE\_Int pcg\_max\_its)

Set the maximal number of iterations for the AMG preconditioned solver.

HYPRE\_Int **HYPRE\_ParCSRHybridSetSetupType**(HYPRE\_Solver solver, HYPRE\_Int setup\_type)

HYPRE\_Int **HYPRE\_ParCSRHybridSetSolverType**(HYPRE\_Solver solver, HYPRE\_Int solver\_type)

Set the desired solver type.

There are the following options:

- 1 : PCG (default)
- 2 : GMRES
- 3 : BiCGSTAB

HYPRE\_Int **HYPRE\_ParCSRHybridSetRecomputeResidual**(HYPRE\_Solver solver, HYPRE\_Int recompute\_residual)

(Optional) Set recompute residual (don't rely on 3-term recurrence).

HYPRE\_Int **HYPRE\_ParCSRHybridGetRecomputeResidual**(HYPRE\_Solver solver, HYPRE\_Int \*recompute\_residual)

(Optional) Get recompute residual option.

HYPRE\_Int **HYPRE\_ParCSRHybridSetRecomputeResidualP**(HYPRE\_Solver solver, HYPRE\_Int  
recompute\_residual\_p)

(Optional) Set recompute residual period (don't rely on 3-term recurrence).

Recomputes residual after every specified number of iterations.

HYPRE\_Int **HYPRE\_ParCSRHybridGetRecomputeResidualP**(HYPRE\_Solver solver, HYPRE\_Int  
\*recompute\_residual\_p)

(Optional) Get recompute residual period option.

HYPRE\_Int **HYPRE\_ParCSRHybridSetKDim**(HYPRE\_Solver solver, HYPRE\_Int k\_dim)

Set the Krylov dimension for restarted GMRES.

The default is 5.

HYPRE\_Int **HYPRE\_ParCSRHybridSetTwoNorm**(HYPRE\_Solver solver, HYPRE\_Int two\_norm)

Set the type of norm for PCG.

HYPRE\_Int **HYPRE\_ParCSRHybridSetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int stop\_crit)  
RE-VISIT.

HYPRE\_Int **HYPRE\_ParCSRHybridSetRelChange**(HYPRE\_Solver solver, HYPRE\_Int rel\_change)

HYPRE\_Int **HYPRE\_ParCSRHybridSetPrecond**(HYPRE\_Solver solver, *HYPRE\_PtrToParSolverFcn*  
precond, *HYPRE\_PtrToParSolverFcn* precondition\_setup,  
HYPRE\_Solver precondition\_solver)

Set preconditioner if wanting to use one that is not set up by the hybrid solver.

HYPRE\_Int **HYPRE\_ParCSRHybridSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

Set logging parameter (default: 0, no logging).

HYPRE\_Int **HYPRE\_ParCSRHybridSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

Set print level (default: 0, no printing) 2 will print residual norms per iteration 10 will print AMG setup information if AMG is used 12 both Setup information and iterations.

HYPRE\_Int **HYPRE\_ParCSRHybridSetStrongThreshold**(HYPRE\_Solver solver, HYPRE\_Real  
strong\_threshold)

(Optional) Sets AMG strength threshold.

The default is 0.25. For elasticity problems, a larger strength threshold, such as 0.7 or 0.8, is often better.

HYPRE\_Int **HYPRE\_ParCSRHybridSetMaxRowSum**(HYPRE\_Solver solver, HYPRE\_Real max\_row\_sum)

(Optional) Sets a parameter to modify the definition of strength for diagonal dominant portions of the matrix.

The default is 0.9. If *max\_row\_sum* is 1, no checking for diagonally dominant rows is performed.

HYPRE\_Int **HYPRE\_ParCSRHybridSetTruncFactor**(HYPRE\_Solver solver, HYPRE\_Real trunc\_factor)

(Optional) Defines a truncation factor for the interpolation.

The default is 0.

HYPRE\_Int **HYPRE\_ParCSRHybridSetPMaxElmts**(HYPRE\_Solver solver, HYPRE\_Int P\_max\_elmts)

(Optional) Defines the maximal number of elements per row for the interpolation.

The default is 0.

HYPRE\_Int **HYPRE\_ParCSRHybridSetMaxLevels**(HYPRE\_Solver solver, HYPRE\_Int max\_levels)

(Optional) Defines the maximal number of levels used for AMG.

The default is 25.

HYPRE\_Int **HYPRE\_ParCSRHybridSetMeasureType**(HYPRE\_Solver solver, HYPRE\_Int measure\_type)

(Optional) Defines whether local or global measures are used.

HYPRE\_Int **HYPRE\_ParCSRHybridSetCoarsenType**(HYPRE\_Solver solver, HYPRE\_Int coarsen\_type)

(Optional) Defines which parallel coarsening algorithm is used.

There are the following options for *coarsen\_type*:

- 0 : CLJP-coarsening (a parallel coarsening algorithm using independent sets).
- 1 : classical Ruge-Stueben coarsening on each processor, no boundary treatment
- 3 : classical Ruge-Stueben coarsening on each processor, followed by a third pass, which adds coarse points on the boundaries
- 6 : Falgout coarsening (uses 1 first, followed by CLJP using the interior coarse points generated by 1 as its first independent set)
- 7 : CLJP-coarsening (using a fixed random vector, for debugging purposes only)
- 8 : PMIS-coarsening (a parallel coarsening algorithm using independent sets with lower complexities than CLJP, might also lead to slower convergence)
- 9 : PMIS-coarsening (using a fixed random vector, for debugging purposes only)
- 10 : HMIS-coarsening (uses one pass Ruge-Stueben on each processor independently, followed by PMIS using the interior C-points as its first independent set)
- 11 : one-pass Ruge-Stueben coarsening on each processor, no boundary treatment

The default is 10.

HYPRE\_Int **HYPRE\_ParCSRHybridSetInterpType**(HYPRE\_Solver solver, HYPRE\_Int interp\_type)

(Optional) Specifies which interpolation operator is used The default is ext+i interpolation truncated to at most 4 elements per row.

HYPRE\_Int **HYPRE\_ParCSRHybridSetCycleType**(HYPRE\_Solver solver, HYPRE\_Int cycle\_type)

(Optional) Defines the type of cycle.

For a V-cycle, set *cycle\_type* to 1, for a W-cycle set *cycle\_type* to 2. The default is 1.

HYPRE\_Int **HYPRE\_ParCSRHybridSetGridRelaxType**(HYPRE\_Solver solver, HYPRE\_Int  
\*grid\_relax\_type)

HYPRE\_Int **HYPRE\_ParCSRHybridSetGridRelaxPoints**(HYPRE\_Solver solver, HYPRE\_Int  
\*\*grid\_relax\_points)

HYPRE\_Int **HYPRE\_ParCSRHybridSetNumSweeps**(HYPRE\_Solver solver, HYPRE\_Int num\_sweeps)

(Optional) Sets the number of sweeps.

On the finest level, the up and the down cycle the number of sweeps are set to *num\_sweeps* and on the coarsest level to 1. The default is 1.

HYPRE\_Int **HYPRE\_ParCSRHybridSetCycleNumSweeps**(HYPRE\_Solver solver, HYPRE\_Int num\_sweeps, HYPRE\_Int k)

(Optional) Sets the number of sweeps at a specified cycle.

There are the following options for *k*:

- 1 : the down cycle
- 2 : the up cycle
- 3 : the coarsest level

HYPRE\_Int **HYPRE\_ParCSRHybridSetRelaxType**(HYPRE\_Solver solver, HYPRE\_Int relax\_type)

(Optional) Defines the smoother to be used.

It uses the given smoother on the fine grid, the up and the down cycle and sets the solver on the coarsest level to Gaussian elimination (9). The default is 11-Gauss-Seidel, forward solve on the down cycle (13) and backward solve on the up cycle (14).

There are the following options for *relax\_type*:

- 0 : Jacobi
- 1 : Gauss-Seidel, sequential (very slow!)
- 2 : Gauss-Seidel, interior points in parallel, boundary sequential (slow!)
- 3 : hybrid Gauss-Seidel or SOR, forward solve
- 4 : hybrid Gauss-Seidel or SOR, backward solve
- 6 : hybrid symmetric Gauss-Seidel or SSOR
- 8 : hybrid symmetric 11-Gauss-Seidel or SSOR
- 13 : 11-Gauss-Seidel, forward solve
- 14 : 11-Gauss-Seidel, backward solve
- 18 : 11-Jacobi
- 9 : Gaussian elimination (only on coarsest level)

HYPRE\_Int **HYPRE\_ParCSRHybridSetCycleRelaxType**(HYPRE\_Solver solver, HYPRE\_Int relax\_type, HYPRE\_Int k)

(Optional) Defines the smoother at a given cycle.

For options of *relax\_type* see description of HYPRE\_BoomerAMGSetRelaxType). Options for *k* are

- 1 : the down cycle
- 2 : the up cycle
- 3 : the coarsest level

HYPRE\_Int **HYPRE\_ParCSRHybridSetRelaxOrder**(HYPRE\_Solver solver, HYPRE\_Int relax\_order)

(Optional) Defines in which order the points are relaxed.

There are the following options for *relax\_order*:

- 0 : the points are relaxed in natural or lexicographic order on each processor
- 1 : CF-relaxation is used, i.e on the fine grid and the down cycle the coarse points are relaxed first, followed by the fine points; on the up cycle the F-points are relaxed first, followed by the C-points. On the coarsest level, if an iterative scheme is used, the points are relaxed in lexicographic order.

The default is 0 (CF-relaxation).

HYPRE\_Int **HYPRE\_ParCSRHybridSetRelaxWt**(HYPRE\_Solver solver, HYPRE\_Real relax\_wt)

(Optional) Defines the relaxation weight for smoothed Jacobi and hybrid SOR on all levels.

Values for *relax\_wt* are

- > 0 : this assigns the given relaxation weight on all levels
- = 0 : the weight is determined on each level with the estimate  $\frac{3}{4\|D^{-1/2}AD^{-1/2}\|}$ , where  $D$  is the diagonal of  $A$  (this should only be used with Jacobi)
- = -k : the relaxation weight is determined with at most k CG steps on each level (this should only be used for symmetric positive definite problems)

The default is 1.

HYPRE\_Int **HYPRE\_ParCSRHybridSetLevelRelaxWt**(HYPRE\_Solver solver, HYPRE\_Real relax\_wt, HYPRE\_Int level)

(Optional) Defines the relaxation weight for smoothed Jacobi and hybrid SOR on the user defined level.

Note that the finest level is denoted 0, the next coarser level 1, etc. For nonpositive *relax\_weight*, the parameter is determined on the given level as described for HYPRE\_BoomerAMGSetRelaxWt. The default is 1.

HYPRE\_Int **HYPRE\_ParCSRHybridSetOuterWt**(HYPRE\_Solver solver, HYPRE\_Real outer\_wt)

(Optional) Defines the outer relaxation weight for hybrid SOR and SSOR on all levels.

Values for *outer\_wt* are

- > 0 : this assigns the same outer relaxation weight omega on each level
- = -k : an outer relaxation weight is determined with at most k CG steps on each level (this only makes sense for symmetric positive definite problems and smoothers such as SSOR)

The default is 1.

HYPRE\_Int **HYPRE\_ParCSRHybridSetLevelOuterWt**(HYPRE\_Solver solver, HYPRE\_Real outer\_wt, HYPRE\_Int level)

(Optional) Defines the outer relaxation weight for hybrid SOR or SSOR on the user defined level.

Note that the finest level is denoted 0, the next coarser level 1, etc. For nonpositive omega, the parameter is determined on the given level as described for HYPRE\_BoomerAMGSetOuterWt. The default is 1.

HYPRE\_Int **HYPRE\_ParCSRHybridSetMaxCoarseSize**(HYPRE\_Solver solver, HYPRE\_Int max\_coarse\_size)

(Optional) Defines the maximal coarse grid size.

The default is 9.

HYPRE\_Int **HYPRE\_ParCSRHybridSetMinCoarseSize**(HYPRE\_Solver solver, HYPRE\_Int min\_coarse\_size)

(Optional) Defines the minimal coarse grid size.

The default is 0.

HYPRE\_Int **HYPRE\_ParCSRHybridSetSeqThreshold**(HYPRE\_Solver solver, HYPRE\_Int seq\_threshold)  
(Optional) enables redundant coarse grid size.

If the system size becomes smaller than seq\_threshold, sequential AMG is used on all remaining processors. The default is 0.

HYPRE\_Int **HYPRE\_ParCSRHybridSetRelaxWeight**(HYPRE\_Solver solver, HYPRE\_Real \*relax\_weight)

HYPRE\_Int **HYPRE\_ParCSRHybridSetOmega**(HYPRE\_Solver solver, HYPRE\_Real \*omega)

HYPRE\_Int **HYPRE\_ParCSRHybridSetFlexibleCycleStruct**(HYPRE\_Solver solver, HYPRE\_Int  
\*cycle\_struct\_flexible)

(Optional) Defines a flexible cycle structure for the BoomerAMG cycle.

Cycle structure is defined by an array of integers with values:

-2 : terminate the cycle -1 : move to the next coarser level 0 : remain on the current level 1 : move to the next finer level

HYPRE\_Int **HYPRE\_ParCSRHybridSetFlexibleRelaxTypes**(HYPRE\_Solver solver, HYPRE\_Int  
\*relax\_types\_flexible)

(Optional) Defines the relaxation types used during a flexible cycle.

Different relaxation methods may be used at each point in the cycle. This must be used in conjunction with a cycle structure defined by HYPRE\_ParCSRHybridSetFlexibleCycleStruct. The array of relax types passed must have the same length as the array defining the cycle structure. See HYPRE\_BoomerAMGSetRelaxType for possible relax type values.

HYPRE\_Int **HYPRE\_ParCSRHybridSetFlexibleRelaxOrders**(HYPRE\_Solver solver, HYPRE\_Int  
\*relax\_orders\_flexible)

(Optional) Defines the relax orders (lexicographical or CF) used during a flexible cycle.

Different relax orders may be used at each point in the cycle. This must be used in conjunction with a cycle structure defined by HYPRE\_ParCSRHybridSetFlexibleCycleStruct. The array of relax orders passed must have the same length as the array defining the cycle structure. See also HYPRE\_BoomerAMGSetRelaxOrder.

HYPRE\_Int **HYPRE\_ParCSRHybridSetFlexibleRelaxWeights**(HYPRE\_Solver solver, HYPRE\_Real  
\*relax\_weights\_flexible)

(Optional) Defines the relaxation weights used during a flexible cycle.

Different relaxation weights may be used at each point in the cycle. This must be used in conjunction with a cycle structure defined by HYPRE\_ParCSRHybridSetFlexibleCycleStruct. The array of relax weights passed must have the same length as the array defining the cycle structure. See also HYPRE\_BoomerAMGSetRelaxWeight.

HYPRE\_Int **HYPRE\_ParCSRHybridSetFlexibleOuterWeights**(HYPRE\_Solver solver, HYPRE\_Real  
\*outer\_weights\_flexible)

(Optional) Defines the outer relaxation weights used during a flexible cycle.

Different outer weights may be used at each point in the cycle. This must be used in conjunction with a cycle structure defined by HYPRE\_ParCSRHybridSetFlexibleCycleStruct. The array of outer weights passed must have the same length as the array defining the cycle structure. See also HYPRE\_BoomerAMGSetOuterWt.

HYPRE\_Int **HYPRE\_ParCSRHybridSetFlexibleCGCScalingFactors**(HYPRE\_Solver solver,  
HYPRE\_Real  
\*cgc\_scaling\_factors\_flexible)

(Optional) Defines the a scaling factor for the coarse-grid correction during a flexible cycle.

Different scaling factors may be used at each point in the cycle. This must be used in conjunction with a cycle structure defined by `HYPRE_ParCSRHybridSetFlexibleCycleStruct`. The array of coarse- grid correction scaling factors passed must have the same length as the array defining the cycle structure. Note, however, that the scaling factors will only be used when the flexible cycle performs a coarse-grid correction, i.e. when the value of the array passed for the cycle structure is 1.

`HYPRE_Int HYPRE_ParCSRHybridSetAggNumLevels`(`HYPRE_Solver solver`, `HYPRE_Int agg_num_levels`)

(Optional) Defines the number of levels of aggressive coarsening, starting with the finest level.

The default is 0, i.e. no aggressive coarsening.

`HYPRE_Int HYPRE_ParCSRHybridSetAggInterpType`(`HYPRE_Solver solver`, `HYPRE_Int agg_interp_type`)

(Optional) Defines the interpolation used on levels of aggressive coarsening The default is 4, i.e. multipass interpolation. The following options exist:

- 1 : 2-stage extended+i interpolation
- 2 : 2-stage standard interpolation
- 3 : 2-stage extended interpolation
- 4 : multipass interpolation
- 5 : 2-stage extended interpolation in matrix-matrix form
- 6 : 2-stage extended+i interpolation in matrix-matrix form
- 7 : 2-stage extended+e interpolation in matrix-matrix form

`HYPRE_Int HYPRE_ParCSRHybridSetNumPaths`(`HYPRE_Solver solver`, `HYPRE_Int num_paths`)

(Optional) Defines the degree of aggressive coarsening.

The default is 1, which leads to the most aggressive coarsening. Setting `num_paths` to 2 will increase complexity somewhat, but can lead to better convergence.

`HYPRE_Int HYPRE_ParCSRHybridSetNumFunctions`(`HYPRE_Solver solver`, `HYPRE_Int num_functions`)

(Optional) Sets the size of the system of PDEs, if using the systems version.

The default is 1.

`HYPRE_Int HYPRE_ParCSRHybridSetDofFunc`(`HYPRE_Solver solver`, `HYPRE_Int *dof_func`)

(Optional) Sets the mapping that assigns the function to each variable, if using the systems version.

If no assignment is made and the number of functions is  $k > 1$ , the mapping generated is  $(0,1,\dots,k-1,0,1,\dots,k-1,\dots)$ .

`HYPRE_Int HYPRE_ParCSRHybridSetNodal`(`HYPRE_Solver solver`, `HYPRE_Int nodal`)

(Optional) Sets whether to use the nodal systems version.

The default is 0 (the unknown based approach).

`HYPRE_Int HYPRE_ParCSRHybridSetKeepTranspose`(`HYPRE_Solver solver`, `HYPRE_Int keepT`)

(Optional) Sets whether to store local transposed interpolation The default is 0 (don't store).

HYPRE\_Int **HYPRE\_ParCSRHybridSetNonGalerkinTol**(HYPRE\_Solver solver, HYPRE\_Int num\_levels, HYPRE\_Real \*nongalerkin\_tol)

(Optional) Sets whether to use non-Galerkin option The default is no non-Galerkin option num\_levels sets the number of levels where to use it nongalerkin\_tol contains the tolerances for <num\_levels> levels

HYPRE\_Int **HYPRE\_ParCSRHybridGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_its)

Retrieves the total number of iterations.

HYPRE\_Int **HYPRE\_ParCSRHybridGetDSCGNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*dscg\_num\_its)

Retrieves the number of iterations used by the diagonally scaled solver.

HYPRE\_Int **HYPRE\_ParCSRHybridGetPCGNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*pcg\_num\_its)

Retrieves the number of iterations used by the AMG preconditioned solver.

HYPRE\_Int **HYPRE\_ParCSRHybridGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

Retrieves the final relative residual norm.

HYPRE\_Int **HYPRE\_ParCSRHybridSetNumGridSweeps**(HYPRE\_Solver solver, HYPRE\_Int \*num\_grid\_sweeps)

HYPRE\_Int **HYPRE\_ParCSRHybridGetSetupSolveTime**(HYPRE\_Solver solver, HYPRE\_Real \*time)

### ParCSR MGR Solver

Parallel multigrid reduction solver and preconditioner.

This solver or preconditioner is designed with systems of PDEs in mind. However, it can also be used for scalar linear systems, particularly for problems where the user can exploit information from the physics of the problem. In this way, the MGR solver could potentially be used as a foundation for a physics-based preconditioner.

HYPRE\_Int **HYPRE\_MGRCreate**(HYPRE\_Solver \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_MGRDestroy**(HYPRE\_Solver solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_MGRSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Setup the MGR solver or preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] object to be set up.
- **A** – [IN] ParCSR matrix used to construct the solver/preconditioner.
- **b** – right-hand-side of the linear system to be solved (Ignored by this function).
- **x** – approximate solution of the linear system to be solved (Ignored by this function).

HYPRE\_Int **HYPRE\_MGRSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Solve the system or apply MGR as a preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

**Parameters**

- **solver** – [IN] solver or preconditioner object to be applied.
- **A** – [IN] ParCSR matrix, matrix of the linear system to be solved
- **b** – [IN] right hand side of the linear system to be solved
- **x** – [OUT] approximated solution of the linear system to be solved

HYPRE\_Int **HYPRE\_MGRSetCpointsByContiguousBlock**(HYPRE\_Solver solver, HYPRE\_Int block\_size, HYPRE\_Int max\_num\_levels, HYPRE\_BigInt \*idx\_array, HYPRE\_Int \*num\_block\_coarse\_points, HYPRE\_Int \*\*block\_coarse\_indexes)

Set the block data assuming that the physical variables are ordered contiguously, i.e.

$p_1, p_2, \dots, p_n, s_1, s_2, \dots, s_n, \dots$

**Parameters**

- **solver** – [IN] solver or preconditioner object
- **block\_size** – [IN] system block size
- **max\_num\_levels** – [IN] maximum number of reduction levels
- **num\_block\_coarse\_points** – [IN] number of coarse points per block per level
- **block\_coarse\_indexes** – [IN] index for each block coarse point per level

HYPRE\_Int **HYPRE\_MGRSetCpointsByBlock**(HYPRE\_Solver solver, HYPRE\_Int block\_size, HYPRE\_Int max\_num\_levels, HYPRE\_Int \*num\_block\_coarse\_points, HYPRE\_Int \*\*block\_coarse\_indexes)

Set the block data (by grid points) and prescribe the coarse indexes per block for each reduction level.

**Parameters**

- **solver** – [IN] solver or preconditioner object
- **block\_size** – [IN] system block size
- **max\_num\_levels** – [IN] maximum number of reduction levels
- **num\_block\_coarse\_points** – [IN] number of coarse points per block per level
- **block\_coarse\_indexes** – [IN] index for each block coarse point per level

HYPRE\_Int **HYPRE\_MGRSetCpointsByPointMarkerArray**(HYPRE\_Solver solver, HYPRE\_Int block\_size, HYPRE\_Int max\_num\_levels, HYPRE\_Int \*num\_block\_coarse\_points, HYPRE\_Int \*\*lvl\_block\_coarse\_indexes, HYPRE\_Int \*point\_marker\_array)

Set the coarse indices for the levels using an array of tags for all the local degrees of freedom.

TODO: Rename the function to make it more descriptive.

**Parameters**

- **solver** – [IN] solver or preconditioner object
- **block\_size** – [IN] system block size
- **max\_num\_levels** – [IN] maximum number of reduction levels
- **num\_block\_coarse\_points** – [IN] number of coarse points per block per level

- **lvl\_block\_coarse\_indexes** – [IN] indices for the coarse points per level
- **point\_marker\_array** – [IN] array of tags for the local degrees of freedom

HYPRE\_Int **HYPRE\_MGRSetNonCpointsToFpoints**(HYPRE\_Solver solver, HYPRE\_Int nonCptToFptFlag)  
 (Optional) Set non C-points to F-points.

This routine determines how the coarse points are selected for the next level reduction. Options for *non-CptToFptFlag* are:

- 0 : Allow points not prescribed as C points to be potentially set as C points using classical AMG coarsening strategies (currently uses CLJP-coarsening).
- 1 : Fix points not prescribed as C points to be F points for the next reduction

HYPRE\_Int **HYPRE\_MGRSetMaxCoarseLevels**(HYPRE\_Solver solver, HYPRE\_Int maxlev)  
 (Optional) Set maximum number of coarsening (or reduction) levels.

The default is 10.

HYPRE\_Int **HYPRE\_MGRSetBlockSize**(HYPRE\_Solver solver, HYPRE\_Int bsize)  
 (Optional) Set the system block size.

This should match the block size set in the MGRSetCpointsByBlock function. The default is 1.

HYPRE\_Int **HYPRE\_MGRSetReservedCoarseNodes**(HYPRE\_Solver solver, HYPRE\_Int reserved\_coarse\_size, HYPRE\_BigInt \*reserved\_coarse\_nodes)

(Optional) Defines indexes of coarse nodes to be kept to the coarsest level.

These indexes are passed down through the MGR hierarchy to the coarsest grid of the coarse grid (Boomer-AMG) solver.

#### Parameters

- **solver** – [IN] solver or preconditioner object
- **reserved\_coarse\_size** – [IN] number of reserved coarse points
- **reserved\_coarse\_nodes** – [IN] (global) indexes of reserved coarse points

HYPRE\_Int **HYPRE\_MGRSetReservedCpointsLevelToKeep**(HYPRE\_Solver solver, HYPRE\_Int level)  
 (Optional) Set the level for reducing the reserved Cpoints before the coarse grid solve.

This is necessary for some applications, such as phase transitions. The default is 0 (no reduction, i.e. keep the reserved cpoints in the coarse grid solve).

The default setup for the reduction is as follows:

- Interpolation type: Jacobi (2)
- Restriction type: Injection (0)
- F-relaxation type: LU factorization with pivoting (99)
- Coarse grid type: galerkin (0)

HYPRE\_Int **HYPRE\_MGRSetRelaxType**(HYPRE\_Solver solver, HYPRE\_Int relax\_type)  
 (Optional) Set the relaxation type for F-relaxation.

Currently supports the following flavors of relaxation types as described in the *BoomerAMGSetRelaxType*: *relax\_type* 0, 3 - 8, 13, 14, 18. Also supports AMG (options 1 and 2) and direct solver variants (9, 99, 199). See *HYPRE\_MGRSetLevelFRelaxType* for details.

HYPRE\_Int **HYPRE\_MGRSetFRelaxMethod**(HYPRE\_Solver solver, HYPRE\_Int relax\_method)

(Optional) Set the strategy for F-relaxation.

Options for *relax\_method* are:

- 0 : Single-level relaxation sweeps for F-relaxation as prescribed by *MGRSetRelaxType*
- 1 : Multi-level relaxation strategy for F-relaxation (V(1,0) cycle currently supported).

NOTE: This function will be removed in favor of *HYPRE\_MGRSetLevelFRelaxType!!*

HYPRE\_Int **HYPRE\_MGRSetLevelFRelaxMethod**(HYPRE\_Solver solver, HYPRE\_Int \*relax\_method)

(Optional) This function is an extension of HYPRE\_MGRSetFRelaxMethod.

It allows setting the F-relaxation strategy for each MGR level.

HYPRE\_Int **HYPRE\_MGRSetLevelFRelaxType**(HYPRE\_Solver solver, HYPRE\_Int \*relax\_type)

(Optional) Set the relaxation type for F-relaxation at each level.

This function takes precedence over, and will replace *HYPRE\_MGRSetFRelaxMethod* and *HYPRE\_MGRSetRelaxType*. Options for *relax\_type* entries are:

- 0, 3 - 8, 13, 14, 18: (as described in *BoomerAMGSetRelaxType*)
- 1 : Multi-level relaxation strategy for F-relaxation (V(1,0) cycle currently supported).
- 2 : AMG
- 29: Sparse direct solver (requires SuperLU\_Dist support)
- 32: ILU
- 9, 99, 199 : Gaussian Elimination variants (GE, GE with pivoting, direct inversion respectively)

HYPRE\_Int **HYPRE\_MGRSetCoarseGridMethod**(HYPRE\_Solver solver, HYPRE\_Int \*cg\_method)

(Optional) Set the strategy for coarse grid computation.

Options for *cg\_method* are:

- 0 : Galerkin coarse grid computation using RAP.
- 1 - 5 : Non-Galerkin coarse grid computation with dropping strategy.
  - 1:  $\text{inv}(A_{FF})$  approximated by its (block) diagonal inverse
  - 2: CPR-like approximation with  $\text{inv}(A_{FF})$  approximated by its diagonal inverse
  - 3: CPR-like approximation with  $\text{inv}(A_{FF})$  approximated by its block diagonal inverse
  - 4:  $\text{inv}(A_{FF})$  approximated by sparse approximate inverse
  - 5:  $\text{inv}(A_{FF})$  is an empty matrix and coarse level matrix is set to  $A_{CC}$

HYPRE\_Int **HYPRE\_MGRSetNonGalerkinMaxElmts**(HYPRE\_Solver solver, HYPRE\_Int max\_elmts)

(Optional) Set the maximum number of nonzeros per row of the coarse grid correction operator computed in the Non-Galerkin approach.

Options for *max\_elmts* are:

- 0: keep only the (block) diagonal portion of the correction matrix (default).
- k > 0: keep the (block) diagonal plus the k-th largest entries per row of the correction matrix.

HYPRE\_Int **HYPRE\_MGRSetLevelNonGalerkinMaxElmts**(HYPRE\_Solver solver, HYPRE\_Int \*max\_elmts)

(Optional) Set the maximum number of nonzeros per row of the coarse grid correction operator computed in the Non-Galerkin approach at each MGR level.

For options, see *HYPRE\_MGRSetNonGalerkinMaxElmts*.

HYPRE\_Int **HYPRE\_MGRSetLevelFRelaxNumFunctions**(HYPRE\_Solver solver, HYPRE\_Int \*num\_functions)

(Optional) Set the number of functions for F-relaxation V-cycle.

For problems like elasticity, one may want to perform coarsening and interpolation for block matrices. The number of functions corresponds to the number of scalar PDEs in the system.

HYPRE\_Int **HYPRE\_MGRSetRestrictType**(HYPRE\_Solver solver, HYPRE\_Int restrict\_type)

(Optional) Set the strategy for computing the MGR restriction operator.

Options for *restrict\_type* are:

- 0 : injection [0I]
- 1 : unscaled (not recommended)
- 2 : diagonal scaling (Jacobi)
- 3 : approximate inverse
- 4 : pAIR distance 1
- 5 : pAIR distance 2
- 12 : Block Jacobi
- 13 : CPR-like restriction operator
- 14 : (Block) Column-lumped restriction
- 15 : partial Column-lumped restriction
- else : use classical modified interpolation

The default is injection.

HYPRE\_Int **HYPRE\_MGRSetLevelRestrictType**(HYPRE\_Solver solver, HYPRE\_Int \*restrict\_type)

(Optional) This function is an extension of *HYPRE\_MGRSetRestrictType*.

It allows setting the restriction operator strategy for each MGR level.

HYPRE\_Int **HYPRE\_MGRSetNumRestrictSweeps**(HYPRE\_Solver solver, HYPRE\_Int nsweeps)

(Optional) Set number of restriction sweeps.

This option is for *restrict\_type* > 2.

HYPRE\_Int **HYPRE\_MGRSetInterpType**(HYPRE\_Solver solver, HYPRE\_Int interp\_type)

(Optional) Set the strategy for computing the MGR interpolation operator.

Options for *interp\_type* are:

- 0 : injection  $[0I]^T$
- 1 : L1-Jacobi
- 2 : diagonal scaling (Jacobi)
- 3 : classical modified interpolation
- 4 : approximate inverse
- 12 : Block Jacobi
- 13 : block row-sum (lumped) with regular sums
- 14 : block row-sum (lumped) with absolute-value sums
- else : classical modified interpolation

The default is diagonal scaling.

HYPRE\_Int **HYPRE\_MGRSetLevelInterpType**(HYPRE\_Solver solver, HYPRE\_Int \*interp\_type)

(Optional) This function is an extension of *HYPRE\_MGRSetInterpType*.

It allows setting the prolongation (interpolation) operator strategy for each MGR level.

HYPRE\_Int **HYPRE\_MGRSetNumRelaxSweeps**(HYPRE\_Solver solver, HYPRE\_Int nsweeps)

(Optional) Set number of relaxation sweeps.

This option is for the “single level” F-relaxation (*relax\_method* = 0).

HYPRE\_Int **HYPRE\_MGRSetLevelNumRelaxSweeps**(HYPRE\_Solver solver, HYPRE\_Int \*nsweeps)

(Optional) This function is an extension of *HYPRE\_MGRSetNumRelaxSweeps*.

It allows setting the number of single-level relaxation sweeps for each MGR level.

HYPRE\_Int **HYPRE\_MGRSetNumInterpSweeps**(HYPRE\_Solver solver, HYPRE\_Int nsweeps)

(Optional) Set number of interpolation sweeps.

This option is for *interp\_type* > 2.

HYPRE\_Int **HYPRE\_MGRSetBlockJacobiBlockSize**(HYPRE\_Solver solver, HYPRE\_Int blk\_size)

(Optional) Set block size for block (global) smoother and interp/restriction.

This option is for *interp\_type/restrict\_type* == 12, and *smooth\_type* == 0 or 1.

HYPRE\_Int **HYPRE\_MGRSetFSolver**(HYPRE\_Solver solver, *HYPRE\_PtrToParSolverFcn*  
fine\_grid\_solver\_solve, *HYPRE\_PtrToParSolverFcn*  
fine\_grid\_solver\_setup, HYPRE\_Solver fsolver)

(Optional) Set the fine grid solver.

#### Parameters

- **solver** – [IN] MGR solver/preconditioner object
- **fine\_grid\_solver\_solve** – [IN] solve routine
- **fine\_grid\_solver\_setup** – [IN] setup routine
- **fine\_grid\_solver** – [IN] fine grid solver object

HYPRE\_Int **HYPRE\_MGRSetFSolverAtLevel**(HYPRE\_Solver solver, HYPRE\_Solver fsolver, HYPRE\_Int level)

(Optional) Set the F-relaxation solver at a given level.

#### Parameters

- **solver** – [IN] MGR solver/preconditioner object
- **fsolver** – [IN] F-relaxation solver object
- **level** – [IN] MGR solver level

HYPRE\_Int **HYPRE\_MGRBuildAff**(HYPRE\_ParCSRMatrix A, HYPRE\_Int \*CF\_marker, HYPRE\_Int debug\_flag, HYPRE\_ParCSRMatrix \*A\_ff)

(Optional) Extract A\_FF block from matrix A.

TODO (VPM): Does this need to be exposed? Move to parcsr\_mv?

HYPRE\_Int **HYPRE\_MGRSetCoarseSolver**(HYPRE\_Solver solver, HYPRE\_PtrToParSolverFcn coarse\_grid\_solver\_solve, HYPRE\_PtrToParSolverFcn coarse\_grid\_solver\_setup, HYPRE\_Solver coarse\_grid\_solver)

(Optional) Set the coarse grid solver.

Currently uses BoomerAMG. The default, if not set, is BoomerAMG with default options.

#### Parameters

- **solver** – [IN] MGR solver/preconditioner object
- **coarse\_grid\_solver\_solve** – [IN] solve routine for BoomerAMG
- **coarse\_grid\_solver\_setup** – [IN] setup routine for BoomerAMG
- **coarse\_grid\_solver** – [IN] coarse grid solver object

HYPRE\_Int **HYPRE\_MGRSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

HYPRE\_Int **HYPRE\_MGRSetFrelaxPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

(Optional) Set the print level of the F-relaxation solver

HYPRE\_Int **HYPRE\_MGRSetCoarseGridPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

(Optional) Set the print level of the coarse grid solver

HYPRE\_Int **HYPRE\_MGRSetTruncateCoarseGridThreshold**(HYPRE\_Solver solver, HYPRE\_Real threshold)

(Optional) Set the threshold for dropping small entries on the coarse grid at each level.

No dropping is applied if *threshold* = 0.0 (default).

HYPRE\_Int **HYPRE\_MGRSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

(Optional) Requests logging of solver diagnostics.

Requests additional computations for diagnostic and similar data to be logged by the user. Default is 0, do nothing. The latest residual will be available if logging > 1.

HYPRE\_Int **HYPRE\_MGRSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations if used as a solver.

Set this to 1 if MGR is used as a preconditioner. The default is 20.

HYPRE\_Int **HYPRE\_MGRSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance for the MGR solver.

Use tol = 0.0 if MGR is used as a preconditioner. The default is 1.e-6.

HYPRE\_Int **HYPRE\_MGRSetMaxGlobalSmoothIters**(HYPRE\_Solver solver, HYPRE\_Int smooth\_iter)

(Optional) Determines how many sweeps of global smoothing to do.

Default is 0 (no global smoothing).

HYPRE\_Int **HYPRE\_MGRSetLevelSmoothIters**(HYPRE\_Solver solver, HYPRE\_Int \*smooth\_iters)

(Optional) Determines how many sweeps of global smoothing to do on each level.

Default is 0 (no global smoothing).

HYPRE\_Int **HYPRE\_MGRSetGlobalSmoothCycle**(HYPRE\_Solver solver, HYPRE\_Int global\_smooth\_cycle)

(Optional) Set the cycle for global smoothing.

Options for *global\_smooth\_cycle* are:

- 0 : no global smoothing
- 1 : pre-smoothing only - down cycle (default)
- 2 : post-smoothing only - up cycle
- 3 : both pre- and post-smoothing (V(1,1))

HYPRE\_Int **HYPRE\_MGRSetFRelaxCycle**(HYPRE\_Solver solver, HYPRE\_Int frelax\_cycle)

(Optional) Controls when F-relaxation sweeps are applied during the cycle.

Options for *frelax\_cycle* are:

- 0 : no F-relaxation
- 1 : pre-relaxation only (default, applied on the way down)
- 2 : post-relaxation only (applied on the way up, after coarse correction)
- 3 : both pre- and post-relaxation

HYPRE\_Int **HYPRE\_MGRSetCycleType**(HYPRE\_Solver solver, HYPRE\_Int cycle\_type)

(Optional) Set MGR's cycling strategy.

Options for *cycle\_type* are:

- 1 : V-cycle (default)
- 2 : W-cycle Note: F-cycles are not supported.

HYPRE\_Int **HYPRE\_MGRSetGlobalSmoothType**(HYPRE\_Solver solver, HYPRE\_Int smooth\_type)

(Optional) Determines type of global smoother.

Options for *smooth\_type* are:

- 0 : block Jacobi (default)
- 1 : block Gauss-Seidel
- 2 : Jacobi
- 3 : Gauss-Seidel, sequential (very slow!)
- 4 : Gauss-Seidel, interior points in parallel, boundary sequential (slow!)
- 5 : hybrid Gauss-Seidel or SOR, forward solve
- 6 : hybrid Gauss-Seidel or SOR, backward solve
- 8 : Euclid (ILU)
- 16 : HYPRE\_ILU
- 18 : L1-Jacobi

HYPRE\_Int **HYPRE\_MGRSetLevelSmoothType**(HYPRE\_Solver solver, HYPRE\_Int \*smooth\_type)

Sets the type of global smoother for each level in the multigrid reduction (MGR) solver.

This function allows the user to specify the type of global smoother to be used at each level of the multigrid reduction process. The types of smoothers available can be found in the documentation for *HYPRE\_MGRSetGlobalSmoothType*. The smoother type for each level is indicated by the *smooth\_type* array, which should have a size equal to *max\_num\_coarse\_levels*.

**Note**

This function does not take ownership of the *smooth\_type* array.

**Note**

If *smooth\_type* is a NULL pointer, a default global smoother (Jacobi) is used for all levels.

**Note**

This call is optional. It is intended for advanced users who need specific control over the smoothing process at different levels of the solver. If not called, the solver will proceed with default smoothing parameters.

**Param**

HYPRE\_Int **HYPRE\_MGRSetGlobalSmootherAtLevel**(HYPRE\_Solver solver, HYPRE\_Solver smoother, HYPRE\_Int level)

Sets the global smoother method for a specified MGR level using a HYPRE solver object.

This function enables solvers within hypr to be used as complex smoothers for a specific level within the multigrid reduction (MGR) scheme. Users can configure the solver options and pass the solver in as the smoother. Currently supported solver options via this interface are ILU and AMG.

**Note**

Unlike some other setup functions that might require an array to set options across multiple levels, this function focuses on a single level, identified by the *level* parameter.

**Warning**

The smoother passed to function takes precedence over the smoother type set for that level in the MGR hierarchy.

**Param**

HYPRE\_Int **HYPRE\_MGRGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

(Optional) Return the number of MGR iterations.

HYPRE\_Int **HYPRE\_MGRGetCoarseGridConvergenceFactor**(HYPRE\_Solver solver, HYPRE\_Real \*conv\_factor)

(Optional) Return the relative residual for the coarse level system.

HYPRE\_Int **HYPRE\_MGRSetPMaxElmts**(HYPRE\_Solver solver, HYPRE\_Int P\_max\_elmts)

(Optional) Set the maximum number of nonzeros per row for interpolation operators.

HYPRE\_Int **HYPRE\_MGRSetLevelPMaxElmts**(HYPRE\_Solver solver, HYPRE\_Int \*P\_max\_elmts)

(Optional) Set the maximum number of nonzeros per row for interpolation operators for each level.

HYPRE\_Int **HYPRE\_MGRGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*res\_norm)

(Optional) Return the norm of the final relative residual.

### ParCSR ILU Solver

(Parallel) Incomplete LU factorization.

HYPRE\_Int **HYPRE\_ILUCreate**(HYPRE\_Solver \*solver)

Create a solver object.

HYPRE\_Int **HYPRE\_ILUDestroy**(HYPRE\_Solver solver)

Destroy a solver object.

HYPRE\_Int **HYPRE\_ILUSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Setup the ILU solver or preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] object to be set up.
- **A** – [IN] ParCSR matrix used to construct the solver/preconditioner.
- **b** – right-hand-side of the linear system to be solved (Ignored by this function).
- **x** – approximate solution of the linear system to be solved (Ignored by this function).

HYPRE\_Int **HYPRE\_ILUSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Solve the system or apply ILU as a preconditioner.

If used as a preconditioner, this function should be passed to the iterative solver *SetPrecond* function.

#### Parameters

- **solver** – [IN] solver or preconditioner object to be applied.
- **A** – [IN] ParCSR matrix, matrix of the linear system to be solved
- **b** – [IN] right hand side of the linear system to be solved
- **x** – [OUT] approximated solution of the linear system to be solved

HYPRE\_Int **HYPRE\_ILUSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations if used as a solver.

Set this to 1 if ILU is used as a preconditioner. The default is 20.

HYPRE\_Int **HYPRE\_ILUSetIterativeSetupType**(HYPRE\_Solver solver, HYPRE\_Int iter\_setup\_type)

(Optional) Set the algorithm type to compute the ILU factorization.

Options are:

- 0 : Non-iterative algorithm (default)
- 1 : Asynchronous with in-place storage (recommended for iterative version)
- 2 : Asynchronous with explicit storage splitting
- 3 : Synchronous with explicit storage splitting
- 4 : Semi-synchronous with explicit storage splitting

Note: Iterative ILU is available only for zero fill-in and it depends on rocSPARSE.

HYPRE\_Int **HYPRE\_ILUSetIterativeSetupOption**(HYPRE\_Solver solver, HYPRE\_Int iter\_setup\_option)

(Optional) Set the compute option for iterative ILU in an additive fashion, i.e.

; multiple options can be turned on by summing their respective numeric codes as given below:

- 2: Use stopping tolerance to finish the algorithm
- 4: Compute correction norms
- 8: Compute residual norms
- 16: Save convergence history
- 32: Use rocSPARSE's internal COO format

The iterative ILU algorithm can terminate based on the maximum number of iterations (default) or a target tolerance (option 2). In the tolerance-based case, the max. number of iterations is still used to terminate the algorithm in case it does not converge to the requested tolerance. In addition, the tolerance-based mode uses residual norms by default (option 8). To use correction norms instead, enable option 4. Lastly, the convergence history for computing the triangular factors can be saved and printed out by enabling option 16.

Note: Iterative ILU is available only for zero fill-in, and it depends on rocSPARSE.

HYPRE\_Int **HYPRE\_ILUSetIterativeSetupMaxIter**(HYPRE\_Solver solver, HYPRE\_Int iter\_setup\_max\_iter)

(Optional) Set the max.

number of iterations for the iterative ILU algorithm.

Note: Iterative ILU is available only for zero fill-in and it depends on rocSPARSE.

HYPRE\_Int **HYPRE\_ILUSetIterativeSetupTolerance**(HYPRE\_Solver solver, HYPRE\_Real iter\_setup\_tolerance)

(Optional) Set the stop tolerance for the iterative ILU algorithm.

Note: Iterative ILU is available only for zero fill-in and it depends on rocSPARSE.

HYPRE\_Int **HYPRE\_ILUSetTriSolve**(HYPRE\_Solver solver, HYPRE\_Int tri\_solve)

(Optional) Set triangular solver type.

Options are:

- 0 : iterative
- 1 : direct (default)

HYPRE\_Int **HYPRE\_ILUSetLowerJacobiIters**(HYPRE\_Solver solver, HYPRE\_Int lower\_jacobi\_iterations)

(Optional) Set number of lower Jacobi iterations for the triangular L solves Set this to integer > 0 when using iterative tri\_solve (0).

The default is 5 iterations.

HYPRE\_Int **HYPRE\_ILUSetUpperJacobiIters**(HYPRE\_Solver solver, HYPRE\_Int upper\_jacobi\_iterations)

(Optional) Set number of upper Jacobi iterations for the triangular U solves Set this to integer > 0 when using iterative tri\_solve (0).

The default is 5 iterations.

HYPRE\_Int **HYPRE\_ILUSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance for ILU.

Use tol = 0.0 if ILU is used as a preconditioner. The default is 1.e-7.

HYPRE\_Int **HYPRE\_ILUSetLevelOfFill**(HYPRE\_Solver solver, HYPRE\_Int lfill)

(Optional) Set the level of fill k, for level-based ILU(k) The default is 0 (for ILU(0)).

HYPRE\_Int **HYPRE\_ILUSetMaxNnzPerRow**(HYPRE\_Solver solver, HYPRE\_Int nzmax)

(Optional) Set the max non-zeros per row in L and U factors (for ILUT) The default is 1000.

HYPRE\_Int **HYPRE\_ILUSetDropThreshold**(HYPRE\_Solver solver, HYPRE\_Real threshold)

(Optional) Set the threshold for dropping in L and U factors (for ILUT).

Any fill-in less than this threshold is dropped in the factorization. The default is 1.0e-2.

HYPRE\_Int **HYPRE\_ILUSetDropThresholdArray**(HYPRE\_Solver solver, HYPRE\_Real \*threshold)

(Optional) Set the array of thresholds for dropping in ILUT.

B, E, and F correspond to upper left, lower left and upper right of 2 x 2 block decomposition respectively. Any fill-in less than threshold is dropped in the factorization.

- threshold[0] : threshold for matrix B.
- threshold[1] : threshold for matrix E and F.
- threshold[2] : threshold for matrix S (Schur Complement). The default is 1.0e-2.

HYPRE\_Int **HYPRE\_ILUSetNSHDropThreshold**(HYPRE\_Solver solver, HYPRE\_Real threshold)

(Optional) Set the threshold for dropping in Newton–Schulz–Hotelling iteration (NSH-ILU).

Any entries less than this threshold are dropped when forming the approximate inverse matrix. The default is 1.0e-2.

HYPRE\_Int **HYPRE\_ILUSetNSHDropThresholdArray**(HYPRE\_Solver solver, HYPRE\_Real \*threshold)

(Optional) Set the array of thresholds for dropping in Newton–Schulz–Hotelling iteration (for NSH-ILU).

Any fill-in less than thresholds is dropped when forming the approximate inverse matrix.

- threshold[0] : threshold for Minimal Residual iteration (initial guess for NSH).
- threshold[1] : threshold for Newton–Schulz–Hotelling iteration.

The default is 1.0e-2.

HYPRE\_Int **HYPRE\_ILUSetSchurMaxIter**(HYPRE\_Solver solver, HYPRE\_Int ss\_max\_iter)

(Optional) Set maximum number of iterations for Schur System Solve.

For GMRES-ILU, this is the maximum number of iterations for GMRES. The Krylov dimension for GMRES is set equal to this value to avoid restart. For NSH-ILU, this is the maximum number of iterations for NSH solve. The default is 5.

HYPRE\_Int **HYPRE\_ILUSetType**(HYPRE\_Solver solver, HYPRE\_Int ilu\_type)

Set the type of ILU factorization.

Options for *ilu\_type* are:

- 0 : BJ with ILU(k) (default, with k = 0)
- 1 : BJ with ILUT
- 10 : GMRES with ILU(k)
- 11 : GMRES with ILUT
- 20 : NSH with ILU(k)
- 21 : NSH with ILUT
- 30 : RAS with ILU(k)
- 31 : RAS with ILUT
- 40 : (nonsymmetric permutation) DDPQ-GMRES with ILU(k)
- 41 : (nonsymmetric permutation) DDPQ-GMRES with ILUT
- 50 : GMRES with RAP-ILU(0) using MILU(0) for P

HYPRE\_Int **HYPRE\_ILUSetLocalReordering**(HYPRE\_Solver solver, HYPRE\_Int reordering\_type)

Set the type of reordering for the local matrix.

Options for *reordering\_type* are:

- 0 : No reordering
- 1 : RCM (default)

HYPRE\_Int **HYPRE\_ILUSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

(Optional) Set the print level to print setup and solve information.

- 0 : no printout (default)
- 1 : print setup information
- 2 : print solve information
- 3 : print both setup and solve information

HYPRE\_Int **HYPRE\_ILUSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

(Optional) Requests logging of solver diagnostics.

Requests additional computations for diagnostic and similar data to be logged by the user. Default is 0, do nothing. The latest residual will be available if logging > 1.

HYPRE\_Int **HYPRE\_ILUGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

(Optional) Return the number of ILU iterations.

HYPRE\_Int **HYPRE\_ILUGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*res\_norm)

(Optional) Return the norm of the final relative residual.

*HYPRE\_ParCSRMatrix* **GenerateLaplacian**(MPI\_Comm comm, HYPRE\_BigInt nx, HYPRE\_BigInt ny, HYPRE\_BigInt nz, HYPRE\_Int P, HYPRE\_Int Q, HYPRE\_Int R, HYPRE\_Int p, HYPRE\_Int q, HYPRE\_Int r, HYPRE\_Real \*value)

*HYPRE\_ParCSRMatrix* **GenerateLaplacian27pt**(MPI\_Comm comm, HYPRE\_BigInt nx, HYPRE\_BigInt ny, HYPRE\_BigInt nz, HYPRE\_Int P, HYPRE\_Int Q, HYPRE\_Int R, HYPRE\_Int p, HYPRE\_Int q, HYPRE\_Int r, HYPRE\_Real \*value)

*HYPRE\_ParCSRMatrix* **GenerateLaplacian9pt**(MPI\_Comm comm, HYPRE\_BigInt nx, HYPRE\_BigInt ny, HYPRE\_Int P, HYPRE\_Int Q, HYPRE\_Int p, HYPRE\_Int q, HYPRE\_Real \*value)

*HYPRE\_ParCSRMatrix* **GenerateDifConv**(MPI\_Comm comm, HYPRE\_BigInt nx, HYPRE\_BigInt ny, HYPRE\_BigInt nz, HYPRE\_Int P, HYPRE\_Int Q, HYPRE\_Int R, HYPRE\_Int p, HYPRE\_Int q, HYPRE\_Int r, HYPRE\_Real \*value)

*HYPRE\_ParCSRMatrix* **GenerateRotate7pt**(MPI\_Comm comm, HYPRE\_BigInt nx, HYPRE\_BigInt ny, HYPRE\_Int P, HYPRE\_Int Q, HYPRE\_Int p, HYPRE\_Int q, HYPRE\_Real alpha, HYPRE\_Real eps)

*HYPRE\_ParCSRMatrix* **GenerateVarDifConv**(MPI\_Comm comm, HYPRE\_BigInt nx, HYPRE\_BigInt ny, HYPRE\_BigInt nz, HYPRE\_Int P, HYPRE\_Int Q, HYPRE\_Int R, HYPRE\_Int p, HYPRE\_Int q, HYPRE\_Int r, HYPRE\_Real eps, *HYPRE\_ParVector* \*rhs\_ptr)

*HYPRE\_ParCSRMatrix* **GenerateRSVarDifConv**(MPI\_Comm comm, HYPRE\_BigInt nx, HYPRE\_BigInt ny, HYPRE\_BigInt nz, HYPRE\_Int P, HYPRE\_Int Q, HYPRE\_Int R, HYPRE\_Int p, HYPRE\_Int q, HYPRE\_Int r, HYPRE\_Real eps, *HYPRE\_ParVector* \*rhs\_ptr, HYPRE\_Int type)

float \***hypr\_GenerateCoordinates**(MPI\_Comm comm, HYPRE\_BigInt nx, HYPRE\_BigInt ny, HYPRE\_BigInt nz, HYPRE\_Int P, HYPRE\_Int Q, HYPRE\_Int R, HYPRE\_Int p, HYPRE\_Int q, HYPRE\_Int r, HYPRE\_Int coorddim)

### ParCSR LOBPCG Eigensolver

These routines should be used in conjunction with the generic interface in *Eigensolvers*.

HYPRE\_Int **HYPRE\_ParCSRSetupInterpreter**(mv\_InterfaceInterpreter \*i)

Load interface interpreter.

Vector part loaded with *hypr\_ParKrylov* functions and multivector part loaded with *mv\_TempMultiVector* functions.

HYPRE\_Int **HYPRE\_ParCSRSetupMatvec**(HYPRE\_MatvecFunctions \*mv)

Load Matvec interpreter with *hypr\_ParKrylov* functions.

HYPRE\_Int **HYPRE\_ParCSRMultiVectorPrint**(void \*x\_, const char \*fileName)

void \***HYPRE\_ParCSRMultiVectorRead**(MPI\_Comm comm, void \*ii\_, const char \*fileName)

HYPRE\_Int **HYPRE\_TempParCSRSetupInterpreter**(mv\_InterfaceInterpreter \*i)

### Functions

HYPRE\_Int **HYPRE\_AMECreate**(HYPRE\_Solver \*esolver)

HYPRE\_Int **HYPRE\_AMEDestroy**(HYPRE\_Solver esolver)

HYPRE\_Int **HYPRE\_AMESetup**(HYPRE\_Solver esolver)

HYPRE\_Int **HYPRE\_AMESolve**(HYPRE\_Solver esolver)

HYPRE\_Int **HYPRE\_AMESetAMSSolver**(HYPRE\_Solver esolver, HYPRE\_Solver ams\_solver)

HYPRE\_Int **HYPRE\_AMESetMassMatrix**(HYPRE\_Solver esolver, *HYPRE\_ParCSRMatrix* M)

HYPRE\_Int **HYPRE\_AMESetBlockSize**(HYPRE\_Solver esolver, HYPRE\_Int block\_size)

HYPRE\_Int **HYPRE\_AMESetMaxIter**(HYPRE\_Solver esolver, HYPRE\_Int maxit)

HYPRE\_Int **HYPRE\_AMESetMaxPCGIter**(HYPRE\_Solver esolver, HYPRE\_Int maxit)

HYPRE\_Int **HYPRE\_AMESetTol**(HYPRE\_Solver esolver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_AMESetRTol**(HYPRE\_Solver esolver, HYPRE\_Real tol)

HYPRE\_Int **HYPRE\_AMESetPrintLevel**(HYPRE\_Solver esolver, HYPRE\_Int print\_level)

HYPRE\_Int **HYPRE\_AMEGetEigenvalues**(HYPRE\_Solver esolver, HYPRE\_Real \*\*eigenvalues)

HYPRE\_Int **HYPRE\_AMEGetEigenvectors**(HYPRE\_Solver esolver, *HYPRE\_ParVector* \*\*eigenvectors)

HYPRE\_Int **HYPRE\_SchwarzCreate**(HYPRE\_Solver \*solver)

HYPRE\_Int **HYPRE\_SchwarzDestroy**(HYPRE\_Solver solver)

HYPRE\_Int **HYPRE\_SchwarzSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

Set up a Schwarz solver.

Overlapping Schwarz variants (10+) currently require host-resident ParCSR matrices and vectors.

HYPRE\_Int **HYPRE\_SchwarzSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A, *HYPRE\_ParVector* b, *HYPRE\_ParVector* x)

HYPRE\_Int **HYPRE\_SchwarzSetVariant**(HYPRE\_Solver solver, HYPRE\_Int variant)

HYPRE\_Int **HYPRE\_SchwarzSetOverlap**(HYPRE\_Solver solver, HYPRE\_Int overlap)

HYPRE\_Int **HYPRE\_SchwarzSetDomainType**(HYPRE\_Solver solver, HYPRE\_Int domain\_type)

HYPRE\_Int **HYPRE\_SchwarzSetRelaxWeight**(HYPRE\_Solver solver, HYPRE\_Real relax\_weight)

HYPRE\_Int **HYPRE\_SchwarzSetDomainStructure**(HYPRE\_Solver solver, HYPRE\_CSRMatrix domain\_structure)

HYPRE\_Int **HYPRE\_SchwarzSetNumFunctions**(HYPRE\_Solver solver, HYPRE\_Int num\_functions)

HYPRE\_Int **HYPRE\_SchwarzSetDofFunc**(HYPRE\_Solver solver, HYPRE\_Int \*dof\_func)

HYPRE\_Int **HYPRE\_SchwarzSetNonSymm**(HYPRE\_Solver solver, HYPRE\_Int use\_nonsymm)

HYPRE\_Int **HYPRE\_SchwarzSetLocalSolverType**(HYPRE\_Solver solver, HYPRE\_Int local\_solver\_type)

Set the local solver type for new overlapping Schwarz variants (10+).

If the solver is still on the default Schwarz variant 0, this call selects the corresponding RAS overlapping Schwarz variant.

#### Parameters

- **solver** – [IN] Schwarz solver
- **local\_solver\_type** – [IN]
  - 0 = ILU(k) (default)
  - 1 = ILUT
  - 2 = AMG
  - 3 = SuperLU\_dist

HYPRE\_Int **HYPRE\_SchwarzSetILUKLevelOfFill**(HYPRE\_Solver solver, HYPRE\_Int level\_of\_fill)

Set the level of fill for ILU(k) local solver.

Default is 0. (Overlapping Schwarz variants 10+ only)

HYPRE\_Int **HYPRE\_SchwarzSetILUTMaxNnzPerRow**(HYPRE\_Solver solver, HYPRE\_Int max\_nnz\_row)

Set the max nonzeros per row for ILUT local solver.

Default is 1000. (Overlapping Schwarz variants 10+ only)

HYPRE\_Int **HYPRE\_SchwarzSetILUTDroptol**(HYPRE\_Solver solver, HYPRE\_Real droptol)

Set the drop tolerance for ILUT local solver.

Default is 1e-2. (Overlapping Schwarz variants 10+ only)

HYPRE\_Int **HYPRE\_SchwarzSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

Set maximum number of iterations.

Default is 1. (Overlapping Schwarz variants 10+ only)

HYPRE\_Int **HYPRE\_SchwarzSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

Set convergence tolerance.

Default is 0.0 (no convergence checking). (Overlapping Schwarz variants 10+ only)

HYPRE\_Int **HYPRE\_SchwarzSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

Set print level.

0=none (default), 1=summary, 2=per-iteration. (Overlapping Schwarz variants 10+ only)

HYPRE\_Int **HYPRE\_SchwarzSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

Set logging level.

0=none (default), 1=store residual norms. (Overlapping Schwarz variants 10+ only)

HYPRE\_Int **HYPRE\_SchwarzGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

Get the number of iterations performed.

HYPRE\_Int **HYPRE\_SchwarzGetFinalResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

Get the final residual norm.

**HYPRE\_Int HYPRE\_ParCSRCreate**(MPI\_Comm comm, HYPRE\_Solver \*solver)  
**HYPRE\_Int HYPRE\_ParCSRDestroy**(HYPRE\_Solver solver)  
**HYPRE\_Int HYPRE\_ParCSRSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A,  
*HYPRE\_ParVector* b, *HYPRE\_ParVector* x)  
**HYPRE\_Int HYPRE\_ParCSRSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A,  
*HYPRE\_ParVector* b, *HYPRE\_ParVector* x)  
**HYPRE\_Int HYPRE\_ParCSRSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)  
**HYPRE\_Int HYPRE\_ParCSRSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)  
**HYPRE\_Int HYPRE\_ParCSRSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)  
**HYPRE\_Int HYPRE\_ParCSRSetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int stop\_crit)  
**HYPRE\_Int HYPRE\_ParCSRSetPrecond**(HYPRE\_Solver solver, *HYPRE\_PtrToParSolverFcn* precond,  
*HYPRE\_PtrToParSolverFcn* precondT,  
*HYPRE\_PtrToParSolverFcn* precond\_setup, HYPRE\_Solver  
precond\_solver)  
**HYPRE\_Int HYPRE\_ParCSRGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data)  
**HYPRE\_Int HYPRE\_ParCSRSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)  
**HYPRE\_Int HYPRE\_ParCSRGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)  
**HYPRE\_Int HYPRE\_ParCSRGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real  
\*norm)  
**HYPRE\_Int HYPRE\_BlockTridiagCreate**(HYPRE\_Solver \*solver)  
**HYPRE\_Int HYPRE\_BlockTridiagDestroy**(HYPRE\_Solver solver)  
**HYPRE\_Int HYPRE\_BlockTridiagSetup**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A,  
*HYPRE\_ParVector* b, *HYPRE\_ParVector* x)  
**HYPRE\_Int HYPRE\_BlockTridiagSolve**(HYPRE\_Solver solver, *HYPRE\_ParCSRMatrix* A,  
*HYPRE\_ParVector* b, *HYPRE\_ParVector* x)  
**HYPRE\_Int HYPRE\_BlockTridiagSetIndexSet**(HYPRE\_Solver solver, HYPRE\_Int n, HYPRE\_Int \*inds)  
**HYPRE\_Int HYPRE\_BlockTridiagSetAMGStrengthThreshold**(HYPRE\_Solver solver, HYPRE\_Real  
thresh)  
**HYPRE\_Int HYPRE\_BlockTridiagSetAMGNumSweeps**(HYPRE\_Solver solver, HYPRE\_Int num\_sweeps)  
**HYPRE\_Int HYPRE\_BlockTridiagSetAMGRelaxType**(HYPRE\_Solver solver, HYPRE\_Int relax\_type)  
**HYPRE\_Int HYPRE\_BlockTridiagSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int print\_level)

## 8.8 Krylov Solvers

### group Krylov Solvers

A basic interface for Krylov solvers.

These solvers support many of the matrix/vector storage schemes in hypr. They should be used in conjunction with the storage-specific interfaces, particularly the specific Create() and Destroy() functions.

### Krylov Solvers

```
typedef HYPRE_Int (*HYPRE_PtrToModifyPCFcn)(HYPRE_Solver, HYPRE_Int, HYPRE_Real)
```

### PCG Solver

```
HYPRE_Int HYPRE_PCGSetup(HYPRE_Solver solver, HYPRE_Matrix A, HYPRE_Vector b, HYPRE_Vector x)
```

Prepare to solve the system.

The coefficient data in  $b$  and  $x$  is ignored here, but information about the layout of the data may be used.

```
HYPRE_Int HYPRE_PCGSolve(HYPRE_Solver solver, HYPRE_Matrix A, HYPRE_Vector b, HYPRE_Vector x)
```

Solve the system.

```
HYPRE_Int HYPRE_PCGSetTol(HYPRE_Solver solver, HYPRE_Real tol)
```

(Optional) Set the relative convergence tolerance.

```
HYPRE_Int HYPRE_PCGSetAbsoluteTol(HYPRE_Solver solver, HYPRE_Real a_tol)
```

(Optional) Set the absolute convergence tolerance (default is 0).

If one desires the convergence test to check the absolute convergence tolerance *only*, then set the relative convergence tolerance to 0.0. (The default convergence test is  $\langle C * r, r \rangle \leq \max(\text{relative\_tolerance}^2 * \langle C * b, b \rangle, \text{absolute\_tolerance}^2)$ .)

```
HYPRE_Int HYPRE_PCGSetResidualTol(HYPRE_Solver solver, HYPRE_Real r_tol)
```

(Optional) Set a residual-based convergence tolerance which checks if  $\|r_{old} - r_{new}\| < r_{tol} \|b\|$ .

This is useful when trying to converge to very low relative and/or absolute tolerances, in order to bail-out before roundoff errors affect the approximation.

```
HYPRE_Int HYPRE_PCGSetAbsoluteTolFactor(HYPRE_Solver solver, HYPRE_Real abstolf)
```

```
HYPRE_Int HYPRE_PCGSetConvergenceFactorTol(HYPRE_Solver solver, HYPRE_Real cf_tol)
```

```
HYPRE_Int HYPRE_PCGSetStopCrit(HYPRE_Solver solver, HYPRE_Int stop_crit)
```

```
HYPRE_Int HYPRE_PCGSetMaxIter(HYPRE_Solver solver, HYPRE_Int max_iter)
```

(Optional) Set maximum number of iterations.

```
HYPRE_Int HYPRE_PCGSetTwoNorm(HYPRE_Solver solver, HYPRE_Int two_norm)
```

(Optional) Use the two-norm in stopping criteria.

```
HYPRE_Int HYPRE_PCGSetRelChange(HYPRE_Solver solver, HYPRE_Int rel_change)
```

(Optional) Additionally require that the relative difference in successive iterates be small.

```
HYPRE_Int HYPRE_PCGSetRecomputeResidual(HYPRE_Solver solver, HYPRE_Int recompute_residual)
```

(Optional) Recompute the residual at the end to double-check convergence.

HYPRE\_Int **HYPRE\_PCGSetRecomputeResidualP**(HYPRE\_Solver solver, HYPRE\_Int  
recompute\_residual\_p)

(Optional) Periodically recompute the residual while iterating.

HYPRE\_Int **HYPRE\_PCGSetFlex**(HYPRE\_Solver solver, HYPRE\_Int flex)

(Optional) Setting this to 1 allows use of Polak-Ribiere Method (flexible) this increases robustness, but adds an additional dot product per iteration

HYPRE\_Int **HYPRE\_PCGSetSkipBreak**(HYPRE\_Solver solver, HYPRE\_Int skip\_break)

(Optional) Skips subnormal alpha, gamma and iprod values in CG.

If set to 0 (default): will break if values are below HYPRE\_REAL\_MIN If set to 1: will break if values are below HYPRE\_REAL\_TRUE\_MIN (requires C11 minimal or will check to HYPRE\_REAL\_MIN) If set to 2: will break if values are <= 0. If set to 3 or larger: will not break at all

HYPRE\_Int **HYPRE\_PCGSetPrecond**(HYPRE\_Solver solver, HYPRE\_PtrToSolverFcn precondition,  
HYPRE\_PtrToSolverFcn precondition\_setup, HYPRE\_Solver  
precond\_solver)

(Optional) Set the preconditioner to use.

HYPRE\_Int **HYPRE\_PCGSetPrecondMatrix**(HYPRE\_Solver solver, HYPRE\_Matrix precondition\_matrix)

HYPRE\_Int **HYPRE\_PCGSetPreconditioner**(HYPRE\_Solver solver, HYPRE\_Solver precondition)

(Optional) Set the preconditioner to use in a generic fashion.

This function does not require explicit input of the setup and solve pointers of the preconditioner object. Instead, it automatically extracts this information from the aforementioned object.

HYPRE\_Int **HYPRE\_PCGSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

(Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_PCGSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int level)

(Optional) Set the amount of printing to do to the screen.

HYPRE\_Int **HYPRE\_PCGGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_PCGGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

HYPRE\_Int **HYPRE\_PCGGetResidual**(HYPRE\_Solver solver, void \*residual)

Return the residual.

HYPRE\_Int **HYPRE\_PCGGetTol**(HYPRE\_Solver solver, HYPRE\_Real \*tol)

HYPRE\_Int **HYPRE\_PCGGetResidualTol**(HYPRE\_Solver solver, HYPRE\_Real \*rtol)

HYPRE\_Int **HYPRE\_PCGGetAbsoluteTolFactor**(HYPRE\_Solver solver, HYPRE\_Real \*abstolf)

HYPRE\_Int **HYPRE\_PCGGetConvergenceFactorTol**(HYPRE\_Solver solver, HYPRE\_Real \*cf\_tol)

HYPRE\_Int **HYPRE\_PCGGetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int \*stop\_crit)

HYPRE\_Int **HYPRE\_PCGGetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int \*max\_iter)

HYPRE\_Int **HYPRE\_PCGGetTwoNorm**(HYPRE\_Solver solver, HYPRE\_Int \*two\_norm)

HYPRE\_Int **HYPRE\_PCGGetRelChange**(HYPRE\_Solver solver, HYPRE\_Int \*rel\_change)

HYPRE\_Int **HYPRE\_PCGGetSkipBreak**(HYPRE\_Solver solver, HYPRE\_Int \*skip\_break)

HYPRE\_Int **HYPRE\_PCGGetFlex**(HYPRE\_Solver solver, HYPRE\_Int \*flex)

HYPRE\_Int **HYPRE\_PCGGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data\_ptr)

HYPRE\_Int **HYPRE\_PCGGetPrecondMatrix**(HYPRE\_Solver solver, HYPRE\_Matrix \*precond\_matrix\_ptr)

HYPRE\_Int **HYPRE\_PCGGetLogging**(HYPRE\_Solver solver, HYPRE\_Int \*logging)

HYPRE\_Int **HYPRE\_PCGGetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int \*level)

HYPRE\_Int **HYPRE\_PCGGetConverged**(HYPRE\_Solver solver, HYPRE\_Int \*converged)

### GMRES Solver

HYPRE\_Int **HYPRE\_GMRESSetup**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b, HYPRE\_Vector x)

Prepare to solve the system.

The coefficient data in  $b$  and  $x$  is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_GMRESSolve**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b, HYPRE\_Vector x)

Solve the system.

HYPRE\_Int **HYPRE\_GMRESSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

(Optional) Set the relative convergence tolerance.

HYPRE\_Int **HYPRE\_GMRESSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real a\_tol)

(Optional) Set the absolute convergence tolerance (default is 0).

If one desires the convergence test to check the absolute convergence tolerance *only*, then set the relative convergence tolerance to 0.0. (The convergence test is  $\|r\| \leq \max(\text{relative\_tolerance} * \|b\|, \text{absolute\_tolerance})$ .)

HYPRE\_Int **HYPRE\_GMRESSetConvergenceFactorTol**(HYPRE\_Solver solver, HYPRE\_Real cf\_tol)

HYPRE\_Int **HYPRE\_GMRESSetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int stop\_crit)

HYPRE\_Int **HYPRE\_GMRESSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_GMRESSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_GMRESSetKDim**(HYPRE\_Solver solver, HYPRE\_Int k\_dim)

(Optional) Set the maximum size of the Krylov space.

HYPRE\_Int **HYPRE\_GMRESSetRelChange**(HYPRE\_Solver solver, HYPRE\_Int rel\_change)

(Optional) Additionally require that the relative difference in successive iterates be small.

HYPRE\_Int **HYPRE\_GMRESSetSkipRealResidualCheck**(HYPRE\_Solver solver, HYPRE\_Int skip\_real\_r\_check)

(Optional) By default, hypr checks for convergence by evaluating the actual residual before returnig from GMRES (with restart if the true residual does not indicate convergence).

This option allows users to skip the evaluation and the check of the actual residual for badly conditioned problems where restart is not expected to be beneficial.

HYPRE\_Int **HYPRE\_GMRESSetPrecond**(HYPRE\_Solver solver, HYPRE\_PtrToSolverFcn precondition, HYPRE\_PtrToSolverFcn precondition\_setup, HYPRE\_Solver precondition\_solver)

(Optional) Set the preconditioner to use.

HYPRE\_Int **HYPRE\_GMRESSetPrecondMatrix**(HYPRE\_Solver solver, HYPRE\_Matrix precondition\_matrix)

HYPRE\_Int **HYPRE\_GMRESSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

(Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_GMRESSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int level)

(Optional) Set the amount of printing to do to the screen.

Options 6-9 are mainly for debugging purposes as they require setting up a reference solution vector via *HYPRE\_GMRESSetRefSolution*

#### Parameters

- **solver** – The solver object
- **level** – The print level:
  - 0: no output
  - 1: print warnings
  - 2: print convergence history for the absolute and relative residual norms
  - 3: print absolute residual norms for each tag in multi-tag vectors
  - 4: print relative residual norms for each tag in multi-tag vectors, where each residual norm is divided by the norm of its corresponding tagged component of the right-hand side vector (RHS).
  - 5: print relative residual norms for each tag in multi-tag vectors, where the residual norm is divided by the norm of the original right-hand side vector (RHS).
  - 6: print convergence history for the absolute and relative error norms
  - 7: print absolute error norms for each tag in multi-tag vectors.
  - 8: print relative error norms for each tag in multi-tag vectors, where each residual norm is divided by the norm of its corresponding tagged component of the initial error vector.
  - 9: print relative error norms for each tag in multi-tag vectors, where the error norms are divided by the norm of the initial error vector.

HYPRE\_Int **HYPRE\_GMRESGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_GMRESGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

HYPRE\_Int **HYPRE\_GMRESGetResidual**(HYPRE\_Solver solver, void \*residual)

Return the residual.

HYPRE\_Int **HYPRE\_GMRESGetSkipRealResidualCheck**(HYPRE\_Solver solver, HYPRE\_Int \*skip\_real\_r\_check)

**HYPRE\_Int HYPRE\_GMRESGetTol**(HYPRE\_Solver solver, HYPRE\_Real \*tol)  
**HYPRE\_Int HYPRE\_GMRESGetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real \*tol)  
**HYPRE\_Int HYPRE\_GMRESGetConvergenceFactorTol**(HYPRE\_Solver solver, HYPRE\_Real \*cf\_tol)  
**HYPRE\_Int HYPRE\_GMRESGetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int \*stop\_crit)  
**HYPRE\_Int HYPRE\_GMRESGetMinIter**(HYPRE\_Solver solver, HYPRE\_Int \*min\_iter)  
**HYPRE\_Int HYPRE\_GMRESGetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int \*max\_iter)  
**HYPRE\_Int HYPRE\_GMRESGetKDim**(HYPRE\_Solver solver, HYPRE\_Int \*k\_dim)  
**HYPRE\_Int HYPRE\_GMRESGetRelChange**(HYPRE\_Solver solver, HYPRE\_Int \*rel\_change)  
**HYPRE\_Int HYPRE\_GMRESGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data\_ptr)  
**HYPRE\_Int HYPRE\_GMRESGetPrecondMatrix**(HYPRE\_Solver solver, HYPRE\_Matrix  
\*precond\_matrix\_ptr)  
**HYPRE\_Int HYPRE\_GMRESGetLogging**(HYPRE\_Solver solver, HYPRE\_Int \*level)  
**HYPRE\_Int HYPRE\_GMRESGetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int \*level)  
**HYPRE\_Int HYPRE\_GMRESGetConverged**(HYPRE\_Solver solver, HYPRE\_Int \*converged)  
**HYPRE\_Int HYPRE\_GMRESSetRefSolution**(HYPRE\_Solver solver, HYPRE\_Vector xref)  
(Optional) Set a reference solution vector for error computation.  
**HYPRE\_Int HYPRE\_GMRESGetRefSolution**(HYPRE\_Solver solver, HYPRE\_Vector \*xref)  
Get the reference solution vector.

### FlexGMRES Solver

**HYPRE\_Int HYPRE\_FlexGMRESSetup**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b,  
HYPRE\_Vector x)  
Prepare to solve the system.  
The coefficient data in  $b$  and  $x$  is ignored here, but information about the layout of the data may be used.

**HYPRE\_Int HYPRE\_FlexGMRESSolve**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b,  
HYPRE\_Vector x)  
Solve the system.

**HYPRE\_Int HYPRE\_FlexGMRESSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)  
(Optional) Set the convergence tolerance.

**HYPRE\_Int HYPRE\_FlexGMRESSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real a\_tol)  
(Optional) Set the absolute convergence tolerance (default is 0).  
If one desires the convergence test to check the absolute convergence tolerance *only*, then set the relative convergence tolerance to 0.0. (The convergence test is  $\|r\| \leq \max(\text{relative\_tolerance} * \|b\|, \text{absolute\_tolerance})$ .)

**HYPRE\_Int HYPRE\_FlexGMRESSetConvergenceFactorTol**(HYPRE\_Solver solver, HYPRE\_Real cf\_tol)  
**HYPRE\_Int HYPRE\_FlexGMRESSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_FlexGMRESsetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)  
 (Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_FlexGMRESsetKDim**(HYPRE\_Solver solver, HYPRE\_Int k\_dim)  
 (Optional) Set the maximum size of the Krylov space.

HYPRE\_Int **HYPRE\_FlexGMRESsetPrecond**(HYPRE\_Solver solver, HYPRE\_PtrToSolverFcn precondition, HYPRE\_PtrToSolverFcn precondition\_setup, HYPRE\_Solver precondition\_solver)  
 (Optional) Set the preconditioner to use.

HYPRE\_Int **HYPRE\_FlexGMRESsetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)  
 (Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_FlexGMRESsetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int level)  
 (Optional) Set the amount of printing to do to the screen.

HYPRE\_Int **HYPRE\_FlexGMRESgetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)  
 Return the number of iterations taken.

HYPRE\_Int **HYPRE\_FlexGMRESgetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)  
 Return the norm of the final relative residual.

HYPRE\_Int **HYPRE\_FlexGMRESgetResidual**(HYPRE\_Solver solver, void \*residual)  
 Return the residual.

HYPRE\_Int **HYPRE\_FlexGMRESgetTol**(HYPRE\_Solver solver, HYPRE\_Real \*tol)

HYPRE\_Int **HYPRE\_FlexGMRESgetConvergenceFactorTol**(HYPRE\_Solver solver, HYPRE\_Real \*cf\_tol)

HYPRE\_Int **HYPRE\_FlexGMRESgetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int \*stop\_crit)

HYPRE\_Int **HYPRE\_FlexGMRESgetMinIter**(HYPRE\_Solver solver, HYPRE\_Int \*min\_iter)

HYPRE\_Int **HYPRE\_FlexGMRESgetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int \*max\_iter)

HYPRE\_Int **HYPRE\_FlexGMRESgetKDim**(HYPRE\_Solver solver, HYPRE\_Int \*k\_dim)

HYPRE\_Int **HYPRE\_FlexGMRESgetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data\_ptr)

HYPRE\_Int **HYPRE\_FlexGMRESgetLogging**(HYPRE\_Solver solver, HYPRE\_Int \*level)

HYPRE\_Int **HYPRE\_FlexGMRESgetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int \*level)

HYPRE\_Int **HYPRE\_FlexGMRESgetConverged**(HYPRE\_Solver solver, HYPRE\_Int \*converged)

HYPRE\_Int **HYPRE\_FlexGMRESsetModifyPC**(HYPRE\_Solver solver, *HYPRE\_PtrToModifyPCFcn* modify\_pc)  
 (Optional) Set a user-defined function to modify solve-time preconditioner attributes.

### LGMRES Solver

HYPRE\_Int **HYPRE\_LGMRESsetup**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b, HYPRE\_Vector x)  
 Prepare to solve the system.  
 The coefficient data in  $b$  and  $x$  is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_LGMRESSolve**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b, HYPRE\_Vector x)

Solve the system.

Details on LGMRES may be found in A. H. Baker, E.R. Jessup, and T.A. Manteuffel, "A technique for accelerating the

convergence of restarted GMRES." SIAM Journal on Matrix Analysis and Applications, 26 (2005), pp. 962-984. LGMRES(m,k) in the paper corresponds to LGMRES(Kdim+AugDim, AugDim).

HYPRE\_Int **HYPRE\_LGMRESSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance.

HYPRE\_Int **HYPRE\_LGMRESSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real a\_tol)

(Optional) Set the absolute convergence tolerance (default is 0).

If one desires the convergence test to check the absolute convergence tolerance *only*, then set the relative convergence tolerance to 0.0. (The convergence test is  $\|r\| \leq \max(\text{relative\_tolerance} * \|b\|, \text{absolute\_tolerance})$ .)

HYPRE\_Int **HYPRE\_LGMRESSetConvergenceFactorTol**(HYPRE\_Solver solver, HYPRE\_Real cf\_tol)

HYPRE\_Int **HYPRE\_LGMRESSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_LGMRESSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_LGMRESSetKDim**(HYPRE\_Solver solver, HYPRE\_Int k\_dim)

(Optional) Set the maximum size of the approximation space (includes the augmentation vectors).

HYPRE\_Int **HYPRE\_LGMRESSetAugDim**(HYPRE\_Solver solver, HYPRE\_Int aug\_dim)

(Optional) Set the number of augmentation vectors (default: 2).

HYPRE\_Int **HYPRE\_LGMRESSetPrecond**(HYPRE\_Solver solver, HYPRE\_PtrToSolverFcn precondition, HYPRE\_PtrToSolverFcn precondition\_setup, HYPRE\_Solver precondition\_solver)

(Optional) Set the preconditioner to use.

HYPRE\_Int **HYPRE\_LGMRESSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

(Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_LGMRESSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int level)

(Optional) Set the amount of printing to do to the screen.

HYPRE\_Int **HYPRE\_LGMRESGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_LGMRESGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

HYPRE\_Int **HYPRE\_LGMRESGetResidual**(HYPRE\_Solver solver, void \*residual)

Return the residual.

HYPRE\_Int **HYPRE\_LGMRESGetTol**(HYPRE\_Solver solver, HYPRE\_Real \*tol)

HYPRE\_Int **HYPRE\_LGMRESGetConvergenceFactorTol**(HYPRE\_Solver solver, HYPRE\_Real \*cf\_tol)

HYPRE\_Int **HYPRE\_LGMRESGetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int \*stop\_crit)  
HYPRE\_Int **HYPRE\_LGMRESGetMinIter**(HYPRE\_Solver solver, HYPRE\_Int \*min\_iter)  
HYPRE\_Int **HYPRE\_LGMRESGetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int \*max\_iter)  
HYPRE\_Int **HYPRE\_LGMRESGetKDim**(HYPRE\_Solver solver, HYPRE\_Int \*k\_dim)  
HYPRE\_Int **HYPRE\_LGMRESGetAugDim**(HYPRE\_Solver solver, HYPRE\_Int \*k\_dim)  
HYPRE\_Int **HYPRE\_LGMRESGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data\_ptr)  
HYPRE\_Int **HYPRE\_LGMRESGetLogging**(HYPRE\_Solver solver, HYPRE\_Int \*level)  
HYPRE\_Int **HYPRE\_LGMRESGetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int \*level)  
HYPRE\_Int **HYPRE\_LGMRESGetConverged**(HYPRE\_Solver solver, HYPRE\_Int \*converged)

### COGMRES Solver

HYPRE\_Int **HYPRE\_COGMRESSetup**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b,  
HYPRE\_Vector x)

Prepare to solve the system.

The coefficient data in  $b$  and  $x$  is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_COGMRESolve**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b,  
HYPRE\_Vector x)

Solve the system.

HYPRE\_Int **HYPRE\_COGMRESSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance.

HYPRE\_Int **HYPRE\_COGMRESSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real a\_tol)

(Optional) Set the absolute convergence tolerance (default is 0).

If one desires the convergence test to check the absolute convergence tolerance *only*, then set the relative convergence tolerance to 0.0. (The convergence test is  $\|r\| \leq \max(\text{relative\_tolerance} * \|b\|, \text{absolute\_tolerance})$ .)

HYPRE\_Int **HYPRE\_COGMRESSetConvergenceFactorTol**(HYPRE\_Solver solver, HYPRE\_Real cf\_tol)

HYPRE\_Int **HYPRE\_COGMRESSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_COGMRESSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_COGMRESSetKDim**(HYPRE\_Solver solver, HYPRE\_Int k\_dim)

(Optional) Set the maximum size of the Krylov space.

HYPRE\_Int **HYPRE\_COGMRESSetUnroll**(HYPRE\_Solver solver, HYPRE\_Int unroll)

(Optional) Set number of unrolling in mass functions in COGMRES Can be 4 or 8.

Default: no unrolling.

HYPRE\_Int **HYPRE\_COGMRESSetCGS**(HYPRE\_Solver solver, HYPRE\_Int cgs)

(Optional) Set the number of orthogonalizations in COGMRES (at most 2).

HYPRE\_Int **HYPRE\_COGMRESSetPrecond**(HYPRE\_Solver solver, HYPRE\_PtrToSolverFcn precondition, HYPRE\_PtrToSolverFcn precondition\_setup, HYPRE\_Solver precondition\_solver)

(Optional) Set the preconditioner to use.

HYPRE\_Int **HYPRE\_COGMRESSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

(Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_COGMRESSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int level)

(Optional) Set the amount of printing to do to the screen.

HYPRE\_Int **HYPRE\_COGMRESGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_COGMRESGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

HYPRE\_Int **HYPRE\_COGMRESGetResidual**(HYPRE\_Solver solver, void \*residual)

Return the residual.

HYPRE\_Int **HYPRE\_COGMRESGetTol**(HYPRE\_Solver solver, HYPRE\_Real \*tol)

HYPRE\_Int **HYPRE\_COGMRESGetConvergenceFactorTol**(HYPRE\_Solver solver, HYPRE\_Real \*cf\_tol)

HYPRE\_Int **HYPRE\_COGMRESGetMinIter**(HYPRE\_Solver solver, HYPRE\_Int \*min\_iter)

HYPRE\_Int **HYPRE\_COGMRESGetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int \*max\_iter)

HYPRE\_Int **HYPRE\_COGMRESGetKDim**(HYPRE\_Solver solver, HYPRE\_Int \*k\_dim)

HYPRE\_Int **HYPRE\_COGMRESGetUnroll**(HYPRE\_Solver solver, HYPRE\_Int \*unroll)

HYPRE\_Int **HYPRE\_COGMRESGetCGS**(HYPRE\_Solver solver, HYPRE\_Int \*cgs)

HYPRE\_Int **HYPRE\_COGMRESGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data\_ptr)

HYPRE\_Int **HYPRE\_COGMRESGetLogging**(HYPRE\_Solver solver, HYPRE\_Int \*level)

HYPRE\_Int **HYPRE\_COGMRESGetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int \*level)

HYPRE\_Int **HYPRE\_COGMRESGetConverged**(HYPRE\_Solver solver, HYPRE\_Int \*converged)

HYPRE\_Int **HYPRE\_COGMRESSetModifyPC**(HYPRE\_Solver solver, *HYPRE\_PtrToModifyPCFcn* modify\_pc)

(Optional) Set a user-defined function to modify solve-time preconditioner attributes.

### BiCGSTAB Solver

HYPRE\_Int **HYPRE\_BiCGSTABDestroy**(HYPRE\_Solver solver)

HYPRE\_Int **HYPRE\_BiCGSTABSetup**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b, HYPRE\_Vector x)

Prepare to solve the system.

The coefficient data in  $b$  and  $x$  is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_BiCGSTABolve**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b, HYPRE\_Vector x)

Solve the system.

HYPRE\_Int **HYPRE\_BiCGSTABSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance.

HYPRE\_Int **HYPRE\_BiCGSTABSetAbsoluteTol**(HYPRE\_Solver solver, HYPRE\_Real a\_tol)

(Optional) Set the absolute convergence tolerance (default is 0).

If one desires the convergence test to check the absolute convergence tolerance *only*, then set the relative convergence tolerance to 0.0. (The convergence test is  $\|r\| \leq \max(\text{relative\_tolerance} * \|b\|, \text{absolute\_tolerance})$ .)

HYPRE\_Int **HYPRE\_BiCGSTABSetConvergenceFactorTol**(HYPRE\_Solver solver, HYPRE\_Real cf\_tol)

HYPRE\_Int **HYPRE\_BiCGSTABSetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int stop\_crit)

HYPRE\_Int **HYPRE\_BiCGSTABSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_BiCGSTABSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_BiCGSTABSetPrecond**(HYPRE\_Solver solver, HYPRE\_PtrToSolverFcn precondition, HYPRE\_PtrToSolverFcn precondition\_setup, HYPRE\_Solver precondition\_solver)

(Optional) Set the preconditioner to use.

HYPRE\_Int **HYPRE\_BiCGSTABSetPrecondMatrix**(HYPRE\_Solver solver, HYPRE\_Matrix precondition\_matrix)

HYPRE\_Int **HYPRE\_BiCGSTABSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

(Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_BiCGSTABSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int level)

(Optional) Set the amount of printing to do to the screen.

HYPRE\_Int **HYPRE\_BiCGSTABGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_BiCGSTABGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

HYPRE\_Int **HYPRE\_BiCGSTABGetResidual**(HYPRE\_Solver solver, void \*residual)

Return the residual.

HYPRE\_Int **HYPRE\_BiCGSTABGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data\_ptr)

HYPRE\_Int **HYPRE\_BiCGSTABGetPrecondMatrix**(HYPRE\_Solver solver, HYPRE\_Matrix \*precond\_matrix\_ptr)

## CGNR Solver

HYPRE\_Int **HYPRE\_CGNRDestroy**(HYPRE\_Solver solver)

HYPRE\_Int **HYPRE\_CGNRSetup**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b, HYPRE\_Vector x)

Prepare to solve the system.

The coefficient data in  $b$  and  $x$  is ignored here, but information about the layout of the data may be used.

HYPRE\_Int **HYPRE\_CGNSolve**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b,  
HYPRE\_Vector x)

Solve the system.

HYPRE\_Int **HYPRE\_CGNSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

(Optional) Set the convergence tolerance.

HYPRE\_Int **HYPRE\_CGNSetStopCrit**(HYPRE\_Solver solver, HYPRE\_Int stop\_crit)

HYPRE\_Int **HYPRE\_CGNSetMinIter**(HYPRE\_Solver solver, HYPRE\_Int min\_iter)

HYPRE\_Int **HYPRE\_CGNSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_CGNSetPrecond**(HYPRE\_Solver solver, HYPRE\_PtrToSolverFcn precondition,  
HYPRE\_PtrToSolverFcn preconditionT, HYPRE\_PtrToSolverFcn  
precondition\_setup, HYPRE\_Solver precondition\_solver)

(Optional) Set the preconditioner to use.

Note that the only preconditioner available in hypr for use with CGNR is currently BoomerAMG. It requires to use Jacobi as a smoother without CF smoothing, i.e. `relax_type` needs to be set to 0 or 7 and `relax_order` needs to be set to 0 by the user, since these are not default values. It can be used with a relaxation weight for Jacobi, which can significantly improve convergence.

HYPRE\_Int **HYPRE\_CGNSetLogging**(HYPRE\_Solver solver, HYPRE\_Int logging)

(Optional) Set the amount of logging to do.

HYPRE\_Int **HYPRE\_CGNGetNumIterations**(HYPRE\_Solver solver, HYPRE\_Int \*num\_iterations)

Return the number of iterations taken.

HYPRE\_Int **HYPRE\_CGNGetFinalRelativeResidualNorm**(HYPRE\_Solver solver, HYPRE\_Real \*norm)

Return the norm of the final relative residual.

HYPRE\_Int **HYPRE\_CGNGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data\_ptr)

## 8.9 Eigensolvers

### *group* Eigensolvers

A basic interface for eigensolvers.

These eigensolvers support many of the matrix/vector storage schemes in hypr. They should be used in conjunction with the storage-specific interfaces.

### LOBPCG Eigensolver

HYPRE\_Int **HYPRE\_LOBPCGCreate**(mv\_InterfaceInterpreter \*interpreter, HYPRE\_MatvecFunctions  
\*mvfunctions, HYPRE\_Solver \*solver)

LOBPCG constructor.

HYPRE\_Int **HYPRE\_LOBPCGDestroy**(HYPRE\_Solver solver)

LOBPCG destructor.

HYPRE\_Int **HYPRE\_LOBPCGSetPrecond**(HYPRE\_Solver solver, HYPRE\_PtrToSolverFcn precondition,  
HYPRE\_PtrToSolverFcn precondition\_setup, HYPRE\_Solver  
precondition\_solver)

(Optional) Set the preconditioner to use.

If not called, preconditioning is not used.

HYPRE\_Int **HYPRE\_LOBPCGGetPrecond**(HYPRE\_Solver solver, HYPRE\_Solver \*precond\_data\_ptr)

HYPRE\_Int **HYPRE\_LOBPCGSetup**(HYPRE\_Solver solver, HYPRE\_Matrix A, HYPRE\_Vector b, HYPRE\_Vector x)

Set up  $A$  and the preconditioner (if there is one).

HYPRE\_Int **HYPRE\_LOBPCGSetupB**(HYPRE\_Solver solver, HYPRE\_Matrix B, HYPRE\_Vector x)

(Optional) Set up  $B$ .

If not called,  $B = I$ .

HYPRE\_Int **HYPRE\_LOBPCGSetupT**(HYPRE\_Solver solver, HYPRE\_Matrix T, HYPRE\_Vector x)

(Optional) Set the preconditioning to be applied to  $Tx = b$ , not  $Ax = b$ .

HYPRE\_Int **HYPRE\_LOBPCGSolve**(HYPRE\_Solver solver, mv\_MultiVectorPtr y, mv\_MultiVectorPtr x, HYPRE\_Real \*lambda)

Solve  $Ax = \lambda Bx$ ,  $y'x = 0$ .

HYPRE\_Int **HYPRE\_LOBPCGSetTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

(Optional) Set the absolute convergence tolerance.

HYPRE\_Int **HYPRE\_LOBPCGSetRTol**(HYPRE\_Solver solver, HYPRE\_Real tol)

(Optional) Set the relative convergence tolerance.

HYPRE\_Int **HYPRE\_LOBPCGSetMaxIter**(HYPRE\_Solver solver, HYPRE\_Int max\_iter)

(Optional) Set maximum number of iterations.

HYPRE\_Int **HYPRE\_LOBPCGSetPrecondUsageMode**(HYPRE\_Solver solver, HYPRE\_Int mode)

Define which initial guess for inner PCG iterations to use: *mode* = 0: use zero initial guess, otherwise use RHS.

HYPRE\_Int **HYPRE\_LOBPCGSetPrintLevel**(HYPRE\_Solver solver, HYPRE\_Int level)

(Optional) Set the amount of printing to do to the screen.

utilities\_FortranMatrix \***HYPRE\_LOBPCGResidualNorms**(HYPRE\_Solver solver)

utilities\_FortranMatrix \***HYPRE\_LOBPCGResidualNormsHistory**(HYPRE\_Solver solver)

utilities\_FortranMatrix \***HYPRE\_LOBPCGEigenvaluesHistory**(HYPRE\_Solver solver)

HYPRE\_Int **HYPRE\_LOBPCGIterations**(HYPRE\_Solver solver)

## 8.10 Utilities

### *group* Utilities

Various utilities available in hypr.

### Multiprecision

enum **HYPRE\_Precision**

Available precisions.

*Values:*

enumerator **HYPRE\_REAL\_SINGLE**

enumerator **HYPRE\_REAL\_DOUBLE**

enumerator **HYPRE\_REAL\_LONGDOUBLE**

**HYPRE\_Int HYPRE\_SetGlobalPrecision**(*HYPRE\_Precision* precision)

Set the global default runtime precision.

**HYPRE\_Int HYPRE\_GetGlobalPrecision**(*HYPRE\_Precision* \*precision)

Get the global default runtime precision.

## Memory Management

enum **\_HYPRE\_MemoryLocation**

Memory location.

*Values:*

enumerator **HYPRE\_MEMORY\_UNDEFINED**

enumerator **HYPRE\_MEMORY\_HOST**

enumerator **HYPRE\_MEMORY\_DEVICE**

typedef enum *\_HYPRE\_MemoryLocation* **HYPRE\_MemoryLocation**

Memory location.

**HYPRE\_Int HYPRE\_SetMemoryLocation**(*HYPRE\_MemoryLocation* memory\_location)

(Optional) Sets the default (abstract) memory location.

**HYPRE\_Int HYPRE\_GetMemoryLocation**(*HYPRE\_MemoryLocation* \*memory\_location)

(Optional) Gets a pointer to the default (abstract) memory location.

## Execution Policy

enum **\_HYPRE\_ExecutionPolicy**

Execution Policy.

*Values:*

enumerator **HYPRE\_EXEC\_UNDEFINED**

enumerator **HYPRE\_EXEC\_HOST**

enumerator **HYPRE\_EXEC\_DEVICE**

typedef enum *HYPRE\_ExecutionPolicy* **HYPRE\_ExecutionPolicy**

Execution Policy.

**HYPRE\_Int HYPRE\_SetExecutionPolicy**(*HYPRE\_ExecutionPolicy* exec\_policy)

(Optional) Sets the default execution policy.

**HYPRE\_Int HYPRE\_GetExecutionPolicy**(*HYPRE\_ExecutionPolicy* \*exec\_policy)

(Optional) Gets a pointer to the default execution policy.

const char \***HYPRE\_GetExecutionPolicyName**(*HYPRE\_ExecutionPolicy* exec\_policy)

(Optional) Returns a string denoting the execution policy passed as input.

## Error Codes

**HYPRE\_Int HYPRE\_GetGlobalError**(MPI\_Comm comm)

Return an aggregate error code representing the collective status of all ranks.

**HYPRE\_Int HYPRE\_GetError**(void)

Return the current hypr error flag.

**HYPRE\_Int HYPRE\_CheckError**(HYPRE\_Int hypr\_ierr, HYPRE\_Int hypr\_error\_code)

Check if the given error flag contains the given error code.

**HYPRE\_Int HYPRE\_GetErrorArg**(void)

Return the index of the argument (counting from 1) where argument error (HYPRE\_ERROR\_ARG) has occurred.

void **HYPRE\_DescribeError**(HYPRE\_Int hypr\_ierr, char \*descr)

Describe the given error flag in the given string.

**HYPRE\_Int HYPRE\_ClearAllErrors**(void)

Clear the hypr error flag.

**HYPRE\_Int HYPRE\_ClearError**(HYPRE\_Int hypr\_error\_code)

Clear the given error code from the hypr error flag.

**HYPRE\_Int HYPRE\_SetPrintErrorMode**(HYPRE\_Int mode)

Set behavior for printing errors: mode 0 = stderr, mode 1 = memory buffer.

**HYPRE\_Int HYPRE\_SetPrintErrorVerbosity**(HYPRE\_Int code, HYPRE\_Int verbosity)

Set which error code messages to record for printing: code is an error code such as HYPRE\_ERROR\_CONV, code -1 = all codes, verbosity 0 = do not record.

**HYPRE\_Int HYPRE\_GetErrorMessages**(char \*\*buffer, HYPRE\_Int \*bufsz)

Return a buffer of error messages and clear them in hypr.

**HYPRE\_Int HYPRE\_PrintErrorMessages**(MPI\_Comm comm)

Print the error messages and clear them in hypr.

**HYPRE\_Int HYPRE\_ClearErrorMessages**(void)

Clear the error messages in hypr and free any related memory allocated.

**HYPRE\_ERROR\_GENERIC**

**HYPRE\_ERROR\_MEMORY**

**HYPRE\_ERROR\_ARG**

**HYPRE\_ERROR\_CONV**

**HYPRE\_MAX\_FILE\_NAME\_LEN**

**HYPRE\_MAX\_MSG\_LEN**

### Initialize and Finalize

**HYPRE\_Int HYPRE\_Initialize**(void)

(Required) Initializes the hypr library.

**HYPRE\_Int HYPRE\_DeviceInitialize**(void)

(Optional) Initializes GPU features in the hypr library.

**HYPRE\_Int HYPRE\_Finalize**(void)

(Required) Finalizes the hypr library.

**HYPRE\_Int HYPRE\_Initialized**(void)

(Optional) Returns true if the hypr library has been initialized but not finalized yet.

**HYPRE\_Int HYPRE\_Finalized**(void)

(Optional) Returns true if the hypr library has been finalized but not re-initialized yet.

**HYPRE\_Init**()

(Required) Initializes the hypr library.

This function is provided for backward compatibility. Please, use `HYPRE_Initialize` instead.

### Miscellaneous Information

**HYPRE\_Int HYPRE\_PrintDeviceInfo**(void)

Print GPU information.

**HYPRE\_Int HYPRE\_MemoryPrintUsage**(MPI\_Comm comm, HYPRE\_Int level, const char \*function, HYPRE\_Int line)

Prints the memory usage of the current process.

This function prints the memory usage details of the process to standard output. It provides information such as the virtual memory size, resident set size, and other related statistics including GPU memory usage for device builds.

#### Note

The function is designed to be platform-independent but may provide different levels of detail depending on the underlying operating system (e.g., Linux, macOS). However, this function does not lead to correct memory usage statistics on Windows platforms.

#### Parameters

- **comm** – [in] The MPI communicator. This parameter allows the function to print memory usage information for the process within the context of an MPI program.
- **level** – [in] The level of detail in the memory statistics output.
  - 1 : Display memory usage statistics for each MPI rank.
  - 2 : Display aggregate memory usage statistics over MPI ranks.
- **function** – [in] The name of the function from which `HYPRE_MemoryPrintUsage` is called. This is typically set to `__func__`, which automatically captures the name of the calling function. This variable can also be used to denote a region name.
- **line** – [in] The line number in the source file where `HYPRE_MemoryPrintUsage` is called. This is typically set to `__LINE__`, which automatically captures the line number. The line number can be omitted by passing a negative value to this variable.

**Returns**

Returns an integer status code. 0 indicates success, while a non-zero value indicates an error occurred.

**Library Version Information**

`HYPRE_Int HYPRE_Version(char **version_ptr)`

Allocates and returns a string with version number information in it.

`HYPRE_Int HYPRE_VersionNumber(HYPRE_Int *major_ptr, HYPRE_Int *minor_ptr, HYPRE_Int *patch_ptr, HYPRE_Int *single_ptr)`

Returns version number information in integer form.

Use 'NULL' for values not needed. The argument `{\tt single}` is a single sortable integer representation of the release number.

**HYPRE\_VERSION****Umpire and GPU Memory Pooling**

`HYPRE_Int HYPRE_SetUmpireDevicePoolSize(size_t nbytes)`

Sets the size of the Umpire device memory pool.

**Parameters**

**nbytes** – [in] The size of the device memory pool in bytes.

**Returns**

Returns hypre's global error code, where 0 indicates success.

`HYPRE_Int HYPRE_SetUmpireUMPoolSize(size_t nbytes)`

Sets the size of the Umpire unified memory pool.

**Parameters**

**nbytes** – [in] The size of the unified memory pool in bytes.

**Returns**

Returns hypre's global error code, where 0 indicates success.

`HYPRE_Int HYPRE_SetUmpireHostPoolSize(size_t nbytes)`

Sets the size of the Umpire host memory pool.

**Parameters**

**nbytes** – [in] The size of the host memory pool in bytes.

**Returns**

Returns hypr's global error code, where 0 indicates success.

HYPRE\_Int **HYPRE\_SetUmpirePinnedPoolSize**(size\_t nbytes)

Sets the size of the Umpire pinned memory pool.

**Parameters**

**nbytes** – [in] The size of the pinned memory pool in bytes.

**Returns**

Returns hypr's global error code, where 0 indicates success.

HYPRE\_Int **HYPRE\_SetUmpireDevicePoolName**(const char \*pool\_name)

Sets the name of the Umpire device memory pool.

**Parameters**

**pool\_name** – [in] The name to assign to the device memory pool.

**Returns**

Returns hypr's global error code, where 0 indicates success.

HYPRE\_Int **HYPRE\_SetUmpireUMPoolName**(const char \*pool\_name)

Sets the name of the Umpire unified memory pool.

**Parameters**

**pool\_name** – [in] The name to assign to the unified memory pool.

**Returns**

Returns hypr's global error code, where 0 indicates success.

HYPRE\_Int **HYPRE\_SetUmpireHostPoolName**(const char \*pool\_name)

Sets the name of the Umpire host memory pool.

**Parameters**

**pool\_name** – [in] The name to assign to the host memory pool.

**Returns**

Returns hypr's global error code, where 0 indicates success.

HYPRE\_Int **HYPRE\_SetUmpirePinnedPoolName**(const char \*pool\_name)

Sets the name of the Umpire pinned memory pool.

**Parameters**

**pool\_name** – [in] The name to assign to the pinned memory pool.

**Returns**

Returns hypr's global error code, where 0 indicates success.

**Miscellaneous**

HYPRE\_Int **HYPRE\_SetLogLevel**(HYPRE\_Int log\_level)

Sets the logging level for the HYPRE library.

The following options are available for *log\_level*:

- 0 : (default) No messaging.
- 1 : Display memory usage statistics for each MPI rank.
- 2 : Display aggregate memory usage statistics over MPI ranks.

**Note**

Log level codes can be combined using bitwise OR to enable multiple logging behaviors simultaneously.

**Parameters**

**log\_level** – The logging level to set.

**Returns**

Returns hypr's global error code, where 0 indicates success.

HYPRE\_Int **HYPRE\_SetSpTransUseVendor**(HYPRE\_Int use\_vendor)

Specifies the algorithm used for sparse matrix transposition in device builds.

The following options are available for *use\_vendor*:

- 0 : Use hypr's internal implementation.
- 1 : (default) Use the vendor library's implementation. This includes:
  - cuSPARSE for CUDA (HYPRE\_USING\_CUSPARSE)
  - rocSPARSE for HIP (HYPRE\_USING\_ROCSPARSE)
  - oneMKL for SYCL (HYPRE\_USING\_ONEMKLSPARSE)

**Parameters**

**use\_vendor** – Indicates whether to use the internal or vendor-provided implementation.

**Returns**

Returns hypr's global error code, where 0 indicates success.

HYPRE\_Int **HYPRE\_SetSpMVUseVendor**(HYPRE\_Int use\_vendor)

Specifies the algorithm used for sparse matrix/vector multiplication in device builds.

The following options are available for *use\_vendor*:

- 0 : Use hypr's internal implementation.
- 1 : (default) Use the vendor library's implementation. This includes:
  - cuSPARSE for CUDA (HYPRE\_USING\_CUSPARSE)
  - rocSPARSE for HIP (HYPRE\_USING\_ROCSPARSE)
  - oneMKL for SYCL (HYPRE\_USING\_ONEMKLSPARSE)

**Parameters**

**use\_vendor** – Indicates whether to use the internal or vendor-provided implementation.

**Returns**

Returns hypr's global error code, where 0 indicates success.

HYPRE\_Int **HYPRE\_SetSpGemmUseVendor**(HYPRE\_Int use\_vendor)

Specifies the algorithm used for sparse matrix/matrix multiplication in device builds.

The following options are available for *use\_vendor*:

- 0 : Use hypr's internal implementation.
- 1 : Use the vendor library's implementation. This includes:
  - cuSPARSE for CUDA (HYPRE\_USING\_CUSPARSE)
  - rocSPARSE for HIP (HYPRE\_USING\_ROCSPARSE)
  - oneMKL for SYCL (HYPRE\_USING\_ONEMKLSPARSE)

**Note**

The default value is 1, except for CUDA builds, which is zero.

**Parameters**

**use\_vendor** – Indicates whether to use the internal or vendor-provided implementation.

**Returns**

Returns hypr's global error code, where 0 indicates success.

HYPRE\_Int **HYPRE\_SetUseGpuRand**(HYPRE\_Int use\_curand)

Specifies the algorithm used for generating random numbers in device builds.

The following options are available for *use\_curand*:

- 0 : random numbers are generated on the host and copied to device memory.
- 1 : (default) Use the vendor library's implementation. This includes:
  - cuSPARSE for CUDA (HYPRE\_USING\_CUSPARSE)
  - rocSPARSE for HIP (HYPRE\_USING\_ROCSPARSE)
  - oneMKL for SYCL (HYPRE\_USING\_ONEMKLSPARSE)

**Parameters**

**use\_curand** – Indicates whether to use the vendor-provided implementation or not.

**Returns**

Returns hypr's global error code, where 0 indicates success.

HYPRE\_Int **HYPRE\_SetGpuAwareMPI**(HYPRE\_Int use\_gpu\_aware\_mpi)

Configures the usage of GPU-aware MPI for communication in device builds.

The following options are available for *use\_gpu\_aware\_mpi*:

- 0 : MPI buffers are transferred between device and host memory. Communication occurs on the host.
- 1 : MPI communication is performed directly from the device using device-resident buffers.

**Note**

This option requires hypr to be configured with GPU-aware MPI support for it to take effect.

**Parameters**

**use\_gpu\_aware\_mpi** – Specifies whether to enable GPU-aware MPI communication or not.

**Returns**

Returns hypr's global error code, where 0 indicates success.

**HYPRE\_SetSpGemmUseCuspars**(use\_vendor)

## BIBLIOGRAPHY

- [AsFa1996] S. F. Ashby and R. D. Falgout. A parallel multigrid preconditioned conjugate gradient algorithm for groundwater flow simulations. *Nuclear Science and Engineering*, 124(1):145–159, September 1996. Also available as LLNL Technical Report UCRL-JC-122359.
- [BFKY2011] A. Baker, R. Falgout, T. Kolev, and U. M. Yang. Multigrid smoothers for ultra-parallel computing. *SIAM J. on Sci. Comp.*, 33:2864–2887, 2011. Also available as LLNL technical report LLLNL-JRNL-473191.
- [BaFY2006] A.H. Baker, R.D. Falgout, and U.M. Yang. An assumed partition algorithm for determining processor inter-communication. *Parallel Computing*, 32:394–414, 2006.
- [BaKY2010] A. Baker, T. Kolev, and U. M. Yang. Improving algebraic multigrid interpolation operators for linear elasticity problems. *Numer. Linear Algebra Appl.*, 17:495–517, 2010. Also available as LLNL technical report LLLNL-JRNL-412928.
- [BKRHSMTY2021] Luc Berger-Vergiat, Brian Kelley, Sivasankaran Rajamanickam, Jonathan Hu, Katarzyna Swirydowicz, Paul Mulleney, Stephen Thomas, Ichitaro Yamazaki. Two-Stage Gauss–Seidel Preconditioners and Smoothers for Krylov Solvers on a GPU cluster. <https://arxiv.org/abs/2104.01196>.
- [BLOPEWeb] BLOPEX, parallel preconditioned eigenvalue solvers. <http://code.google.com/p/blopex/>.
- [BrFJ2000] P. N. Brown, R. D. Falgout, and J. E. Jones. Semicoarsening multigrid on distributed memory machines. *SIAM J. Sci. Comput.*, 21(5):1823–1834, 2000. Special issue on the Fifth Copper Mountain Conference on Iterative Methods. Also available as LLNL technical report UCRL-JC-130720.
- [Chow2000] E. Chow. A priori sparsity patterns for parallel sparse approximate inverse preconditioners. *SIAM J. Sci. Comput.*, 21:1804–1822, 2000.
- [CMakeWeb] CMake, a cross-platform open-source build system. <http://www.cmake.org/>.
- [DFNY2008] H. De Sterck, R. Falgout, J. Nolting, and U. M. Yang. Distance-two interpolation for parallel algebraic multigrid. *Numer. Linear Algebra Appl.*, 15:115–139, 2008. Also available as LLNL technical report UCRL-JRNL-230844.
- [DeYH2004] H. De Sterck, U. M. Yang, and J. Heys. Reducing complexity in parallel algebraic multigrid preconditioners. *SIAM Journal on Matrix Analysis and Applications*, 27:1019–1039, 2006. Also available as LLNL technical report UCRL-JRNL-206780.
- [FaJo2000] R. D. Falgout and J. E. Jones. Multigrid on massively parallel architectures. In E. Dick, K. Rienslagh, and J. Vierendeels, editors, *Multigrid Methods VI*, volume 14 of *Lecture Notes in Computational Science and Engineering*, pages 101–107, Berlin, 2000. Springer. Proc. of the Sixth European Multigrid Conference held in Gent, Belgium, September 27–30, 1999. Also available as LLNL technical report UCRL-JC-133948.
- [FaJY2004] R. D. Falgout, J. E. Jones, and U. M. Yang. The design and implementation of hypre, a library of parallel high performance preconditioners. In A. M. Bruaset and A. Tveito, editors, *Numerical Solution*

- of Partial Differential Equations on Parallel Computers*, pages 267–294. Springer–Verlag, 2006. Also available as LLNL technical report UCRL-JRNL-205459.
- [FaJY2005] R. D. Falgout, J. E. Jones, and U. M. Yang. Conceptual interfaces in hypr. *Future Generation Computer Systems*, 22:239–251, 2006. Special issue on PDE software. Also available as LLNL technical report UCRL-JC-148957.
- [FaSc2014] Robert D. Falgout and Jacob B. Schroder. Non-galerkin coarse grids for algebraic multigrid. *SIAM J. Sci. Comput.*, 36(3):309–334, 2014.
- [GrKo2015] A. Grayver and Tz. Kolev. Large-scale 3D geoelectromagnetic modeling using parallel adaptive high-order finite element method. *Geophysics*, 80(6):E277–E291, 2015. Also available as LLNL technical report LLNL-JRNL-665742.
- [GrMS2006a] M. Griebel, B. Metsch, and M. A. Schweitzer. Coarse grid classification: A parallel coarsening scheme for algebraic multigrid methods. *Numerical Linear Algebra with Applications*, 13(2–3):193–214, 2006. Also available as SFB 611 preprint No. 225, Universität Bonn, 2005.
- [GrMS2006b] M. Griebel, B. Metsch, and M. A. Schweitzer. Coarse grid classification - Part II: Automatic coarse grid agglomeration for parallel AMG. Preprint No. 271, Sonderforschungsbereich 611, Universität Bonn, 2006.
- [HeYa2002] V. E. Henson and U. M. Yang. BoomerAMG: a parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics*, 41(5):155–177, 2002. Also available as LLNL technical report UCRL-JC-141495.
- [HiXu2006] R. Hiptmair and J. Xu. Nodal auxiliary space preconditioning in  $H(\text{curl})$  and  $H(\text{div})$  spaces. *Numer. Math.*, 2006.
- [HyPo1999] D. Hysom and A. Pothen. Efficient parallel computation of ILU(k) preconditioners. In *Proceedings of Supercomputing '99*. ACM, November 1999. Published on CDROM, ISBN #1-58113-091-0, ACM Order #415990, IEEE Computer Society Press Order # RS00197.
- [HyPo2001] D. Hysom and A. Pothen. A scalable parallel algorithm for incomplete factor preconditioning. *SIAM J. Sci. Comput.*, 22(6):2194–2215, 2001.
- [KaKu1998] G. Karypis and V. Kumar. Parallel threshold-based ILU factorization. Technical Report 061, University of Minnesota, Department of Computer Science/Army HPC Research Center, Minneapolis, MN 5455, 1998.
- [Knya2001] A. Knyazev. Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. *SIAM J. Sci. Comput.*, 23(2):517–541, 2001.
- [KLAO2007] A. Knyazev, I. Lashuk, M. Argentati, and E. Ovchinnikov. Block locally optimal preconditioned eigenvalue solvers (blopex) in hypr and petsc. *SIAM J. Sci. Comput.*, 25(5):2224–2239, 2007.
- [KoVa2009] Tz. Kolev and P. Vassilevski. Parallel auxiliary space AMG for  $H(\text{curl})$  problems. *J. Comput. Math.*, 27:604–623, 2009. Special issue on Adaptive and Multilevel Methods in Electromagnetics. UCRL-JRNL-237306.
- [KoYe1993] L. Yu. Kolotilina and A. Yu. Yeremin. Factorized Sparse Approximate Inverse Preconditionings I. Theory. *SIAM J. Matrix Anal. A.*, 14(1):45–58, 1993. <https://doi.org/10.1137/0614004>.
- [LiSY2021] R. Li, B. Sjögreen and U. M. Yang. A new class of AMG interpolation methods based on matrix-matrix multiplications. *SIAM J. Sci. Comput.*, 43(5), S540–S564.
- [JaFe2015] C. Janna, M. Ferronato, F. Sartoretto and G. Gambolati. FSAIPACK: A Software Package for High-Performance Factored Sparse Approximate Inverse Preconditioning. *ACM T. Math. Software*, 41(2):1–26, 2015. <https://doi.org/10.1145/2629475>.
- [McCo1989] S. F. McCormick. *Multilevel Adaptive Methods for Partial Differential Equations*, volume 6 of *Frontiers in Applied Mathematics*. SIAM Books, Philadelphia, 1989.

- [MoRS1998] J.E. Morel, Randy M. Roberts, and Mikhail J. Shashkov. A local support-operators diffusion discretization scheme for quadrilateral r-z meshes. *J. Comp. Physics*, 144:17–51, 1998.
- [PaFa2019] V. A. Paludetto Magri, A. Franceschini and C. Janna. A novel algebraic multigrid approach based on adaptive smoothing and prolongation for ill-conditioned systems. *SIAM J. Sci. Comput.*, 41(1):A190–A219, 2019. <https://doi.org/10.1137/17M1161178>.
- [RuSt1987] J. W. Ruge and K. Stüben. Algebraic multigrid (AMG). In S. F. McCormick, editor, *Multigrid Methods*, volume 3 of *Frontiers in Applied Mathematics*, pages 73–130. SIAM, Philadelphia, PA, 1987.
- [Scha1998] S. Schaffer. A semi-coarsening multigrid method for elliptic partial differential equations with highly discontinuous and anisotropic coefficients. *SIAM J. Sci. Comput.*, 20(1):228–242, 1998.
- [MaFaYa23] V. A. P. Magri, R. Falgout and U. M. Yang. A new semistructured algebraic multigrid method. *SIAM J. Sci. Comput.*, 45(3):S439-S460, 2023.
- [Stue1999] K. Stüben. Algebraic multigrid (AMG): an introduction with applications. In U. Trottenberg, C. Oosterlee, and A. Schüller, editors, *Multigrid*. Academic Press, 2001.
- [Umpire] Umpire: Managing Heterogeneous Memory Resources. <https://github.com/LLNL/Umpire>.
- [VaMB1996] P. Vaněk, J. Mandel, and M. Brezina. Algebraic multigrid based on smoothed aggregation for second and fourth order problems. *Computing*, 56:179–196, 1996.
- [VaBM2001] P. Vaněk, M. Brezina, and J. Mandel. Convergence of algebraic multigrid based on smoothed aggregation. *Numerische Mathematik*, 88:559–579, 2001.
- [VaYa2014] P. Vassilevski and U. M. Yang. Reducing communication in algebraic multigrid using additive variants. *Numer. Linear Algebra Appl.*, 21:275–296, 2014. Also available as LLNL technical report LLLNL-JRNL-637872.
- [Yang2004] U. M. Yang. On the use of relaxation parameters in hybrid smoothers. *Numerical Linear Algebra with Applications*, 11:155–172, 2004.
- [Yang2005] U. M. Yang. Parallel algebraic multigrid methods - high performance preconditioners. In A. M. Bruaset and A. Tveito, editors, *Numerical Solution of Partial Differential Equations on Parallel Computers*, pages 209–236. Springer-Verlag, 2006. Also available as LLNL technical report UCRL-BOOK-208032.
- [Yang2010] U. M. Yang. On long range interpolation operators for aggressive coarsening. *Numer. Linear Algebra Appl.*, 17:453–472, 2010. Also available as LLNL technical report LLLNL-JRNL-417371.



## Symbols

- `_HYPRE_ExecutionPolicy` (C++ *enum*), 213
  - `_HYPRE_ExecutionPolicy::HYPRE_EXEC_DEVICE` (C++ *enumerator*), 213
  - `_HYPRE_ExecutionPolicy::HYPRE_EXEC_HOST` (C++ *enumerator*), 213
  - `_HYPRE_ExecutionPolicy::HYPRE_EXEC_UNDEFINED` (C++ *enumerator*), 213
  - `_HYPRE_MemoryLocation` (C++ *enum*), 213
  - `_HYPRE_MemoryLocation::HYPRE_MEMORY_DEVICE` (C++ *enumerator*), 213
  - `_HYPRE_MemoryLocation::HYPRE_MEMORY_HOST` (C++ *enumerator*), 213
  - `_HYPRE_MemoryLocation::HYPRE_MEMORY_UNDEFINED` (C++ *enumerator*), 213
- G**
- `GenerateDifConv` (C++ *function*), 197
  - `GenerateLaplacian` (C++ *function*), 197
  - `GenerateLaplacian27pt` (C++ *function*), 197
  - `GenerateLaplacian9pt` (C++ *function*), 197
  - `GenerateRotate7pt` (C++ *function*), 197
  - `GenerateRSVarDifConv` (C++ *function*), 197
  - `GenerateVarDifConv` (C++ *function*), 197
- H**
- `HYPRE_ADSCreate` (C++ *function*), 168
  - `HYPRE_ADSDestroy` (C++ *function*), 168
  - `HYPRE_ADSSetFinalRelativeResidualNorm` (C++ *function*), 171
  - `HYPRE_ADSSetNumIterations` (C++ *function*), 171
  - `HYPRE_ADSSetAMGOptions` (C++ *function*), 171
  - `HYPRE_ADSSetAMGOptions` (C++ *function*), 171
  - `HYPRE_ADSSetChebySmoothingOptions` (C++ *function*), 171
  - `HYPRE_ADSSetCoordinateVectors` (C++ *function*), 169
  - `HYPRE_ADSSetCycleType` (C++ *function*), 170
  - `HYPRE_ADSSetDiscreteCurl` (C++ *function*), 169
  - `HYPRE_ADSSetDiscreteGradient` (C++ *function*), 169
  - `HYPRE_ADSSetInterpolations` (C++ *function*), 169
  - `HYPRE_ADSSetMaxIter` (C++ *function*), 170
  - `HYPRE_ADSSetPrintLevel` (C++ *function*), 170
  - `HYPRE_ADSSetSmoothingOptions` (C++ *function*), 170
  - `HYPRE_ADSSetTol` (C++ *function*), 170
  - `HYPRE_ADSSetup` (C++ *function*), 168
  - `HYPRE_ADSSolve` (C++ *function*), 169
  - `HYPRE_AMECreate` (C++ *function*), 198
  - `HYPRE_AMEDestroy` (C++ *function*), 198
  - `HYPRE_AMEGetEigenvalues` (C++ *function*), 198
  - `HYPRE_AMEGetEigenvectors` (C++ *function*), 198
  - `HYPRE_AMESetAMSSolver` (C++ *function*), 198
  - `HYPRE_AMESetBlockSize` (C++ *function*), 198
  - `HYPRE_AMESetMassMatrix` (C++ *function*), 198
  - `HYPRE_AMESetMaxIter` (C++ *function*), 198
  - `HYPRE_AMESetMaxPCGIter` (C++ *function*), 198
  - `HYPRE_AMESetPrintLevel` (C++ *function*), 198
  - `HYPRE_AMESetRTol` (C++ *function*), 198
  - `HYPRE_AMESetTol` (C++ *function*), 198
  - `HYPRE_AMESetup` (C++ *function*), 198
  - `HYPRE_AMESolve` (C++ *function*), 198
  - `HYPRE_AMSConstructDiscreteGradient` (C++ *function*), 168
  - `HYPRE_AMSCreate` (C++ *function*), 164
  - `HYPRE_AMSDestroy` (C++ *function*), 164
  - `HYPRE_AMSGetFinalRelativeResidualNorm` (C++ *function*), 168
  - `HYPRE_AMSGetNumIterations` (C++ *function*), 168
  - `HYPRE_AMSProjectOutGradients` (C++ *function*), 168
  - `HYPRE_AMSSetAlphaAMGCoarseRelaxType` (C++ *function*), 167
  - `HYPRE_AMSSetAlphaAMGOptions` (C++ *function*), 167
  - `HYPRE_AMSSetAlphaPoissonMatrix` (C++ *function*), 166
  - `HYPRE_AMSSetBetaAMGCoarseRelaxType` (C++ *function*), 168
  - `HYPRE_AMSSetBetaAMGOptions` (C++ *function*), 167
  - `HYPRE_AMSSetBetaPoissonMatrix` (C++ *function*), 166
  - `HYPRE_AMSSetCoordinateVectors` (C++ *function*), 165
  - `HYPRE_AMSSetCycleType` (C++ *function*), 166
  - `HYPRE_AMSSetDimension` (C++ *function*), 165
  - `HYPRE_AMSSetDiscreteGradient` (C++ *function*), 165

- HYPRE\_AMS**SetEdgeConstantVectors** (C++ *function*), 165
- HYPRE\_AMS**SetInteriorNodes** (C++ *function*), 166
- HYPRE\_AMS**SetInterpolations** (C++ *function*), 165
- HYPRE\_AMS**SetMaxIter** (C++ *function*), 166
- HYPRE\_AMS**SetPrintLevel** (C++ *function*), 167
- HYPRE\_AMS**SetProjectionFrequency** (C++ *function*), 166
- HYPRE\_AMS**SetSmoothingOptions** (C++ *function*), 167
- HYPRE\_AMS**SetTol** (C++ *function*), 166
- HYPRE\_AMS**Setup** (C++ *function*), 164
- HYPRE\_AMS**Solve** (C++ *function*), 165
- HYPRE\_BiCGSTAB**Destroy** (C++ *function*), 209
- HYPRE\_BiCGSTAB**GetFinalRelativeResidualNorm** (C++ *function*), 210
- HYPRE\_BiCGSTAB**GetNumIterations** (C++ *function*), 210
- HYPRE\_BiCGSTAB**GetPrecond** (C++ *function*), 210
- HYPRE\_BiCGSTAB**GetPrecondMatrix** (C++ *function*), 210
- HYPRE\_BiCGSTAB**GetResidual** (C++ *function*), 210
- HYPRE\_BiCGSTAB**SetAbsoluteTol** (C++ *function*), 210
- HYPRE\_BiCGSTAB**SetConvergenceFactorTol** (C++ *function*), 210
- HYPRE\_BiCGSTAB**SetLogging** (C++ *function*), 210
- HYPRE\_BiCGSTAB**SetMaxIter** (C++ *function*), 210
- HYPRE\_BiCGSTAB**SetMinIter** (C++ *function*), 210
- HYPRE\_BiCGSTAB**SetPrecond** (C++ *function*), 210
- HYPRE\_BiCGSTAB**SetPrecondMatrix** (C++ *function*), 210
- HYPRE\_BiCGSTAB**SetPrintLevel** (C++ *function*), 210
- HYPRE\_BiCGSTAB**SetStopCrit** (C++ *function*), 210
- HYPRE\_BiCGSTAB**SetTol** (C++ *function*), 209
- HYPRE\_BiCGSTAB**Setup** (C++ *function*), 209
- HYPRE\_BiCGSTAB**Solve** (C++ *function*), 209
- HYPRE\_BlockTridiag**Create** (C++ *function*), 200
- HYPRE\_BlockTridiag**Destroy** (C++ *function*), 200
- HYPRE\_BlockTridiag**SetAMGNumSweeps** (C++ *function*), 200
- HYPRE\_BlockTridiag**SetAMGRelaxType** (C++ *function*), 200
- HYPRE\_BlockTridiag**SetAMGStrengthThreshold** (C++ *function*), 200
- HYPRE\_BlockTridiag**SetIndexSet** (C++ *function*), 200
- HYPRE\_BlockTridiag**SetPrintLevel** (C++ *function*), 200
- HYPRE\_BlockTridiag**Setup** (C++ *function*), 200
- HYPRE\_BlockTridiag**Solve** (C++ *function*), 200
- HYPRE\_BoomerAMG**Create** (C++ *function*), 132
- HYPRE\_BoomerAMG**DDCreate** (C++ *function*), 154
- HYPRE\_BoomerAMG**DDDestroy** (C++ *function*), 154
- HYPRE\_BoomerAMG**DDGetAMG** (C++ *function*), 155
- HYPRE\_BoomerAMG**DDGetFinalRelativeResidualNorm** (C++ *function*), 155
- HYPRE\_BoomerAMG**DDGetNumIterations** (C++ *function*), 155
- HYPRE\_BoomerAMG**DDSetFACCycleType** (C++ *function*), 154
- HYPRE\_BoomerAMG**DDSetFACNumCycles** (C++ *function*), 154
- HYPRE\_BoomerAMG**DDSetFACNumRelax** (C++ *function*), 154
- HYPRE\_BoomerAMG**DDSetFACRelaxType** (C++ *function*), 155
- HYPRE\_BoomerAMG**DDSetFACRelaxWeight** (C++ *function*), 155
- HYPRE\_BoomerAMG**DDSetNumGhostLayers** (C++ *function*), 155
- HYPRE\_BoomerAMG**DDSetPadding** (C++ *function*), 155
- HYPRE\_BoomerAMG**DDSetStartLevel** (C++ *function*), 155
- HYPRE\_BoomerAMG**DDSetup** (C++ *function*), 154
- HYPRE\_BoomerAMG**DDSetUserFACRelaxation** (C++ *function*), 155
- HYPRE\_BoomerAMG**DSolve** (C++ *function*), 154
- HYPRE\_BoomerAMG**Destroy** (C++ *function*), 132
- HYPRE\_BoomerAMG**GetAdditive** (C++ *function*), 140
- HYPRE\_BoomerAMG**GetCoarsenCutFactor** (C++ *function*), 134
- HYPRE\_BoomerAMG**GetCoarsenType** (C++ *function*), 136
- HYPRE\_BoomerAMG**GetConvergeType** (C++ *function*), 134
- HYPRE\_BoomerAMG**GetCumNnzAP** (C++ *function*), 133
- HYPRE\_BoomerAMG**GetCumNumIterations** (C++ *function*), 133
- HYPRE\_BoomerAMG**GetCycleNumSweeps** (C++ *function*), 142
- HYPRE\_BoomerAMG**GetCycleRelaxType** (C++ *function*), 143
- HYPRE\_BoomerAMG**GetCycleType** (C++ *function*), 140
- HYPRE\_BoomerAMG**GetDebugFlag** (C++ *function*), 151
- HYPRE\_BoomerAMG**GetDomainType** (C++ *function*), 147
- HYPRE\_BoomerAMG**GetFCycle** (C++ *function*), 140
- HYPRE\_BoomerAMG**GetFilterFunctions** (C++ *function*), 133
- HYPRE\_BoomerAMG**GetFilterThresholdR** (C++ *function*), 135
- HYPRE\_BoomerAMG**GetFinalRelativeResidualNorm** (C++ *function*), 133
- HYPRE\_BoomerAMG**GetGridHierarchy** (C++ *function*), 152
- HYPRE\_BoomerAMG**GetJacobiTruncThreshold** (C++ *function*), 153
- HYPRE\_BoomerAMG**GetLogging** (C++ *function*), 151
- HYPRE\_BoomerAMG**GetMaxCoarseSize** (C++ *function*),

- 134
- HYPRE\_BoomerAMGGetMaxIter (C++ function), 134
- HYPRE\_BoomerAMGGetMaxLevels (C++ function), 134
- HYPRE\_BoomerAMGGetMaxRowSum (C++ function), 153
- HYPRE\_BoomerAMGGetMeasureType (C++ function), 136
- HYPRE\_BoomerAMGGetMinCoarseSize (C++ function), 134
- HYPRE\_BoomerAMGGetMultAdditive (C++ function), 140
- HYPRE\_BoomerAMGGetNumFunctions (C++ function), 133
- HYPRE\_BoomerAMGGetNumIterations (C++ function), 133
- HYPRE\_BoomerAMGGetOverlap (C++ function), 147
- HYPRE\_BoomerAMGGetPMaxElmts (C++ function), 138
- HYPRE\_BoomerAMGGetPostInterpType (C++ function), 153
- HYPRE\_BoomerAMGGetPrintLevel (C++ function), 151
- HYPRE\_BoomerAMGGetRedundant (C++ function), 141
- HYPRE\_BoomerAMGGetResidual (C++ function), 132
- HYPRE\_BoomerAMGGetSchwarzRlxWeight (C++ function), 148
- HYPRE\_BoomerAMGGetSeqThreshold (C++ function), 141
- HYPRE\_BoomerAMGGetSimple (C++ function), 140
- HYPRE\_BoomerAMGGetSmoothNumLevels (C++ function), 146
- HYPRE\_BoomerAMGGetSmoothNumSweeps (C++ function), 146
- HYPRE\_BoomerAMGGetSmoothType (C++ function), 145
- HYPRE\_BoomerAMGGetStrongThreshold (C++ function), 134
- HYPRE\_BoomerAMGGetStrongThresholdR (C++ function), 135
- HYPRE\_BoomerAMGGetTol (C++ function), 134
- HYPRE\_BoomerAMGGetTruncFactor (C++ function), 138
- HYPRE\_BoomerAMGGetVariant (C++ function), 147
- HYPRE\_BoomerAMGInitGridRelaxation (C++ function), 151
- HYPRE\_BoomerAMGSetAdditive (C++ function), 140
- HYPRE\_BoomerAMGSetAddLastLvl (C++ function), 140
- HYPRE\_BoomerAMGSetAddPMaxElmts (C++ function), 141
- HYPRE\_BoomerAMGSetAddRelaxType (C++ function), 141
- HYPRE\_BoomerAMGSetAddRelaxWt (C++ function), 141
- HYPRE\_BoomerAMGSetAddTruncFactor (C++ function), 141
- HYPRE\_BoomerAMGSetADropTol (C++ function), 151
- HYPRE\_BoomerAMGSetADropType (C++ function), 151
- HYPRE\_BoomerAMGSetAggInterpType (C++ function), 138
- HYPRE\_BoomerAMGSetAggNumLevels (C++ function), 136
- HYPRE\_BoomerAMGSetAggP12MaxElmts (C++ function), 139
- HYPRE\_BoomerAMGSetAggP12TruncFactor (C++ function), 139
- HYPRE\_BoomerAMGSetAggPMaxElmts (C++ function), 139
- HYPRE\_BoomerAMGSetAggTruncFactor (C++ function), 139
- HYPRE\_BoomerAMGSetCGCIts (C++ function), 137
- HYPRE\_BoomerAMGSetChebyEigEst (C++ function), 145
- HYPRE\_BoomerAMGSetChebyFraction (C++ function), 145
- HYPRE\_BoomerAMGSetChebyOrder (C++ function), 145
- HYPRE\_BoomerAMGSetChebyScale (C++ function), 145
- HYPRE\_BoomerAMGSetChebyVariant (C++ function), 145
- HYPRE\_BoomerAMGSetCoarsenCutFactor (C++ function), 134
- HYPRE\_BoomerAMGSetCoarsenType (C++ function), 135
- HYPRE\_BoomerAMGSetConvergeType (C++ function), 133
- HYPRE\_BoomerAMGSetCoordDim (C++ function), 152
- HYPRE\_BoomerAMGSetCoordinates (C++ function), 152
- HYPRE\_BoomerAMGSetCPoints (C++ function), 152
- HYPRE\_BoomerAMGSetCpointsToKeep (C++ function), 152
- HYPRE\_BoomerAMGSetCRRate (C++ function), 153
- HYPRE\_BoomerAMGSetCRStrongTh (C++ function), 153
- HYPRE\_BoomerAMGSetCRUseCG (C++ function), 154
- HYPRE\_BoomerAMGSetCumNnzAP (C++ function), 133
- HYPRE\_BoomerAMGSetCycleNumSweeps (C++ function), 142
- HYPRE\_BoomerAMGSetCycleRelaxType (C++ function), 143
- HYPRE\_BoomerAMGSetCycleType (C++ function), 140
- HYPRE\_BoomerAMGSetDebugFlag (C++ function), 151
- HYPRE\_BoomerAMGSetDofFunc (C++ function), 133
- HYPRE\_BoomerAMGSetDomainType (C++ function), 147
- HYPRE\_BoomerAMGSetDropTol (C++ function), 148
- HYPRE\_BoomerAMGSetEuBJ (C++ function), 148
- HYPRE\_BoomerAMGSetEuclidFile (C++ function), 148
- HYPRE\_BoomerAMGSetEuLevel (C++ function), 148
- HYPRE\_BoomerAMGSetEuSparseA (C++ function), 148
- HYPRE\_BoomerAMGSetFCycle (C++ function), 140
- HYPRE\_BoomerAMGSetFilter (C++ function), 148
- HYPRE\_BoomerAMGSetFilterFunctions (C++ function), 133
- HYPRE\_BoomerAMGSetFilterThresholdR (C++ function), 135

HYPRE\_BoomerAMGSetFlexibleCGCScalingFactors (C++ function), 147  
 HYPRE\_BoomerAMGSetFlexibleCycleStruct (C++ function), 146  
 HYPRE\_BoomerAMGSetFlexibleOuterWeights (C++ function), 146  
 HYPRE\_BoomerAMGSetFlexibleRelaxOrders (C++ function), 146  
 HYPRE\_BoomerAMGSetFlexibleRelaxTypes (C++ function), 146  
 HYPRE\_BoomerAMGSetFlexibleRelaxWeights (C++ function), 146  
 HYPRE\_BoomerAMGSetFPoints (C++ function), 153  
 HYPRE\_BoomerAMGSetFSAIAlgoType (C++ function), 150  
 HYPRE\_BoomerAMGSetFSAIEigMaxIters (C++ function), 150  
 HYPRE\_BoomerAMGSetFSAIKapTolerance (C++ function), 150  
 HYPRE\_BoomerAMGSetFSAILocalSolveType (C++ function), 150  
 HYPRE\_BoomerAMGSetFSAIMaxNnzRow (C++ function), 150  
 HYPRE\_BoomerAMGSetFSAIMaxSteps (C++ function), 150  
 HYPRE\_BoomerAMGSetFSAIMaxStepSize (C++ function), 150  
 HYPRE\_BoomerAMGSetFSAINumLevels (C++ function), 150  
 HYPRE\_BoomerAMGSetFSAIThreshold (C++ function), 150  
 HYPRE\_BoomerAMGSetGMRESSwitchR (C++ function), 151  
 HYPRE\_BoomerAMGSetGridRelaxPoints (C++ function), 144  
 HYPRE\_BoomerAMGSetGridRelaxType (C++ function), 142  
 HYPRE\_BoomerAMGSetGSMG (C++ function), 140  
 HYPRE\_BoomerAMGSetILUDroptol (C++ function), 149  
 HYPRE\_BoomerAMGSetILUIterSetupMaxIter (C++ function), 149  
 HYPRE\_BoomerAMGSetILUIterSetupOption (C++ function), 149  
 HYPRE\_BoomerAMGSetILUIterSetupTolerance (C++ function), 149  
 HYPRE\_BoomerAMGSetILUIterSetupType (C++ function), 149  
 HYPRE\_BoomerAMGSetILULevel (C++ function), 149  
 HYPRE\_BoomerAMGSetILULocalReordering (C++ function), 149  
 HYPRE\_BoomerAMGSetILULowerJacobiIters (C++ function), 149  
 HYPRE\_BoomerAMGSetILUMaxIter (C++ function), 149  
 HYPRE\_BoomerAMGSetILUMaxRowNnz (C++ function), 149  
 HYPRE\_BoomerAMGSetILUTriSolve (C++ function), 149  
 HYPRE\_BoomerAMGSetILUType (C++ function), 149  
 HYPRE\_BoomerAMGSetILUUpperJacobiIters (C++ function), 149  
 HYPRE\_BoomerAMGSetInterpRefine (C++ function), 138  
 HYPRE\_BoomerAMGSetInterpType (C++ function), 137  
 HYPRE\_BoomerAMGSetInterpVecAbsQTrunc (C++ function), 139  
 HYPRE\_BoomerAMGSetInterpVecQMax (C++ function), 139  
 HYPRE\_BoomerAMGSetInterpVectors (C++ function), 139  
 HYPRE\_BoomerAMGSetInterpVecVariant (C++ function), 139  
 HYPRE\_BoomerAMGSetIsolatedFPoints (C++ function), 153  
 HYPRE\_BoomerAMGSetIsTriangular (C++ function), 151  
 HYPRE\_BoomerAMGSetIStype (C++ function), 154  
 HYPRE\_BoomerAMGSetJacobiTruncThreshold (C++ function), 153  
 HYPRE\_BoomerAMGSetKeepSameSign (C++ function), 137  
 HYPRE\_BoomerAMGSetKeepTranspose (C++ function), 152  
 HYPRE\_BoomerAMGSetLevel (C++ function), 148  
 HYPRE\_BoomerAMGSetLevelNonGalerkinTol (C++ function), 136  
 HYPRE\_BoomerAMGSetLevelOuterWt (C++ function), 145  
 HYPRE\_BoomerAMGSetLevelRelaxWt (C++ function), 144  
 HYPRE\_BoomerAMGSetLogging (C++ function), 151  
 HYPRE\_BoomerAMGSetMaxCoarseSize (C++ function), 134  
 HYPRE\_BoomerAMGSetMaxIter (C++ function), 134  
 HYPRE\_BoomerAMGSetMaxLevels (C++ function), 134  
 HYPRE\_BoomerAMGSetMaxNzPerRow (C++ function), 148  
 HYPRE\_BoomerAMGSetMaxRowSum (C++ function), 135  
 HYPRE\_BoomerAMGSetMeasureType (C++ function), 136  
 HYPRE\_BoomerAMGSetMinCoarseSize (C++ function), 134  
 HYPRE\_BoomerAMGSetMinIter (C++ function), 134  
 HYPRE\_BoomerAMGSetModuleRAP2 (C++ function), 152  
 HYPRE\_BoomerAMGSetMultAdditive (C++ function), 140  
 HYPRE\_BoomerAMGSetMultAddPMaxElmts (C++ function), 141  
 HYPRE\_BoomerAMGSetMultAddTruncFactor (C++

- function*), 141
- HYPRE\_BoomerAMGSetNodal (C++ *function*), 137
- HYPRE\_BoomerAMGSetNodalDiag (C++ *function*), 137
- HYPRE\_BoomerAMGSetNodalLevels (C++ *function*), 137
- HYPRE\_BoomerAMGSetNonGalerkinTol (C++ *function*), 136
- HYPRE\_BoomerAMGSetNonGalerkTol (C++ *function*), 136
- HYPRE\_BoomerAMGSetNumCRRelaxSteps (C++ *function*), 153
- HYPRE\_BoomerAMGSetNumFunctions (C++ *function*), 133
- HYPRE\_BoomerAMGSetNumGridSweeps (C++ *function*), 141
- HYPRE\_BoomerAMGSetNumPaths (C++ *function*), 136
- HYPRE\_BoomerAMGSetNumSamples (C++ *function*), 140
- HYPRE\_BoomerAMGSetNumSweeps (C++ *function*), 142
- HYPRE\_BoomerAMGSetOldDefault (C++ *function*), 132
- HYPRE\_BoomerAMGSetOmega (C++ *function*), 144
- HYPRE\_BoomerAMGSetOuterWt (C++ *function*), 144
- HYPRE\_BoomerAMGSetOverlap (C++ *function*), 147
- HYPRE\_BoomerAMGSetPlotFileName (C++ *function*), 152
- HYPRE\_BoomerAMGSetPlotGrids (C++ *function*), 152
- HYPRE\_BoomerAMGSetPMaxElmts (C++ *function*), 138
- HYPRE\_BoomerAMGSetPostInterpType (C++ *function*), 153
- HYPRE\_BoomerAMGSetPrintFileName (C++ *function*), 151
- HYPRE\_BoomerAMGSetPrintLevel (C++ *function*), 151
- HYPRE\_BoomerAMGSetRAP2 (C++ *function*), 152
- HYPRE\_BoomerAMGSetRedundant (C++ *function*), 141
- HYPRE\_BoomerAMGSetRelaxOrder (C++ *function*), 144
- HYPRE\_BoomerAMGSetRelaxType (C++ *function*), 142
- HYPRE\_BoomerAMGSetRelaxWeight (C++ *function*), 144
- HYPRE\_BoomerAMGSetRelaxWt (C++ *function*), 144
- HYPRE\_BoomerAMGSetRestriction (C++ *function*), 150
- HYPRE\_BoomerAMGSetSabs (C++ *function*), 153
- HYPRE\_BoomerAMGSetSchwarzRlxWeight (C++ *function*), 147
- HYPRE\_BoomerAMGSetSchwarzUseNonSymm (C++ *function*), 148
- HYPRE\_BoomerAMGSetSCommPkgSwitch (C++ *function*), 135
- HYPRE\_BoomerAMGSetSepWeight (C++ *function*), 138
- HYPRE\_BoomerAMGSetSeqThreshold (C++ *function*), 141
- HYPRE\_BoomerAMGSetSetupType (C++ *function*), 136
- HYPRE\_BoomerAMGSetSimple (C++ *function*), 140
- HYPRE\_BoomerAMGSetSmoothInterpVectors (C++ *function*), 139
- HYPRE\_BoomerAMGSetSmoothNumLevels (C++ *function*), 145
- HYPRE\_BoomerAMGSetSmoothNumSweeps (C++ *function*), 146
- HYPRE\_BoomerAMGSetSmoothType (C++ *function*), 145
- HYPRE\_BoomerAMGSetStrongThreshold (C++ *function*), 134
- HYPRE\_BoomerAMGSetStrongThresholdR (C++ *function*), 135
- HYPRE\_BoomerAMGSetSym (C++ *function*), 148
- HYPRE\_BoomerAMGSetThreshold (C++ *function*), 148
- HYPRE\_BoomerAMGSetTol (C++ *function*), 134
- HYPRE\_BoomerAMGSetTruncFactor (C++ *function*), 138
- HYPRE\_BoomerAMGSetup (C++ *function*), 132
- HYPRE\_BoomerAMGSetVariant (C++ *function*), 147
- HYPRE\_BoomerAMGSolve (C++ *function*), 132
- HYPRE\_BoomerAMGSolveT (C++ *function*), 132
- HYPRE\_CGNRDestroy (C++ *function*), 210
- HYPRE\_CGNRGetFinalRelativeResidualNorm (C++ *function*), 211
- HYPRE\_CGNRGetNumIterations (C++ *function*), 211
- HYPRE\_CGNRGetPrecond (C++ *function*), 211
- HYPRE\_CGNRSetLogging (C++ *function*), 211
- HYPRE\_CGNRSetMaxIter (C++ *function*), 211
- HYPRE\_CGNRSetMinIter (C++ *function*), 211
- HYPRE\_CGNRSetPrecond (C++ *function*), 211
- HYPRE\_CGNRSetStopCrit (C++ *function*), 211
- HYPRE\_CGNRSetTol (C++ *function*), 211
- HYPRE\_CGNRSetup (C++ *function*), 210
- HYPRE\_CGNRsolve (C++ *function*), 210
- HYPRE\_CheckError (C++ *function*), 214
- HYPRE\_ClearAllErrors (C++ *function*), 214
- HYPRE\_ClearError (C++ *function*), 214
- HYPRE\_ClearErrorMessage (C++ *function*), 214
- HYPRE\_COGMRESGetCGS (C++ *function*), 209
- HYPRE\_COGMRESGetConverged (C++ *function*), 209
- HYPRE\_COGMRESGetConvergenceFactorTol (C++ *function*), 209
- HYPRE\_COGMRESGetFinalRelativeResidualNorm (C++ *function*), 209
- HYPRE\_COGMRESGetKDim (C++ *function*), 209
- HYPRE\_COGMRESGetLogging (C++ *function*), 209
- HYPRE\_COGMRESGetMaxIter (C++ *function*), 209
- HYPRE\_COGMRESGetMinIter (C++ *function*), 209
- HYPRE\_COGMRESGetNumIterations (C++ *function*), 209
- HYPRE\_COGMRESGetPrecond (C++ *function*), 209
- HYPRE\_COGMRESGetPrintLevel (C++ *function*), 209
- HYPRE\_COGMRESGetResidual (C++ *function*), 209
- HYPRE\_COGMRESGetTol (C++ *function*), 209
- HYPRE\_COGMRESGetUnroll (C++ *function*), 209
- HYPRE\_COGMRESSetAbsoluteTol (C++ *function*), 208
- HYPRE\_COGMRESSetCGS (C++ *function*), 208

HYPRE\_COGMRESGetConvergenceFactorTol (C++ function), 208  
 HYPRE\_COGMRESGetKDim (C++ function), 208  
 HYPRE\_COGMRESGetLogging (C++ function), 209  
 HYPRE\_COGMRESGetMaxIter (C++ function), 208  
 HYPRE\_COGMRESGetMinIter (C++ function), 208  
 HYPRE\_COGMRESGetModifyPC (C++ function), 209  
 HYPRE\_COGMRESGetPrecond (C++ function), 208  
 HYPRE\_COGMRESGetPrintLevel (C++ function), 209  
 HYPRE\_COGMRESGetTol (C++ function), 208  
 HYPRE\_COGMRESGetUnroll (C++ function), 208  
 HYPRE\_COGMRESGetSetup (C++ function), 208  
 HYPRE\_COGMRESGetSolve (C++ function), 208  
 HYPRE\_DescribeError (C++ function), 214  
 HYPRE\_DeviceInitialize (C++ function), 215  
 HYPRE\_ERROR\_ARG (C macro), 215  
 HYPRE\_ERROR\_CONV (C macro), 215  
 HYPRE\_ERROR\_GENERIC (C macro), 214  
 HYPRE\_ERROR\_MEMORY (C macro), 214  
 HYPRE\_EuclidCreate (C++ function), 162  
 HYPRE\_EuclidDestroy (C++ function), 162  
 HYPRE\_EuclidSetBJ (C++ function), 163  
 HYPRE\_EuclidSetILUT (C++ function), 164  
 HYPRE\_EuclidSetLevel (C++ function), 163  
 HYPRE\_EuclidSetMem (C++ function), 163  
 HYPRE\_EuclidSetParams (C++ function), 163  
 HYPRE\_EuclidSetParamsFromFile (C++ function), 163  
 HYPRE\_EuclidSetRowScale (C++ function), 164  
 HYPRE\_EuclidSetSparseA (C++ function), 163  
 HYPRE\_EuclidSetStats (C++ function), 163  
 HYPRE\_EuclidSetup (C++ function), 162  
 HYPRE\_EuclidSolve (C++ function), 163  
 HYPRE\_ExecutionPolicy (C++ type), 214  
 HYPRE\_Finalize (C++ function), 215  
 HYPRE\_Finalized (C++ function), 215  
 HYPRE\_FlexGMRESGetConverged (C++ function), 206  
 HYPRE\_FlexGMRESGetConvergenceFactorTol (C++ function), 206  
 HYPRE\_FlexGMRESGetFinalRelativeResidualNorm (C++ function), 206  
 HYPRE\_FlexGMRESGetKDim (C++ function), 206  
 HYPRE\_FlexGMRESGetLogging (C++ function), 206  
 HYPRE\_FlexGMRESGetMaxIter (C++ function), 206  
 HYPRE\_FlexGMRESGetMinIter (C++ function), 206  
 HYPRE\_FlexGMRESGetNumIterations (C++ function), 206  
 HYPRE\_FlexGMRESGetPrecond (C++ function), 206  
 HYPRE\_FlexGMRESGetPrintLevel (C++ function), 206  
 HYPRE\_FlexGMRESGetResidual (C++ function), 206  
 HYPRE\_FlexGMRESGetStopCrit (C++ function), 206  
 HYPRE\_FlexGMRESGetTol (C++ function), 206  
 HYPRE\_FlexGMRESSetAbsoluteTol (C++ function), 205  
 HYPRE\_FlexGMRESSetConvergenceFactorTol (C++ function), 205  
 HYPRE\_FlexGMRESSetKDim (C++ function), 206  
 HYPRE\_FlexGMRESSetLogging (C++ function), 206  
 HYPRE\_FlexGMRESSetMaxIter (C++ function), 205  
 HYPRE\_FlexGMRESSetMinIter (C++ function), 205  
 HYPRE\_FlexGMRESSetModifyPC (C++ function), 206  
 HYPRE\_FlexGMRESSetPrecond (C++ function), 206  
 HYPRE\_FlexGMRESSetPrintLevel (C++ function), 206  
 HYPRE\_FlexGMRESSetTol (C++ function), 205  
 HYPRE\_FlexGMRESSetup (C++ function), 205  
 HYPRE\_FlexGMRESolve (C++ function), 205  
 HYPRE\_FSAICreate (C++ function), 155  
 HYPRE\_FSAIDestroy (C++ function), 156  
 HYPRE\_FSAISetAlgoType (C++ function), 156  
 HYPRE\_FSAISetEigMaxIters (C++ function), 157  
 HYPRE\_FSAISetKapTolerance (C++ function), 157  
 HYPRE\_FSAISetLocalSolveType (C++ function), 156  
 HYPRE\_FSAISetMaxIterations (C++ function), 157  
 HYPRE\_FSAISetMaxNnzRow (C++ function), 157  
 HYPRE\_FSAISetMaxSteps (C++ function), 156  
 HYPRE\_FSAISetMaxStepSize (C++ function), 156  
 HYPRE\_FSAISetNumLevels (C++ function), 157  
 HYPRE\_FSAISetOmega (C++ function), 157  
 HYPRE\_FSAISetPrintLevel (C++ function), 157  
 HYPRE\_FSAISetThreshold (C++ function), 157  
 HYPRE\_FSAISetTolerance (C++ function), 157  
 HYPRE\_FSAISetup (C++ function), 156  
 HYPRE\_FSAISetZeroGuess (C++ function), 157  
 HYPRE\_FSAISolve (C++ function), 156  
 hypr\_GenerateCoordinates (C++ function), 197  
 HYPRE\_GetError (C++ function), 214  
 HYPRE\_GetErrorArg (C++ function), 214  
 HYPRE\_GetErrorMessage (C++ function), 214  
 HYPRE\_GetExecutionPolicy (C++ function), 214  
 HYPRE\_GetExecutionPolicyName (C++ function), 214  
 HYPRE\_GetGlobalError (C++ function), 214  
 HYPRE\_GetGlobalPrecision (C++ function), 213  
 HYPRE\_GetMemoryLocation (C++ function), 213  
 HYPRE\_GMRESGetAbsoluteTol (C++ function), 205  
 HYPRE\_GMRESGetConverged (C++ function), 205  
 HYPRE\_GMRESGetConvergenceFactorTol (C++ function), 205  
 HYPRE\_GMRESGetFinalRelativeResidualNorm (C++ function), 204  
 HYPRE\_GMRESGetKDim (C++ function), 205  
 HYPRE\_GMRESGetLogging (C++ function), 205  
 HYPRE\_GMRESGetMaxIter (C++ function), 205  
 HYPRE\_GMRESGetMinIter (C++ function), 205  
 HYPRE\_GMRESGetNumIterations (C++ function), 204  
 HYPRE\_GMRESGetPrecond (C++ function), 205  
 HYPRE\_GMRESGetPrecondMatrix (C++ function), 205  
 HYPRE\_GMRESGetPrintLevel (C++ function), 205  
 HYPRE\_GMRESGetRefSolution (C++ function), 205

- HYPRE\_GMRESGetRelChange (C++ function), 205  
 HYPRE\_GMRESGetResidual (C++ function), 204  
 HYPRE\_GMRESGetSkipRealResidualCheck (C++ function), 204  
 HYPRE\_GMRESGetStopCrit (C++ function), 205  
 HYPRE\_GMRESGetTol (C++ function), 204  
 HYPRE\_GMRESSetAbsoluteTol (C++ function), 203  
 HYPRE\_GMRESSetConvergenceFactorTol (C++ function), 203  
 HYPRE\_GMRESSetKDim (C++ function), 203  
 HYPRE\_GMRESSetLogging (C++ function), 204  
 HYPRE\_GMRESSetMaxIter (C++ function), 203  
 HYPRE\_GMRESSetMinIter (C++ function), 203  
 HYPRE\_GMRESSetPrecond (C++ function), 203  
 HYPRE\_GMRESSetPrecondMatrix (C++ function), 204  
 HYPRE\_GMRESSetPrintLevel (C++ function), 204  
 HYPRE\_GMRESSetRefSolution (C++ function), 205  
 HYPRE\_GMRESSetRelChange (C++ function), 203  
 HYPRE\_GMRESSetSkipRealResidualCheck (C++ function), 203  
 HYPRE\_GMRESSetStopCrit (C++ function), 203  
 HYPRE\_GMRESSetTol (C++ function), 203  
 HYPRE\_GMRESSetup (C++ function), 203  
 HYPRE\_GMRESolve (C++ function), 203  
 HYPRE\_IJMatrix (C++ type), 99  
 HYPRE\_IJMatrixAdd (C++ function), 103  
 HYPRE\_IJMatrixAddToValues (C++ function), 100  
 HYPRE\_IJMatrixAddToValues2 (C++ function), 100  
 HYPRE\_IJMatrixAssemble (C++ function), 100  
 HYPRE\_IJMatrixCreate (C++ function), 99  
 HYPRE\_IJMatrixDestroy (C++ function), 99  
 HYPRE\_IJMatrixGetGlobalInfo (C++ function), 101  
 HYPRE\_IJMatrixGetLocalRange (C++ function), 101  
 HYPRE\_IJMatrixGetObject (C++ function), 102  
 HYPRE\_IJMatrixGetObjectType (C++ function), 101  
 HYPRE\_IJMatrixGetRowCounts (C++ function), 100  
 HYPRE\_IJMatrixGetValues (C++ function), 101  
 HYPRE\_IJMatrixGetValues2 (C++ function), 101  
 HYPRE\_IJMatrixGetValuesAndZeroOut (C++ function), 101  
 HYPRE\_IJMatrixInitialize (C++ function), 99  
 HYPRE\_IJMatrixInitialize\_v2 (C++ function), 99  
 HYPRE\_IJMatrixMigrate (C++ function), 103  
 HYPRE\_IJMatrixNorm (C++ function), 103  
 HYPRE\_IJMatrixPartialClone (C++ function), 104  
 HYPRE\_IJMatrixPrint (C++ function), 103  
 HYPRE\_IJMatrixPrintBinary (C++ function), 103  
 HYPRE\_IJMatrixRead (C++ function), 103  
 HYPRE\_IJMatrixReadBinary (C++ function), 103  
 HYPRE\_IJMatrixReadMM (C++ function), 103  
 HYPRE\_IJMatrixSetConstantValues (C++ function), 100  
 HYPRE\_IJMatrixSetDiagOffdSizes (C++ function), 102  
 HYPRE\_IJMatrixSetEarlyAssemble (C++ function), 102  
 HYPRE\_IJMatrixSetGrowFactor (C++ function), 103  
 HYPRE\_IJMatrixSetInitAllocation (C++ function), 102  
 HYPRE\_IJMatrixSetMaxOffProcElmts (C++ function), 102  
 HYPRE\_IJMatrixSetObjectType (C++ function), 101  
 HYPRE\_IJMatrixSetOMPFlag (C++ function), 103  
 HYPRE\_IJMatrixSetPrintLevel (C++ function), 103  
 HYPRE\_IJMatrixSetRowSizes (C++ function), 102  
 HYPRE\_IJMatrixSetValues (C++ function), 99  
 HYPRE\_IJMatrixSetValues2 (C++ function), 100  
 HYPRE\_IJMatrixTranspose (C++ function), 103  
 HYPRE\_IJVector (C++ type), 104  
 HYPRE\_IJVectorAddToValues (C++ function), 105  
 HYPRE\_IJVectorAssemble (C++ function), 106  
 HYPRE\_IJVectorCreate (C++ function), 104  
 HYPRE\_IJVectorDestroy (C++ function), 104  
 HYPRE\_IJVectorGetLocalRange (C++ function), 106  
 HYPRE\_IJVectorGetObject (C++ function), 106  
 HYPRE\_IJVectorGetObjectType (C++ function), 106  
 HYPRE\_IJVectorGetValues (C++ function), 106  
 HYPRE\_IJVectorInitialize (C++ function), 105  
 HYPRE\_IJVectorInitialize\_v2 (C++ function), 105  
 HYPRE\_IJVectorInitializeShell (C++ function), 104  
 HYPRE\_IJVectorInnerProd (C++ function), 107  
 HYPRE\_IJVectorMigrate (C++ function), 107  
 HYPRE\_IJVectorPrint (C++ function), 107  
 HYPRE\_IJVectorPrintBinary (C++ function), 107  
 HYPRE\_IJVectorRead (C++ function), 107  
 HYPRE\_IJVectorReadBinary (C++ function), 107  
 HYPRE\_IJVectorSetComponent (C++ function), 105  
 HYPRE\_IJVectorSetConstantValues (C++ function), 105  
 HYPRE\_IJVectorSetData (C++ function), 104  
 HYPRE\_IJVectorSetMaxOffProcElmts (C++ function), 105  
 HYPRE\_IJVectorSetNumComponents (C++ function), 105  
 HYPRE\_IJVectorSetObjectType (C++ function), 106  
 HYPRE\_IJVectorSetPrintLevel (C++ function), 106  
 HYPRE\_IJVectorSetTags (C++ function), 104  
 HYPRE\_IJVectorSetValues (C++ function), 105  
 HYPRE\_IJVectorUpdateValues (C++ function), 106  
 HYPRE\_ILUCreate (C++ function), 193  
 HYPRE\_ILUDestroy (C++ function), 193  
 HYPRE\_ILUGetFinalRelativeResidualNorm (C++ function), 196  
 HYPRE\_ILUGetNumIterations (C++ function), 196  
 HYPRE\_ILUSetDropThreshold (C++ function), 195  
 HYPRE\_ILUSetDropThresholdArray (C++ function), 195

- HYPRE\_ILUSetIterativeSetupMaxIter (C++ function), 194
- HYPRE\_ILUSetIterativeSetupOption (C++ function), 194
- HYPRE\_ILUSetIterativeSetupTolerance (C++ function), 194
- HYPRE\_ILUSetIterativeSetupType (C++ function), 193
- HYPRE\_ILUSetLevelOfFill (C++ function), 195
- HYPRE\_ILUSetLocalReordering (C++ function), 196
- HYPRE\_ILUSetLogging (C++ function), 196
- HYPRE\_ILUSetLowerJacobiIters (C++ function), 195
- HYPRE\_ILUSetMaxIter (C++ function), 193
- HYPRE\_ILUSetMaxNnzPerRow (C++ function), 195
- HYPRE\_ILUSetNSHDropThreshold (C++ function), 195
- HYPRE\_ILUSetNSHDropThresholdArray (C++ function), 195
- HYPRE\_ILUSetPrintLevel (C++ function), 196
- HYPRE\_ILUSetSchurMaxIter (C++ function), 195
- HYPRE\_ILUSetTol (C++ function), 195
- HYPRE\_ILUSetTriSolve (C++ function), 194
- HYPRE\_ILUSetType (C++ function), 196
- HYPRE\_ILUSetup (C++ function), 193
- HYPRE\_ILUSetUpperJacobiIters (C++ function), 195
- HYPRE\_ILUSolve (C++ function), 193
- HYPRE\_Init (C macro), 215
- HYPRE\_Initialize (C++ function), 215
- HYPRE\_Initialized (C++ function), 215
- HYPRE\_Jacobi (C macro), 126
- HYPRE\_LGMRESGetAugDim (C++ function), 208
- HYPRE\_LGMRESGetConverged (C++ function), 208
- HYPRE\_LGMRESGetConvergenceFactorTol (C++ function), 207
- HYPRE\_LGMRESGetFinalRelativeResidualNorm (C++ function), 207
- HYPRE\_LGMRESGetKDim (C++ function), 208
- HYPRE\_LGMRESGetLogging (C++ function), 208
- HYPRE\_LGMRESGetMaxIter (C++ function), 208
- HYPRE\_LGMRESGetMinIter (C++ function), 208
- HYPRE\_LGMRESGetNumIterations (C++ function), 207
- HYPRE\_LGMRESGetPrecond (C++ function), 208
- HYPRE\_LGMRESGetPrintLevel (C++ function), 208
- HYPRE\_LGMRESGetResidual (C++ function), 207
- HYPRE\_LGMRESGetStopCrit (C++ function), 207
- HYPRE\_LGMRESGetTol (C++ function), 207
- HYPRE\_LGMRESSetAbsoluteTol (C++ function), 207
- HYPRE\_LGMRESSetAugDim (C++ function), 207
- HYPRE\_LGMRESSetConvergenceFactorTol (C++ function), 207
- HYPRE\_LGMRESSetKDim (C++ function), 207
- HYPRE\_LGMRESSetLogging (C++ function), 207
- HYPRE\_LGMRESSetMaxIter (C++ function), 207
- HYPRE\_LGMRESSetMinIter (C++ function), 207
- HYPRE\_LGMRESSetPrecond (C++ function), 207
- HYPRE\_LGMRESSetPrintLevel (C++ function), 207
- HYPRE\_LGMRESSetTol (C++ function), 207
- HYPRE\_LGMRESSetup (C++ function), 206
- HYPRE\_LGMRESSolve (C++ function), 206
- HYPRE\_LOBPCGCreate (C++ function), 211
- HYPRE\_LOBPCGDestroy (C++ function), 211
- HYPRE\_LOBPCGEigenvaluesHistory (C++ function), 212
- HYPRE\_LOBPCGGetPrecond (C++ function), 212
- HYPRE\_LOBPCGIterations (C++ function), 212
- HYPRE\_LOBPCGResidualNorms (C++ function), 212
- HYPRE\_LOBPCGResidualNormsHistory (C++ function), 212
- HYPRE\_LOBPCGSetMaxIter (C++ function), 212
- HYPRE\_LOBPCGSetPrecond (C++ function), 211
- HYPRE\_LOBPCGSetPrecondUsageMode (C++ function), 212
- HYPRE\_LOBPCGSetPrintLevel (C++ function), 212
- HYPRE\_LOBPCGSetRTol (C++ function), 212
- HYPRE\_LOBPCGSetTol (C++ function), 212
- HYPRE\_LOBPCGSetup (C++ function), 212
- HYPRE\_LOBPCGSetupB (C++ function), 212
- HYPRE\_LOBPCGSetupT (C++ function), 212
- HYPRE\_LOBPCGSolve (C++ function), 212
- HYPRE\_MAX\_FILE\_NAME\_LEN (C macro), 215
- HYPRE\_MAX\_MSG\_LEN (C macro), 215
- HYPRE\_MemoryLocation (C++ type), 213
- HYPRE\_MemoryPrintUsage (C++ function), 215
- HYPRE\_MGRBuildAff (C++ function), 190
- HYPRE\_MGRCreate (C++ function), 184
- HYPRE\_MGRDestroy (C++ function), 184
- HYPRE\_MGRGetCoarseGridConvergenceFactor (C++ function), 192
- HYPRE\_MGRGetFinalRelativeResidualNorm (C++ function), 193
- HYPRE\_MGRGetNumIterations (C++ function), 192
- HYPRE\_MGRSetBlockJacobiBlockSize (C++ function), 189
- HYPRE\_MGRSetBlockSize (C++ function), 186
- HYPRE\_MGRSetCoarseGridMethod (C++ function), 187
- HYPRE\_MGRSetCoarseGridPrintLevel (C++ function), 190
- HYPRE\_MGRSetCoarseSolver (C++ function), 190
- HYPRE\_MGRSetCpointsByBlock (C++ function), 185
- HYPRE\_MGRSetCpointsByContiguousBlock (C++ function), 185
- HYPRE\_MGRSetCpointsByPointMarkerArray (C++ function), 185
- HYPRE\_MGRSetCycleType (C++ function), 191
- HYPRE\_MGRSetFRelaxCycle (C++ function), 191
- HYPRE\_MGRSetFRelaxMethod (C++ function), 186
- HYPRE\_MGRSetFRelaxPrintLevel (C++ function), 190
- HYPRE\_MGRSetFSolver (C++ function), 189
- HYPRE\_MGRSetFSolverAtLevel (C++ function), 189

- HYPRE\_MGRSetGlobalSmoothCycle (C++ function), 191
- HYPRE\_MGRSetGlobalSmootherAtLevel (C++ function), 192
- HYPRE\_MGRSetGlobalSmoothType (C++ function), 191
- HYPRE\_MGRSetInterpType (C++ function), 188
- HYPRE\_MGRSetLevelFRelaxMethod (C++ function), 187
- HYPRE\_MGRSetLevelFRelaxNumFunctions (C++ function), 188
- HYPRE\_MGRSetLevelFRelaxType (C++ function), 187
- HYPRE\_MGRSetLevelInterpType (C++ function), 189
- HYPRE\_MGRSetLevelNonGalerkinMaxElmts (C++ function), 188
- HYPRE\_MGRSetLevelNumRelaxSweeps (C++ function), 189
- HYPRE\_MGRSetLevelPMaxElmts (C++ function), 193
- HYPRE\_MGRSetLevelRestrictType (C++ function), 188
- HYPRE\_MGRSetLevelSmoothIters (C++ function), 190
- HYPRE\_MGRSetLevelSmoothType (C++ function), 191
- HYPRE\_MGRSetLogging (C++ function), 190
- HYPRE\_MGRSetMaxCoarseLevels (C++ function), 186
- HYPRE\_MGRSetMaxGlobalSmoothIters (C++ function), 190
- HYPRE\_MGRSetMaxIter (C++ function), 190
- HYPRE\_MGRSetNonCpointsToFpoints (C++ function), 186
- HYPRE\_MGRSetNonGalerkinMaxElmts (C++ function), 187
- HYPRE\_MGRSetNumInterpSweeps (C++ function), 189
- HYPRE\_MGRSetNumRelaxSweeps (C++ function), 189
- HYPRE\_MGRSetNumRestrictSweeps (C++ function), 188
- HYPRE\_MGRSetPMaxElmts (C++ function), 193
- HYPRE\_MGRSetPrintLevel (C++ function), 190
- HYPRE\_MGRSetRelaxType (C++ function), 186
- HYPRE\_MGRSetReservedCoarseNodes (C++ function), 186
- HYPRE\_MGRSetReservedCpointsLevelToKeep (C++ function), 186
- HYPRE\_MGRSetRestrictType (C++ function), 188
- HYPRE\_MGRSetTol (C++ function), 190
- HYPRE\_MGRSetTruncateCoarseGridThreshold (C++ function), 190
- HYPRE\_MGRSetup (C++ function), 184
- HYPRE\_MGRSolve (C++ function), 184
- HYPRE\_ParaSailsBuildIJMatrix (C++ function), 161
- HYPRE\_ParaSailsCreate (C++ function), 159
- HYPRE\_ParaSailsDestroy (C++ function), 159
- HYPRE\_ParaSailsGetFilter (C++ function), 162
- HYPRE\_ParaSailsGetLoadbal (C++ function), 162
- HYPRE\_ParaSailsGetLogging (C++ function), 162
- HYPRE\_ParaSailsGetNlevels (C++ function), 162
- HYPRE\_ParaSailsGetReuse (C++ function), 162
- HYPRE\_ParaSailsGetSym (C++ function), 162
- HYPRE\_ParaSailsGetThresh (C++ function), 161
- HYPRE\_ParaSailsSetFilter (C++ function), 160
- HYPRE\_ParaSailsSetLoadbal (C++ function), 160
- HYPRE\_ParaSailsSetLogging (C++ function), 161
- HYPRE\_ParaSailsSetNlevels (C++ function), 162
- HYPRE\_ParaSailsSetParams (C++ function), 160
- HYPRE\_ParaSailsSetReuse (C++ function), 161
- HYPRE\_ParaSailsSetSym (C++ function), 160
- HYPRE\_ParaSailsSetThresh (C++ function), 161
- HYPRE\_ParaSailsSetup (C++ function), 159
- HYPRE\_ParaSailsSolve (C++ function), 160
- HYPRE\_ParChebyCreate (C++ function), 158
- HYPRE\_ParChebyDestroy (C++ function), 158
- HYPRE\_ParChebyGetMinMaxEigEst (C++ function), 159
- HYPRE\_ParChebySetEigEst (C++ function), 159
- HYPRE\_ParChebySetEigRatio (C++ function), 159
- HYPRE\_ParChebySetLogging (C++ function), 158
- HYPRE\_ParChebySetMaxIterations (C++ function), 158
- HYPRE\_ParChebySetMinMaxEigEst (C++ function), 159
- HYPRE\_ParChebySetOrder (C++ function), 158
- HYPRE\_ParChebySetPrintLevel (C++ function), 158
- HYPRE\_ParChebySetScale (C++ function), 159
- HYPRE\_ParChebySetTolerance (C++ function), 158
- HYPRE\_ParChebySetup (C++ function), 158
- HYPRE\_ParChebySetVariant (C++ function), 159
- HYPRE\_ParChebySetZeroGuess (C++ function), 159
- HYPRE\_ParChebySolve (C++ function), 158
- HYPRE\_ParCSRBiCGSTABCreate (C++ function), 176
- HYPRE\_ParCSRBiCGSTABDestroy (C++ function), 176
- HYPRE\_ParCSRBiCGSTABGetFinalRelativeResidualNorm (C++ function), 176
- HYPRE\_ParCSRBiCGSTABGetNumIterations (C++ function), 176
- HYPRE\_ParCSRBiCGSTABGetPrecond (C++ function), 176
- HYPRE\_ParCSRBiCGSTABGetResidual (C++ function), 176
- HYPRE\_ParCSRBiCGSTABSetAbsoluteTol (C++ function), 176
- HYPRE\_ParCSRBiCGSTABSetLogging (C++ function), 176
- HYPRE\_ParCSRBiCGSTABSetMaxIter (C++ function), 176
- HYPRE\_ParCSRBiCGSTABSetMinIter (C++ function), 176
- HYPRE\_ParCSRBiCGSTABSetPrecond (C++ function), 176
- HYPRE\_ParCSRBiCGSTABSetPrintLevel (C++ function), 176

HYPRE\_ParCSRBiCGSTABSetStopCrit (C++ function), 176  
 HYPRE\_ParCSRBiCGSTABSetTol (C++ function), 176  
 HYPRE\_ParCSRBiCGSTABSetup (C++ function), 176  
 HYPRE\_ParCSRBiCGSTABsolve (C++ function), 176  
 HYPRE\_ParCSRCreate (C++ function), 199  
 HYPRE\_ParCSRDestroy (C++ function), 200  
 HYPRE\_ParCSRGetFinalRelativeResidualNorm (C++ function), 200  
 HYPRE\_ParCSRGetNumIterations (C++ function), 200  
 HYPRE\_ParCSRGetPrecond (C++ function), 200  
 HYPRE\_ParCSRSetLogging (C++ function), 200  
 HYPRE\_ParCSRSetMaxIter (C++ function), 200  
 HYPRE\_ParCSRSetMinIter (C++ function), 200  
 HYPRE\_ParCSRSetPrecond (C++ function), 200  
 HYPRE\_ParCSRSetStopCrit (C++ function), 200  
 HYPRE\_ParCSRSetTol (C++ function), 200  
 HYPRE\_ParCSRSetup (C++ function), 200  
 HYPRE\_ParCSRSolve (C++ function), 200  
 HYPRE\_ParCSRCreate (C++ function), 173  
 HYPRE\_ParCSRDestroy (C++ function), 173  
 HYPRE\_ParCSRGetFinalRelativeResidualNorm (C++ function), 174  
 HYPRE\_ParCSRGetNumIterations (C++ function), 174  
 HYPRE\_ParCSRGetPrecond (C++ function), 174  
 HYPRE\_ParCSRGetResidual (C++ function), 174  
 HYPRE\_ParCSRSetAbsoluteTol (C++ function), 173  
 HYPRE\_ParCSRSetCGS (C++ function), 173  
 HYPRE\_ParCSRSetKDim (C++ function), 173  
 HYPRE\_ParCSRSetLogging (C++ function), 174  
 HYPRE\_ParCSRSetMaxIter (C++ function), 174  
 HYPRE\_ParCSRSetMinIter (C++ function), 173  
 HYPRE\_ParCSRSetPrecond (C++ function), 174  
 HYPRE\_ParCSRSetPrintLevel (C++ function), 174  
 HYPRE\_ParCSRSetTol (C++ function), 173  
 HYPRE\_ParCSRSetUnroll (C++ function), 173  
 HYPRE\_ParCSRSetup (C++ function), 173  
 HYPRE\_ParCSRSolve (C++ function), 173  
 HYPRE\_ParCSRDiagScale (C++ function), 172  
 HYPRE\_ParCSRDiagScaleSetup (C++ function), 172  
 HYPRE\_ParCSRFlexGMRESCreate (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESDestroy (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESGetFinalRelativeResidualNorm (C++ function), 175  
 HYPRE\_ParCSRFlexGMRESGetNumIterations (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESGetPrecond (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESGetResidual (C++ function), 175  
 HYPRE\_ParCSRFlexGMRESSetAbsoluteTol (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESSetKDim (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESSetLogging (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESSetMaxIter (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESSetMinIter (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESSetModifyPC (C++ function), 175  
 HYPRE\_ParCSRFlexGMRESSetPrecond (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESSetPrintLevel (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESSetTol (C++ function), 174  
 HYPRE\_ParCSRFlexGMRESSetup (C++ function), 174  
 HYPRE\_ParCSRGMRESCreate (C++ function), 172  
 HYPRE\_ParCSRGMRESDestroy (C++ function), 172  
 HYPRE\_ParCSRGMRESGetFinalRelativeResidualNorm (C++ function), 173  
 HYPRE\_ParCSRGMRESGetNumIterations (C++ function), 173  
 HYPRE\_ParCSRGMRESGetPrecond (C++ function), 173  
 HYPRE\_ParCSRGMRESGetRefSolution (C++ function), 173  
 HYPRE\_ParCSRGMRESGetResidual (C++ function), 173  
 HYPRE\_ParCSRGMRESSetAbsoluteTol (C++ function), 173  
 HYPRE\_ParCSRGMRESSetKDim (C++ function), 173  
 HYPRE\_ParCSRGMRESSetLogging (C++ function), 173  
 HYPRE\_ParCSRGMRESSetMaxIter (C++ function), 173  
 HYPRE\_ParCSRGMRESSetMinIter (C++ function), 173  
 HYPRE\_ParCSRGMRESSetPrecond (C++ function), 173  
 HYPRE\_ParCSRGMRESSetPrintLevel (C++ function), 173  
 HYPRE\_ParCSRGMRESSetRefSolution (C++ function), 172  
 HYPRE\_ParCSRGMRESSetStopCrit (C++ function), 173  
 HYPRE\_ParCSRGMRESSetTol (C++ function), 173  
 HYPRE\_ParCSRGMRESSetup (C++ function), 172  
 HYPRE\_ParCSRGMRESsolve (C++ function), 172  
 HYPRE\_ParCSRHybridCreate (C++ function), 176  
 HYPRE\_ParCSRHybridDestroy (C++ function), 176  
 HYPRE\_ParCSRHybridGetDSCGNumIterations (C++ function), 184  
 HYPRE\_ParCSRHybridGetFinalRelativeResidualNorm

- (C++ function), 184
- HYPRE\_ParCSRHybridGetNumIterations (C++ function), 184
- HYPRE\_ParCSRHybridGetPCGNumIterations (C++ function), 184
- HYPRE\_ParCSRHybridGetRecomputeResidual (C++ function), 177
- HYPRE\_ParCSRHybridGetRecomputeResidualP (C++ function), 178
- HYPRE\_ParCSRHybridGetSetupSolveTime (C++ function), 184
- HYPRE\_ParCSRHybridSetAbsoluteTol (C++ function), 177
- HYPRE\_ParCSRHybridSetAggInterpType (C++ function), 183
- HYPRE\_ParCSRHybridSetAggNumLevels (C++ function), 183
- HYPRE\_ParCSRHybridSetCoarsenType (C++ function), 179
- HYPRE\_ParCSRHybridSetConvergenceTol (C++ function), 177
- HYPRE\_ParCSRHybridSetCycleNumSweeps (C++ function), 179
- HYPRE\_ParCSRHybridSetCycleRelaxType (C++ function), 180
- HYPRE\_ParCSRHybridSetCycleType (C++ function), 179
- HYPRE\_ParCSRHybridSetDofFunc (C++ function), 183
- HYPRE\_ParCSRHybridSetDSCGMaxIter (C++ function), 177
- HYPRE\_ParCSRHybridSetFlexibleCGCScalingFactors (C++ function), 182
- HYPRE\_ParCSRHybridSetFlexibleCycleStruct (C++ function), 182
- HYPRE\_ParCSRHybridSetFlexibleOuterWeights (C++ function), 182
- HYPRE\_ParCSRHybridSetFlexibleRelaxOrders (C++ function), 182
- HYPRE\_ParCSRHybridSetFlexibleRelaxTypes (C++ function), 182
- HYPRE\_ParCSRHybridSetFlexibleRelaxWeights (C++ function), 182
- HYPRE\_ParCSRHybridSetGridRelaxPoints (C++ function), 179
- HYPRE\_ParCSRHybridSetGridRelaxType (C++ function), 179
- HYPRE\_ParCSRHybridSetInterpType (C++ function), 179
- HYPRE\_ParCSRHybridSetKDim (C++ function), 178
- HYPRE\_ParCSRHybridSetKeepTranspose (C++ function), 183
- HYPRE\_ParCSRHybridSetLevelOuterWt (C++ function), 181
- HYPRE\_ParCSRHybridSetLevelRelaxWt (C++ function), 181
- HYPRE\_ParCSRHybridSetLogging (C++ function), 178
- HYPRE\_ParCSRHybridSetMaxCoarseSize (C++ function), 181
- HYPRE\_ParCSRHybridSetMaxLevels (C++ function), 178
- HYPRE\_ParCSRHybridSetMaxRowSum (C++ function), 178
- HYPRE\_ParCSRHybridSetMeasureType (C++ function), 179
- HYPRE\_ParCSRHybridSetMinCoarseSize (C++ function), 181
- HYPRE\_ParCSRHybridSetNodal (C++ function), 183
- HYPRE\_ParCSRHybridSetNonGalerkinTol (C++ function), 183
- HYPRE\_ParCSRHybridSetNumFunctions (C++ function), 183
- HYPRE\_ParCSRHybridSetNumGridSweeps (C++ function), 184
- HYPRE\_ParCSRHybridSetNumPaths (C++ function), 183
- HYPRE\_ParCSRHybridSetNumSweeps (C++ function), 179
- HYPRE\_ParCSRHybridSetOmega (C++ function), 182
- HYPRE\_ParCSRHybridSetOuterWt (C++ function), 181
- HYPRE\_ParCSRHybridSetPCGMaxIter (C++ function), 177
- HYPRE\_ParCSRHybridSetPMaxElmts (C++ function), 178
- HYPRE\_ParCSRHybridSetPrecond (C++ function), 178
- HYPRE\_ParCSRHybridSetPrintLevel (C++ function), 178
- HYPRE\_ParCSRHybridSetRecomputeResidual (C++ function), 177
- HYPRE\_ParCSRHybridSetRecomputeResidualP (C++ function), 177
- HYPRE\_ParCSRHybridSetRelaxOrder (C++ function), 180
- HYPRE\_ParCSRHybridSetRelaxType (C++ function), 180
- HYPRE\_ParCSRHybridSetRelaxWeight (C++ function), 182
- HYPRE\_ParCSRHybridSetRelaxWt (C++ function), 181
- HYPRE\_ParCSRHybridSetRelChange (C++ function), 178
- HYPRE\_ParCSRHybridSetSeqThreshold (C++ function), 181
- HYPRE\_ParCSRHybridSetSetupType (C++ function), 177
- HYPRE\_ParCSRHybridSetSolverType (C++ function), 177
- HYPRE\_ParCSRHybridSetStopCrit (C++ function), 178
- HYPRE\_ParCSRHybridSetStrongThreshold (C++ function), 181

- function*), 178
- HYPRE\_ParCSRHybridSetTol (C++ *function*), 177
- HYPRE\_ParCSRHybridSetTruncFactor (C++ *function*), 178
- HYPRE\_ParCSRHybridSetTwoNorm (C++ *function*), 178
- HYPRE\_ParCSRHybridSetup (C++ *function*), 176
- HYPRE\_ParCSRHybridSolve (C++ *function*), 177
- HYPRE\_ParCSRLGMRESCreate (C++ *function*), 175
- HYPRE\_ParCSRLGMRESDestroy (C++ *function*), 175
- HYPRE\_ParCSRLGMRESGetFinalRelativeResidualNorm (C++ *function*), 175
- HYPRE\_ParCSRLGMRESGetNumIterations (C++ *function*), 175
- HYPRE\_ParCSRLGMRESGetPrecond (C++ *function*), 175
- HYPRE\_ParCSRLGMRESGetResidual (C++ *function*), 175
- HYPRE\_ParCSRLGMRESSetAbsoluteTol (C++ *function*), 175
- HYPRE\_ParCSRLGMRESSetAugDim (C++ *function*), 175
- HYPRE\_ParCSRLGMRESSetKDim (C++ *function*), 175
- HYPRE\_ParCSRLGMRESSetLogging (C++ *function*), 175
- HYPRE\_ParCSRLGMRESSetMaxIter (C++ *function*), 175
- HYPRE\_ParCSRLGMRESSetMinIter (C++ *function*), 175
- HYPRE\_ParCSRLGMRESSetPrecond (C++ *function*), 175
- HYPRE\_ParCSRLGMRESSetPrintLevel (C++ *function*), 175
- HYPRE\_ParCSRLGMRESSetTol (C++ *function*), 175
- HYPRE\_ParCSRLGMRESSetup (C++ *function*), 175
- HYPRE\_ParCSRLGMRESSolve (C++ *function*), 175
- HYPRE\_ParCSRMatrix (C++ *type*), 107
- HYPRE\_ParCSRMatrixCreate (C++ *function*), 107
- HYPRE\_ParCSRMatrixDestroy (C++ *function*), 107
- HYPRE\_ParCSRMatrixInitialize (C++ *function*), 108
- HYPRE\_ParCSRMatrixMatmat (C++ *function*), 109
- HYPRE\_ParCSRMatrixMatvec (C++ *function*), 108
- HYPRE\_ParCSRMatrixMatvecT (C++ *function*), 108
- HYPRE\_ParCSRMatrixPrint (C++ *function*), 108
- HYPRE\_ParCSRMatrixRead (C++ *function*), 108
- HYPRE\_ParCSRMultiVectorPrint (C++ *function*), 197
- HYPRE\_ParCSRMultiVectorRead (C++ *function*), 197
- HYPRE\_ParCSROnProcTriSetup (C++ *function*), 172
- HYPRE\_ParCSROnProcTriSolve (C++ *function*), 172
- HYPRE\_ParCSRParaSailsCreate (C++ *function*), 161
- HYPRE\_ParCSRParaSailsDestroy (C++ *function*), 161
- HYPRE\_ParCSRParaSailsSetFilter (C++ *function*), 162
- HYPRE\_ParCSRParaSailsSetLoadbal (C++ *function*), 162
- HYPRE\_ParCSRParaSailsSetLogging (C++ *function*), 162
- HYPRE\_ParCSRParaSailsSetParams (C++ *function*), 161
- HYPRE\_ParCSRParaSailsSetReuse (C++ *function*), 162
- HYPRE\_ParCSRParaSailsSetSym (C++ *function*), 162
- HYPRE\_ParCSRParaSailsSetup (C++ *function*), 161
- HYPRE\_ParCSRParaSailsSolve (C++ *function*), 161
- HYPRE\_ParCSRPCGCreate (C++ *function*), 171
- HYPRE\_ParCSRPCGDestroy (C++ *function*), 171
- HYPRE\_ParCSRPCGGetFinalRelativeResidualNorm (C++ *function*), 172
- HYPRE\_ParCSRPCGGetNumIterations (C++ *function*), 172
- HYPRE\_ParCSRPCGGetPrecond (C++ *function*), 172
- HYPRE\_ParCSRPCGGetResidual (C++ *function*), 172
- HYPRE\_ParCSRPCGSetAbsoluteTol (C++ *function*), 171
- HYPRE\_ParCSRPCGSetLogging (C++ *function*), 172
- HYPRE\_ParCSRPCGSetMaxIter (C++ *function*), 171
- HYPRE\_ParCSRPCGSetPrecond (C++ *function*), 172
- HYPRE\_ParCSRPCGSetPreconditioner (C++ *function*), 172
- HYPRE\_ParCSRPCGSetPrintLevel (C++ *function*), 172
- HYPRE\_ParCSRPCGSetRelChange (C++ *function*), 172
- HYPRE\_ParCSRPCGSetStopCrit (C++ *function*), 171
- HYPRE\_ParCSRPCGSetTol (C++ *function*), 171
- HYPRE\_ParCSRPCGSetTwoNorm (C++ *function*), 172
- HYPRE\_ParCSRPCGSetup (C++ *function*), 171
- HYPRE\_ParCSRPCGSolve (C++ *function*), 171
- HYPRE\_ParCSRPilutCreate (C++ *function*), 164
- HYPRE\_ParCSRPilutDestroy (C++ *function*), 164
- HYPRE\_ParCSRPilutSetDropTolerance (C++ *function*), 164
- HYPRE\_ParCSRPilutSetFactorRowSize (C++ *function*), 164
- HYPRE\_ParCSRPilutSetLogging (C++ *function*), 164
- HYPRE\_ParCSRPilutSetMaxIter (C++ *function*), 164
- HYPRE\_ParCSRPilutSetup (C++ *function*), 164
- HYPRE\_ParCSRPilutSolve (C++ *function*), 164
- HYPRE\_ParCSRSetupInterpreter (C++ *function*), 197
- HYPRE\_ParCSRSetupMatvec (C++ *function*), 197
- HYPRE\_ParVector (C++ *type*), 108
- HYPRE\_ParVectorAxy (C++ *function*), 108
- HYPRE\_ParVectorCopy (C++ *function*), 108
- HYPRE\_ParVectorCreate (C++ *function*), 108
- HYPRE\_ParVectorDestroy (C++ *function*), 108
- HYPRE\_ParVectorInitialize (C++ *function*), 108
- HYPRE\_ParVectorInnerProd (C++ *function*), 108
- HYPRE\_ParVectorPrint (C++ *function*), 108
- HYPRE\_ParVectorRead (C++ *function*), 108
- HYPRE\_ParVectorScale (C++ *function*), 108
- HYPRE\_PCGGetAbsoluteTolFactor (C++ *function*), 202
- HYPRE\_PCGGetConverged (C++ *function*), 203
- HYPRE\_PCGGetConvergenceFactorTol (C++ *function*), 202
- HYPRE\_PCGGetFinalRelativeResidualNorm (C++ *function*), 202

- HYPRE\_PCGGetFlex (C++ function), 203  
 HYPRE\_PCGGetLogging (C++ function), 203  
 HYPRE\_PCGGetMaxIter (C++ function), 202  
 HYPRE\_PCGGetNumIterations (C++ function), 202  
 HYPRE\_PCGGetPrecond (C++ function), 203  
 HYPRE\_PCGGetPrecondMatrix (C++ function), 203  
 HYPRE\_PCGGetPrintLevel (C++ function), 203  
 HYPRE\_PCGGetRelChange (C++ function), 202  
 HYPRE\_PCGGetResidual (C++ function), 202  
 HYPRE\_PCGGetResidualTol (C++ function), 202  
 HYPRE\_PCGGetSkipBreak (C++ function), 202  
 HYPRE\_PCGGetStopCrit (C++ function), 202  
 HYPRE\_PCGGetTol (C++ function), 202  
 HYPRE\_PCGGetTwoNorm (C++ function), 202  
 HYPRE\_PCGSetAbsoluteTol (C++ function), 201  
 HYPRE\_PCGSetAbsoluteTolFactor (C++ function), 201  
 HYPRE\_PCGSetConvergenceFactorTol (C++ function), 201  
 HYPRE\_PCGSetFlex (C++ function), 202  
 HYPRE\_PCGSetLogging (C++ function), 202  
 HYPRE\_PCGSetMaxIter (C++ function), 201  
 HYPRE\_PCGSetPrecond (C++ function), 202  
 HYPRE\_PCGSetPreconditioner (C++ function), 202  
 HYPRE\_PCGSetPrecondMatrix (C++ function), 202  
 HYPRE\_PCGSetPrintLevel (C++ function), 202  
 HYPRE\_PCGSetRecomputeResidual (C++ function), 201  
 HYPRE\_PCGSetRecomputeResidualP (C++ function), 201  
 HYPRE\_PCGSetRelChange (C++ function), 201  
 HYPRE\_PCGSetResidualTol (C++ function), 201  
 HYPRE\_PCGSetSkipBreak (C++ function), 202  
 HYPRE\_PCGSetStopCrit (C++ function), 201  
 HYPRE\_PCGSetTol (C++ function), 201  
 HYPRE\_PCGSetTwoNorm (C++ function), 201  
 HYPRE\_PCGSetup (C++ function), 201  
 HYPRE\_PCGSolve (C++ function), 201  
 HYPRE\_PFMG (C macro), 126  
 HYPRE\_Precision (C++ enum), 212  
 HYPRE\_Precision::HYPRE\_REAL\_DOUBLE (C++ enumerator), 213  
 HYPRE\_Precision::HYPRE\_REAL\_LONGDOUBLE (C++ enumerator), 213  
 HYPRE\_Precision::HYPRE\_REAL\_SINGLE (C++ enumerator), 213  
 HYPRE\_PrintDeviceInfo (C++ function), 215  
 HYPRE\_PrintErrorMessages (C++ function), 214  
 HYPRE\_PtrToModifyPCFcn (C++ type), 201  
 HYPRE\_PtrToParSolverFcn (C++ type), 131  
 HYPRE\_PtrToSStructSolverFcn (C++ type), 121  
 HYPRE\_PtrToStructSolverFcn (C++ type), 109  
 HYPRE\_SchwarzCreate (C++ function), 198  
 HYPRE\_SchwarzDestroy (C++ function), 198  
 HYPRE\_SchwarzGetFinalResidualNorm (C++ function), 199  
 HYPRE\_SchwarzGetNumIterations (C++ function), 199  
 HYPRE\_SchwarzSetDofFunc (C++ function), 198  
 HYPRE\_SchwarzSetDomainStructure (C++ function), 198  
 HYPRE\_SchwarzSetDomainType (C++ function), 198  
 HYPRE\_SchwarzSetILUKLevelOfFill (C++ function), 199  
 HYPRE\_SchwarzSetILUTDroptol (C++ function), 199  
 HYPRE\_SchwarzSetILUTMaxNnzPerRow (C++ function), 199  
 HYPRE\_SchwarzSetLocalSolverType (C++ function), 199  
 HYPRE\_SchwarzSetLogging (C++ function), 199  
 HYPRE\_SchwarzSetMaxIter (C++ function), 199  
 HYPRE\_SchwarzSetNonSymm (C++ function), 199  
 HYPRE\_SchwarzSetNumFunctions (C++ function), 198  
 HYPRE\_SchwarzSetOverlap (C++ function), 198  
 HYPRE\_SchwarzSetPrintLevel (C++ function), 199  
 HYPRE\_SchwarzSetRelaxWeight (C++ function), 198  
 HYPRE\_SchwarzSetTol (C++ function), 199  
 HYPRE\_SchwarzSetup (C++ function), 198  
 HYPRE\_SchwarzSetVariant (C++ function), 198  
 HYPRE\_SchwarzSolve (C++ function), 198  
 HYPRE\_SetExecutionPolicy (C++ function), 214  
 HYPRE\_SetGlobalPrecision (C++ function), 213  
 HYPRE\_SetGpuAwareMPI (C++ function), 219  
 HYPRE\_SetLogLevel (C++ function), 217  
 HYPRE\_SetMemoryLocation (C++ function), 213  
 HYPRE\_SetPrintErrorMode (C++ function), 214  
 HYPRE\_SetPrintErrorVerbosity (C++ function), 214  
 HYPRE\_SetSpGemmUseCuspars (C macro), 220  
 HYPRE\_SetSpGemmUseVendor (C++ function), 218  
 HYPRE\_SetSpMVUseVendor (C++ function), 218  
 HYPRE\_SetSpTransUseVendor (C++ function), 218  
 HYPRE\_SetUmpireDevicePoolName (C++ function), 217  
 HYPRE\_SetUmpireDevicePoolSize (C++ function), 216  
 HYPRE\_SetUmpireHostPoolName (C++ function), 217  
 HYPRE\_SetUmpireHostPoolSize (C++ function), 216  
 HYPRE\_SetUmpirePinnedPoolName (C++ function), 217  
 HYPRE\_SetUmpirePinnedPoolSize (C++ function), 217  
 HYPRE\_SetUmpireUMPoolName (C++ function), 217  
 HYPRE\_SetUmpireUMPoolSize (C++ function), 216  
 HYPRE\_SetUseGpuRand (C++ function), 219  
 HYPRE\_SMG (C macro), 126  
 HYPRE\_SSTRUCT\_VARIABLE\_CELL (C macro), 88  
 HYPRE\_SSTRUCT\_VARIABLE\_NODE (C macro), 88  
 HYPRE\_SSTRUCT\_VARIABLE\_UNDEFINED (C macro), 88

- HYPRE\_SSTRUCT\_VARIABLE\_XEDGE (*C macro*), 88  
HYPRE\_SSTRUCT\_VARIABLE\_XFACE (*C macro*), 88  
HYPRE\_SSTRUCT\_VARIABLE\_YEDGE (*C macro*), 88  
HYPRE\_SSTRUCT\_VARIABLE\_YFACE (*C macro*), 88  
HYPRE\_SSTRUCT\_VARIABLE\_ZEDGE (*C macro*), 88  
HYPRE\_SSTRUCT\_VARIABLE\_ZFACE (*C macro*), 88  
HYPRE\_SStructBiCGSTABCreate (*C++ function*), 130  
HYPRE\_SStructBiCGSTABDestroy (*C++ function*), 130  
HYPRE\_SStructBiCGSTABGetFinalRelativeResidualNorm (*C++ function*), 131  
HYPRE\_SStructBiCGSTABGetNumIterations (*C++ function*), 131  
HYPRE\_SStructBiCGSTABGetResidual (*C++ function*), 131  
HYPRE\_SStructBiCGSTABSetAbsoluteTol (*C++ function*), 130  
HYPRE\_SStructBiCGSTABSetLogging (*C++ function*), 131  
HYPRE\_SStructBiCGSTABSetMaxIter (*C++ function*), 131  
HYPRE\_SStructBiCGSTABSetMinIter (*C++ function*), 131  
HYPRE\_SStructBiCGSTABSetPrecond (*C++ function*), 131  
HYPRE\_SStructBiCGSTABSetPrintLevel (*C++ function*), 131  
HYPRE\_SStructBiCGSTABSetStopCrit (*C++ function*), 131  
HYPRE\_SStructBiCGSTABSetTol (*C++ function*), 130  
HYPRE\_SStructBiCGSTABSetup (*C++ function*), 130  
HYPRE\_SStructBiCGSTABSolve (*C++ function*), 130  
HYPRE\_SStructDiagScale (*C++ function*), 127  
HYPRE\_SStructDiagScaleSetup (*C++ function*), 127  
HYPRE\_SStructFlexGMRESCreate (*C++ function*), 128  
HYPRE\_SStructFlexGMRESDestroy (*C++ function*), 128  
HYPRE\_SStructFlexGMRESGetFinalRelativeResidualNorm (*C++ function*), 129  
HYPRE\_SStructFlexGMRESGetNumIterations (*C++ function*), 129  
HYPRE\_SStructFlexGMRESGetResidual (*C++ function*), 129  
HYPRE\_SStructFlexGMRESSetAbsoluteTol (*C++ function*), 129  
HYPRE\_SStructFlexGMRESSetKDim (*C++ function*), 129  
HYPRE\_SStructFlexGMRESSetLogging (*C++ function*), 129  
HYPRE\_SStructFlexGMRESSetMaxIter (*C++ function*), 129  
HYPRE\_SStructFlexGMRESSetMinIter (*C++ function*), 129  
HYPRE\_SStructFlexGMRESSetModifyPC (*C++ function*), 129  
HYPRE\_SStructFlexGMRESSetPrecond (*C++ function*), 129  
HYPRE\_SStructFlexGMRESSetPrintLevel (*C++ function*), 129  
HYPRE\_SStructFlexGMRESSetStopCrit (*C++ function*), 128  
HYPRE\_SStructFlexGMRESSetTol (*C++ function*), 128  
HYPRE\_SStructFlexGMRESSetup (*C++ function*), 128  
HYPRE\_SStructFlexGMRESolve (*C++ function*), 128  
HYPRE\_SStructGraph (*C++ type*), 89  
HYPRE\_SStructGraphAddEntries (*C++ function*), 89  
HYPRE\_SStructGraphAssemble (*C++ function*), 89  
HYPRE\_SStructGraphCreate (*C++ function*), 89  
HYPRE\_SStructGraphDestroy (*C++ function*), 89  
HYPRE\_SStructGraphPrint (*C++ function*), 90  
HYPRE\_SStructGraphRead (*C++ function*), 90  
HYPRE\_SStructGraphSetDomainGrid (*C++ function*), 89  
HYPRE\_SStructGraphSetFEM (*C++ function*), 89  
HYPRE\_SStructGraphSetFEMSparsity (*C++ function*), 89  
HYPRE\_SStructGraphSetObjectType (*C++ function*), 89  
HYPRE\_SStructGraphSetStencil (*C++ function*), 89  
HYPRE\_SStructGrid (*C++ type*), 82  
HYPRE\_SStructGridAddUnstructuredPart (*C++ function*), 87  
HYPRE\_SStructGridAddVariables (*C++ function*), 83  
HYPRE\_SStructGridAssemble (*C++ function*), 87  
HYPRE\_SStructGridCoarsen (*C++ function*), 88  
HYPRE\_SStructGridCreate (*C++ function*), 83  
HYPRE\_SStructGridDestroy (*C++ function*), 83  
HYPRE\_SStructGridGetGradient (*C++ function*), 87

- HYPRE\_SStructGridGetVariableBox (C++ function), 87
- HYPRE\_SStructGridPrintGLVis (C++ function), 88
- HYPRE\_SStructGridProjectBox (C++ function), 88
- HYPRE\_SStructGridSetAMRInterp (C++ function), 86
- HYPRE\_SStructGridSetAMRPart (C++ function), 85
- HYPRE\_SStructGridSetAMRRefInterp (C++ function), 86
- HYPRE\_SStructGridSetAMRRefRestrictT (C++ function), 86
- HYPRE\_SStructGridSetAMRRefSlaves (C++ function), 85
- HYPRE\_SStructGridSetAMRRestrictT (C++ function), 86
- HYPRE\_SStructGridSetExtents (C++ function), 83
- HYPRE\_SStructGridSetFEMOrdering (C++ function), 83
- HYPRE\_SStructGridSetNeighborPart (C++ function), 84
- HYPRE\_SStructGridSetNumGhost (C++ function), 87
- HYPRE\_SStructGridSetObjectType (C++ function), 87
- HYPRE\_SStructGridSetPeriodic (C++ function), 87
- HYPRE\_SStructGridSetSharedPart (C++ function), 84
- HYPRE\_SStructGridSetVariables (C++ function), 83
- HYPRE\_SStructLGMRESCreate (C++ function), 129
- HYPRE\_SStructLGMRESDestroy (C++ function), 129
- HYPRE\_SStructLGMRESGetFinalRelativeResidualNorm (C++ function), 130
- HYPRE\_SStructLGMRESGetNumIterations (C++ function), 130
- HYPRE\_SStructLGMRESGetResidual (C++ function), 130
- HYPRE\_SStructLGMRESSetAbsoluteTol (C++ function), 130
- HYPRE\_SStructLGMRESSetAugDim (C++ function), 130
- HYPRE\_SStructLGMRESSetKDim (C++ function), 130
- HYPRE\_SStructLGMRESSetLogging (C++ function), 130
- HYPRE\_SStructLGMRESSetMaxIter (C++ function), 130
- HYPRE\_SStructLGMRESSetMinIter (C++ function), 130
- HYPRE\_SStructLGMRESSetPrecond (C++ function), 130
- HYPRE\_SStructLGMRESSetPrintLevel (C++ function), 130
- HYPRE\_SStructLGMRESSetTol (C++ function), 130
- HYPRE\_SStructLGMRESSetup (C++ function), 129
- HYPRE\_SStructLGMRESSolve (C++ function), 130
- HYPRE\_SStructMatrix (C++ type), 90
- HYPRE\_SStructMatrixAddFEMBoxValues (C++ function), 93
- HYPRE\_SStructMatrixAddFEMValues (C++ function), 91
- HYPRE\_SStructMatrixAddToBoxValues (C++ function), 92
- HYPRE\_SStructMatrixAddToBoxValues2 (C++ function), 93
- HYPRE\_SStructMatrixAddToValues (C++ function), 91
- HYPRE\_SStructMatrixAssemble (C++ function), 93
- HYPRE\_SStructMatrixCreate (C++ function), 90
- HYPRE\_SStructMatrixDestroy (C++ function), 90
- HYPRE\_SStructMatrixGetBoxValues (C++ function), 93
- HYPRE\_SStructMatrixGetBoxValues2 (C++ function), 93
- HYPRE\_SStructMatrixGetFEMBoxValues (C++ function), 93
- HYPRE\_SStructMatrixGetFEMValues (C++ function), 92
- HYPRE\_SStructMatrixGetGrid (C++ function), 94
- HYPRE\_SStructMatrixGetObject (C++ function), 94
- HYPRE\_SStructMatrixGetValues (C++ function), 91
- HYPRE\_SStructMatrixInitialize (C++ function), 90
- HYPRE\_SStructMatrixMatmat (C++ function), 98
- HYPRE\_SStructMatrixMatvec (C++ function), 98
- HYPRE\_SStructMatrixPrint (C++ function), 94
- HYPRE\_SStructMatrixRead (C++ function), 94
- HYPRE\_SStructMatrixScale (C++ function), 98
- HYPRE\_SStructMatrixSetBoxValues (C++ function), 92
- HYPRE\_SStructMatrixSetBoxValues2 (C++ function), 92
- HYPRE\_SStructMatrixSetConstantEntries (C++ function), 90
- HYPRE\_SStructMatrixSetDomainStride (C++ function), 90
- HYPRE\_SStructMatrixSetEarlyAssemble (C++ function), 90
- HYPRE\_SStructMatrixSetNSSymmetric (C++ function), 94
- HYPRE\_SStructMatrixSetObjectType (C++ function), 94
- HYPRE\_SStructMatrixSetRangeStride (C++ function), 90
- HYPRE\_SStructMatrixSetSymmetric (C++ function), 93
- HYPRE\_SStructMatrixSetValues (C++ function), 91
- HYPRE\_SStructMatrixToIJMatrix (C++ function), 94
- HYPRE\_SStructPCGCreate (C++ function), 127
- HYPRE\_SStructPCGDestroy (C++ function), 127
- HYPRE\_SStructPCGGetFinalRelativeResidualNorm (C++ function), 127
- HYPRE\_SStructPCGGetNumIterations (C++ function), 127

- HYPRE\_SStructPCGGetResidual (C++ function), 127  
 HYPRE\_SStructPCGSetAbsoluteTol (C++ function), 127  
 HYPRE\_SStructPCGSetLogging (C++ function), 127  
 HYPRE\_SStructPCGSetMaxIter (C++ function), 127  
 HYPRE\_SStructPCGSetPrecond (C++ function), 127  
 HYPRE\_SStructPCGSetPrintLevel (C++ function), 127  
 HYPRE\_SStructPCGSetRelChange (C++ function), 127  
 HYPRE\_SStructPCGSetTol (C++ function), 127  
 HYPRE\_SStructPCGSetTwoNorm (C++ function), 127  
 HYPRE\_SStructPCGSetup (C++ function), 127  
 HYPRE\_SStructPCGSolve (C++ function), 127  
 HYPRE\_SStructSetupInterpreter (C++ function), 131  
 HYPRE\_SStructSetupMatvec (C++ function), 131  
 HYPRE\_SStructSolver (C++ type), 121  
 HYPRE\_SStructSplitCreate (C++ function), 125  
 HYPRE\_SStructSplitDestroy (C++ function), 125  
 HYPRE\_SStructSplitGetFinalRelativeResidualNorm (C++ function), 126  
 HYPRE\_SStructSplitGetNumIterations (C++ function), 126  
 HYPRE\_SStructSplitPrintLogging (C++ function), 126  
 HYPRE\_SStructSplitSetLogging (C++ function), 126  
 HYPRE\_SStructSplitSetMaxIter (C++ function), 126  
 HYPRE\_SStructSplitSetNonZeroGuess (C++ function), 126  
 HYPRE\_SStructSplitSetPrintLevel (C++ function), 126  
 HYPRE\_SStructSplitSetStructSolver (C++ function), 126  
 HYPRE\_SStructSplitSetTol (C++ function), 126  
 HYPRE\_SStructSplitSetup (C++ function), 125  
 HYPRE\_SStructSplitSetZeroGuess (C++ function), 126  
 HYPRE\_SStructSplitSolve (C++ function), 125  
 HYPRE\_SStructSSAMGCreate (C++ function), 123  
 HYPRE\_SStructSSAMGDestroy (C++ function), 123  
 HYPRE\_SStructSSAMGGetFinalRelativeResidualNorm (C++ function), 125  
 HYPRE\_SStructSSAMGGetNumIterations (C++ function), 125  
 HYPRE\_SStructSSAMGSetCoarseSolverType (C++ function), 124  
 HYPRE\_SStructSSAMGSetDxyz (C++ function), 125  
 HYPRE\_SStructSSAMGSetInterpType (C++ function), 124  
 HYPRE\_SStructSSAMGSetLogging (C++ function), 125  
 HYPRE\_SStructSSAMGSetMaxCoarseSize (C++ function), 124  
 HYPRE\_SStructSSAMGSetMaxIter (C++ function), 123  
 HYPRE\_SStructSSAMGSetMaxLevels (C++ function), 123  
 HYPRE\_SStructSSAMGSetNonGalerkinRAP (C++ function), 123  
 HYPRE\_SStructSSAMGSetNonZeroGuess (C++ function), 124  
 HYPRE\_SStructSSAMGSetNumCoarseRelax (C++ function), 124  
 HYPRE\_SStructSSAMGSetNumPostRelax (C++ function), 124  
 HYPRE\_SStructSSAMGSetNumPreRelax (C++ function), 124  
 HYPRE\_SStructSSAMGSetPrintFreq (C++ function), 125  
 HYPRE\_SStructSSAMGSetPrintLevel (C++ function), 125  
 HYPRE\_SStructSSAMGSetRelaxType (C++ function), 124  
 HYPRE\_SStructSSAMGSetRelaxWeight (C++ function), 124  
 HYPRE\_SStructSSAMGSetRelChange (C++ function), 123  
 HYPRE\_SStructSSAMGSetSkipRelax (C++ function), 124  
 HYPRE\_SStructSSAMGSetTol (C++ function), 123  
 HYPRE\_SStructSSAMGSetup (C++ function), 123  
 HYPRE\_SStructSSAMGSetZeroGuess (C++ function), 123  
 HYPRE\_SStructSSAMGSolve (C++ function), 123  
 HYPRE\_SStructStencil (C++ type), 88  
 HYPRE\_SStructStencilCreate (C++ function), 88  
 HYPRE\_SStructStencilDestroy (C++ function), 88  
 HYPRE\_SStructStencilPrint (C++ function), 88  
 HYPRE\_SStructStencilRead (C++ function), 89  
 HYPRE\_SStructStencilSetEntry (C++ function), 88  
 HYPRE\_SStructSysPFMGCreate (C++ function), 121  
 HYPRE\_SStructSysPFMGDestroy (C++ function), 121  
 HYPRE\_SStructSysPFMGGetFinalRelativeResidualNorm (C++ function), 122  
 HYPRE\_SStructSysPFMGGetNumIterations (C++ function), 122  
 HYPRE\_SStructSysPFMGSetDxyz (C++ function), 122  
 HYPRE\_SStructSysPFMGSetJacobiWeight (C++ function), 122  
 HYPRE\_SStructSysPFMGSetLogging (C++ function), 122  
 HYPRE\_SStructSysPFMGSetMaxIter (C++ function), 121  
 HYPRE\_SStructSysPFMGSetNonZeroGuess (C++ function), 122  
 HYPRE\_SStructSysPFMGSetNumPostRelax (C++ function), 122  
 HYPRE\_SStructSysPFMGSetNumPreRelax (C++ function), 122  
 HYPRE\_SStructSysPFMGSetPrintLevel (C++ function), 122

- tion), 122
- HYPRE\_SStructSysPFMGSetRelaxType (C++ function), 122
- HYPRE\_SStructSysPFMGSetRelChange (C++ function), 121
- HYPRE\_SStructSysPFMGSetSkipRelax (C++ function), 122
- HYPRE\_SStructSysPFMGSetTol (C++ function), 121
- HYPRE\_SStructSysPFMGSetup (C++ function), 121
- HYPRE\_SStructSysPFMGSetZeroGuess (C++ function), 121
- HYPRE\_SStructSysPFMGsolve (C++ function), 121
- HYPRE\_SStructVariable (C++ type), 82
- HYPRE\_SStructVector (C++ type), 95
- HYPRE\_SStructVectorAddFEMBoxValues (C++ function), 96
- HYPRE\_SStructVectorAddFEMValues (C++ function), 95
- HYPRE\_SStructVectorAddToBoxValues (C++ function), 96
- HYPRE\_SStructVectorAddToBoxValues2 (C++ function), 96
- HYPRE\_SStructVectorAddToValues (C++ function), 95
- HYPRE\_SStructVectorAssemble (C++ function), 97
- HYPRE\_SStructVectorAxy (C++ function), 98
- HYPRE\_SStructVectorCopy (C++ function), 98
- HYPRE\_SStructVectorCreate (C++ function), 95
- HYPRE\_SStructVectorDestroy (C++ function), 95
- HYPRE\_SStructVectorGather (C++ function), 97
- HYPRE\_SStructVectorGetBoxValues (C++ function), 97
- HYPRE\_SStructVectorGetBoxValues2 (C++ function), 97
- HYPRE\_SStructVectorGetFEMBoxValues (C++ function), 97
- HYPRE\_SStructVectorGetFEMValues (C++ function), 96
- HYPRE\_SStructVectorGetObject (C++ function), 97
- HYPRE\_SStructVectorGetValues (C++ function), 95
- HYPRE\_SStructVectorInitialize (C++ function), 95
- HYPRE\_SStructVectorInnerProd (C++ function), 98
- HYPRE\_SStructVectorPrint (C++ function), 98
- HYPRE\_SStructVectorPrintGLVis (C++ function), 98
- HYPRE\_SStructVectorRead (C++ function), 98
- HYPRE\_SStructVectorScale (C++ function), 98
- HYPRE\_SStructVectorSetBoxValues (C++ function), 96
- HYPRE\_SStructVectorSetBoxValues2 (C++ function), 96
- HYPRE\_SStructVectorSetConstantValues (C++ function), 95
- HYPRE\_SStructVectorSetObjectType (C++ function), 97
- HYPRE\_SStructVectorSetRandomValues (C++ function), 95
- HYPRE\_SStructVectorSetValues (C++ function), 95
- HYPRE\_StructBiCGSTABCreate (C++ function), 118
- HYPRE\_StructBiCGSTABDestroy (C++ function), 118
- HYPRE\_StructBiCGSTABGetFinalRelativeResidualNorm (C++ function), 118
- HYPRE\_StructBiCGSTABGetNumIterations (C++ function), 118
- HYPRE\_StructBiCGSTABGetResidual (C++ function), 118
- HYPRE\_StructBiCGSTABSetAbsoluteTol (C++ function), 118
- HYPRE\_StructBiCGSTABSetLogging (C++ function), 118
- HYPRE\_StructBiCGSTABSetMaxIter (C++ function), 118
- HYPRE\_StructBiCGSTABSetPrecond (C++ function), 118
- HYPRE\_StructBiCGSTABSetPrintLevel (C++ function), 118
- HYPRE\_StructBiCGSTABSetTol (C++ function), 118
- HYPRE\_StructBiCGSTABSetup (C++ function), 118
- HYPRE\_StructBiCGSTABsolve (C++ function), 118
- HYPRE\_StructCycRedCreate (C++ function), 114
- HYPRE\_StructCycRedDestroy (C++ function), 114
- HYPRE\_StructCycRedSetBase (C++ function), 115
- HYPRE\_StructCycRedSetTDim (C++ function), 114
- HYPRE\_StructCycRedSetup (C++ function), 114
- HYPRE\_StructCycRedSolve (C++ function), 114
- HYPRE\_StructDiagScale (C++ function), 115
- HYPRE\_StructDiagScaleSetup (C++ function), 115
- HYPRE\_StructFlexGMRESCreate (C++ function), 116
- HYPRE\_StructFlexGMRESDestroy (C++ function), 116
- HYPRE\_StructFlexGMRESGetFinalRelativeResidualNorm (C++ function), 117
- HYPRE\_StructFlexGMRESGetNumIterations (C++ function), 117
- HYPRE\_StructFlexGMRESGetResidual (C++ function), 117
- HYPRE\_StructFlexGMRESSetAbsoluteTol (C++ function), 117
- HYPRE\_StructFlexGMRESSetKDim (C++ function), 117
- HYPRE\_StructFlexGMRESSetLogging (C++ function), 117
- HYPRE\_StructFlexGMRESSetMaxIter (C++ function), 117
- HYPRE\_StructFlexGMRESSetModifyPC (C++ function), 117
- HYPRE\_StructFlexGMRESSetPrecond (C++ function), 117
- HYPRE\_StructFlexGMRESSetPrintLevel (C++ function), 117
- HYPRE\_StructFlexGMRESSetTol (C++ function), 117

- HYPRE\_StructFlexGMRESSetup (C++ function), 116  
 HYPRE\_StructFlexGMRESSolve (C++ function), 116  
 HYPRE\_StructGMRESCreate (C++ function), 116  
 HYPRE\_StructGMRESDestroy (C++ function), 116  
 HYPRE\_StructGMRESGetFinalRelativeResidualNorm (C++ function), 116  
 HYPRE\_StructGMRESGetNumIterations (C++ function), 116  
 HYPRE\_StructGMRESGetResidual (C++ function), 116  
 HYPRE\_StructGMRESSetAbsoluteTol (C++ function), 116  
 HYPRE\_StructGMRESSetKDim (C++ function), 116  
 HYPRE\_StructGMRESSetLogging (C++ function), 116  
 HYPRE\_StructGMRESSetMaxIter (C++ function), 116  
 HYPRE\_StructGMRESSetPrecond (C++ function), 116  
 HYPRE\_StructGMRESSetPrintLevel (C++ function), 116  
 HYPRE\_StructGMRESSetTol (C++ function), 116  
 HYPRE\_StructGMRESSetup (C++ function), 116  
 HYPRE\_StructGMRESSolve (C++ function), 116  
 HYPRE\_StructGrid (C++ type), 75  
 HYPRE\_StructGridAssemble (C++ function), 75  
 HYPRE\_StructGridCoarsen (C++ function), 76  
 HYPRE\_StructGridCreate (C++ function), 75  
 HYPRE\_StructGridDestroy (C++ function), 75  
 HYPRE\_StructGridPrintVTK (C++ function), 75  
 HYPRE\_StructGridProjectBox (C++ function), 76  
 HYPRE\_StructGridSetExtents (C++ function), 75  
 HYPRE\_StructGridSetNumGhost (C++ function), 75  
 HYPRE\_StructGridSetPeriodic (C++ function), 75  
 HYPRE\_StructHybridCreate (C++ function), 119  
 HYPRE\_StructHybridDestroy (C++ function), 119  
 HYPRE\_StructHybridGetDSCGNumIterations (C++ function), 120  
 HYPRE\_StructHybridGetFinalRelativeResidualNorm (C++ function), 120  
 HYPRE\_StructHybridGetNumIterations (C++ function), 120  
 HYPRE\_StructHybridGetPCGNumIterations (C++ function), 120  
 HYPRE\_StructHybridGetRecomputeResidual (C++ function), 119  
 HYPRE\_StructHybridGetRecomputeResidualP (C++ function), 120  
 HYPRE\_StructHybridSetConvergenceTol (C++ function), 119  
 HYPRE\_StructHybridSetDSCGMaxIter (C++ function), 119  
 HYPRE\_StructHybridSetKDim (C++ function), 120  
 HYPRE\_StructHybridSetLogging (C++ function), 120  
 HYPRE\_StructHybridSetPCGAbsoluteTolFactor (C++ function), 120  
 HYPRE\_StructHybridSetPCGMaxIter (C++ function), 119  
 HYPRE\_StructHybridSetPrecond (C++ function), 120  
 HYPRE\_StructHybridSetPrintLevel (C++ function), 120  
 HYPRE\_StructHybridSetRecomputeResidual (C++ function), 119  
 HYPRE\_StructHybridSetRecomputeResidualP (C++ function), 120  
 HYPRE\_StructHybridSetRelChange (C++ function), 119  
 HYPRE\_StructHybridSetSolverType (C++ function), 119  
 HYPRE\_StructHybridSetStopCrit (C++ function), 119  
 HYPRE\_StructHybridSetTol (C++ function), 119  
 HYPRE\_StructHybridSetTwoNorm (C++ function), 119  
 HYPRE\_StructHybridSetup (C++ function), 119  
 HYPRE\_StructHybridSolve (C++ function), 119  
 HYPRE\_StructJacobiCreate (C++ function), 109  
 HYPRE\_StructJacobiDestroy (C++ function), 109  
 HYPRE\_StructJacobiGetFinalRelativeResidualNorm (C++ function), 110  
 HYPRE\_StructJacobiGetMaxIter (C++ function), 109  
 HYPRE\_StructJacobiGetNumIterations (C++ function), 110  
 HYPRE\_StructJacobiGetTol (C++ function), 109  
 HYPRE\_StructJacobiGetZeroGuess (C++ function), 110  
 HYPRE\_StructJacobiSetMaxIter (C++ function), 109  
 HYPRE\_StructJacobiSetNonZeroGuess (C++ function), 110  
 HYPRE\_StructJacobiSetTol (C++ function), 109  
 HYPRE\_StructJacobiSetup (C++ function), 109  
 HYPRE\_StructJacobiSetZeroGuess (C++ function), 110  
 HYPRE\_StructJacobiSolve (C++ function), 109  
 HYPRE\_StructLGMRESCreate (C++ function), 117  
 HYPRE\_StructLGMRESDestroy (C++ function), 117  
 HYPRE\_StructLGMRESGetFinalRelativeResidualNorm (C++ function), 118  
 HYPRE\_StructLGMRESGetNumIterations (C++ function), 118  
 HYPRE\_StructLGMRESGetResidual (C++ function), 118  
 HYPRE\_StructLGMRESSetAbsoluteTol (C++ function), 117  
 HYPRE\_StructLGMRESSetAugDim (C++ function), 117  
 HYPRE\_StructLGMRESSetKDim (C++ function), 117  
 HYPRE\_StructLGMRESSetLogging (C++ function), 118  
 HYPRE\_StructLGMRESSetMaxIter (C++ function), 117  
 HYPRE\_StructLGMRESSetPrecond (C++ function), 117  
 HYPRE\_StructLGMRESSetPrintLevel (C++ function), 118  
 HYPRE\_StructLGMRESSetTol (C++ function), 117  
 HYPRE\_StructLGMRESSetup (C++ function), 117

- HYPRE\_StructLGMRESSolve (C++ function), 117  
 HYPRE\_StructMatrix (C++ type), 76  
 HYPRE\_StructMatrixAddToBoxValues (C++ function), 78  
 HYPRE\_StructMatrixAddToBoxValues2 (C++ function), 78  
 HYPRE\_StructMatrixAddToConstantValues (C++ function), 77  
 HYPRE\_StructMatrixAddToValues (C++ function), 77  
 HYPRE\_StructMatrixAssemble (C++ function), 78  
 HYPRE\_StructMatrixCreate (C++ function), 76  
 HYPRE\_StructMatrixDestroy (C++ function), 76  
 HYPRE\_StructMatrixGetBoxValues (C++ function), 78  
 HYPRE\_StructMatrixGetBoxValues2 (C++ function), 78  
 HYPRE\_StructMatrixGetValues (C++ function), 78  
 HYPRE\_StructMatrixInitialize (C++ function), 77  
 HYPRE\_StructMatrixMatmat (C++ function), 82  
 HYPRE\_StructMatrixMatvec (C++ function), 81  
 HYPRE\_StructMatrixMatvecT (C++ function), 81  
 HYPRE\_StructMatrixPrint (C++ function), 79  
 HYPRE\_StructMatrixRead (C++ function), 79  
 HYPRE\_StructMatrixSetBoxValues (C++ function), 77  
 HYPRE\_StructMatrixSetBoxValues2 (C++ function), 78  
 HYPRE\_StructMatrixSetConstantEntries (C++ function), 79  
 HYPRE\_StructMatrixSetConstantValues (C++ function), 77  
 HYPRE\_StructMatrixSetDomainStride (C++ function), 77  
 HYPRE\_StructMatrixSetNumGhost (C++ function), 79  
 HYPRE\_StructMatrixSetRangeStride (C++ function), 76  
 HYPRE\_StructMatrixSetSymmetric (C++ function), 78  
 HYPRE\_StructMatrixSetTranspose (C++ function), 79  
 HYPRE\_StructMatrixSetValues (C++ function), 77  
 HYPRE\_StructPCGCreate (C++ function), 115  
 HYPRE\_StructPCGDestroy (C++ function), 115  
 HYPRE\_StructPCGGetFinalRelativeResidualNorm (C++ function), 115  
 HYPRE\_StructPCGGetNumIterations (C++ function), 115  
 HYPRE\_StructPCGGetResidual (C++ function), 115  
 HYPRE\_StructPCGSetAbsoluteTol (C++ function), 115  
 HYPRE\_StructPCGSetLogging (C++ function), 115  
 HYPRE\_StructPCGSetMaxIter (C++ function), 115  
 HYPRE\_StructPCGSetPrecond (C++ function), 115  
 HYPRE\_StructPCGSetPrintLevel (C++ function), 115  
 HYPRE\_StructPCGSetRelChange (C++ function), 115  
 HYPRE\_StructPCGSetTol (C++ function), 115  
 HYPRE\_StructPCGSetTwoNorm (C++ function), 115  
 HYPRE\_StructPCGSetup (C++ function), 115  
 HYPRE\_StructPCGSolve (C++ function), 115  
 HYPRE\_StructPFMGCreate (C++ function), 110  
 HYPRE\_StructPFMGDestroy (C++ function), 110  
 HYPRE\_StructPFMGGetFinalRelativeResidualNorm (C++ function), 112  
 HYPRE\_StructPFMGGetJacobiWeight (C++ function), 111  
 HYPRE\_StructPFMGGetLogging (C++ function), 112  
 HYPRE\_StructPFMGGetMatmultType (C++ function), 112  
 HYPRE\_StructPFMGGetMaxIter (C++ function), 110  
 HYPRE\_StructPFMGGetMaxLevels (C++ function), 111  
 HYPRE\_StructPFMGGetNumIterations (C++ function), 112  
 HYPRE\_StructPFMGGetNumPostRelax (C++ function), 112  
 HYPRE\_StructPFMGGetNumPreRelax (C++ function), 112  
 HYPRE\_StructPFMGGetPrintLevel (C++ function), 112  
 HYPRE\_StructPFMGGetRAPType (C++ function), 111  
 HYPRE\_StructPFMGGetRelaxType (C++ function), 111  
 HYPRE\_StructPFMGGetRelChange (C++ function), 111  
 HYPRE\_StructPFMGGetSkipRelax (C++ function), 112  
 HYPRE\_StructPFMGGetTol (C++ function), 110  
 HYPRE\_StructPFMGGetZeroGuess (C++ function), 111  
 HYPRE\_StructPFMGSetDxyz (C++ function), 112  
 HYPRE\_StructPFMGSetJacobiWeight (C++ function), 111  
 HYPRE\_StructPFMGSetLogging (C++ function), 112  
 HYPRE\_StructPFMGSetMatmultType (C++ function), 111  
 HYPRE\_StructPFMGSetMaxIter (C++ function), 110  
 HYPRE\_StructPFMGSetMaxLevels (C++ function), 110  
 HYPRE\_StructPFMGSetNonZeroGuess (C++ function), 111  
 HYPRE\_StructPFMGSetNumPostRelax (C++ function), 112  
 HYPRE\_StructPFMGSetNumPreRelax (C++ function), 112  
 HYPRE\_StructPFMGSetPrintLevel (C++ function), 112  
 HYPRE\_StructPFMGSetRAPType (C++ function), 111  
 HYPRE\_StructPFMGSetRelaxType (C++ function), 111  
 HYPRE\_StructPFMGSetRelChange (C++ function), 111  
 HYPRE\_StructPFMGSetSkipRelax (C++ function), 112  
 HYPRE\_StructPFMGSetTol (C++ function), 110  
 HYPRE\_StructPFMGSetup (C++ function), 110  
 HYPRE\_StructPFMGSetZeroGuess (C++ function), 111  
 HYPRE\_StructPFMGSolve (C++ function), 110

HYPRE\_StructSetupInterpreter (C++ function), 120  
 HYPRE\_StructSetupMatvec (C++ function), 120  
 HYPRE\_StructSMGCreate (C++ function), 113  
 HYPRE\_StructSMGDestroy (C++ function), 113  
 HYPRE\_StructSMGGetFinalRelativeResidualNorm  
 (C++ function), 114  
 HYPRE\_StructSMGGetLogging (C++ function), 114  
 HYPRE\_StructSMGGetMaxIter (C++ function), 113  
 HYPRE\_StructSMGGetMemoryUse (C++ function), 113  
 HYPRE\_StructSMGGetNumIterations (C++ function),  
 114  
 HYPRE\_StructSMGGetNumPostRelax (C++ function),  
 114  
 HYPRE\_StructSMGGetNumPreRelax (C++ function),  
 114  
 HYPRE\_StructSMGGetPrintLevel (C++ function), 114  
 HYPRE\_StructSMGGetRelChange (C++ function), 113  
 HYPRE\_StructSMGGetTol (C++ function), 113  
 HYPRE\_StructSMGGetZeroGuess (C++ function), 113  
 HYPRE\_StructSMGSetLogging (C++ function), 114  
 HYPRE\_StructSMGSetMaxIter (C++ function), 113  
 HYPRE\_StructSMGSetMemoryUse (C++ function), 113  
 HYPRE\_StructSMGSetNonZeroGuess (C++ function),  
 113  
 HYPRE\_StructSMGSetNumPostRelax (C++ function),  
 114  
 HYPRE\_StructSMGSetNumPreRelax (C++ function),  
 113  
 HYPRE\_StructSMGSetPrintLevel (C++ function), 114  
 HYPRE\_StructSMGSetRelChange (C++ function), 113  
 HYPRE\_StructSMGSetTol (C++ function), 113  
 HYPRE\_StructSMGSetup (C++ function), 113  
 HYPRE\_StructSMGSetZeroGuess (C++ function), 113  
 HYPRE\_StructSMGSolve (C++ function), 113  
 HYPRE\_StructSolver (C++ type), 109  
 HYPRE\_StructStencil (C++ type), 76  
 HYPRE\_StructStencilCreate (C++ function), 76  
 HYPRE\_StructStencilDestroy (C++ function), 76  
 HYPRE\_StructStencilSetElement (C++ function), 76  
 HYPRE\_StructStencilSetEntry (C++ function), 76  
 HYPRE\_StructVectorAddToBoxValues (C++ func-  
 tion), 80  
 HYPRE\_StructVectorAddToBoxValues2 (C++ func-  
 tion), 80  
 HYPRE\_StructVectorAddToValues (C++ function), 80  
 HYPRE\_StructVectorAssemble (C++ function), 80  
 HYPRE\_StructVectorAxy (C++ function), 81  
 HYPRE\_StructVectorClone (C++ function), 81  
 HYPRE\_StructVectorCopy (C++ function), 81  
 HYPRE\_StructVectorCreate (C++ function), 79  
 HYPRE\_StructVectorDestroy (C++ function), 79  
 HYPRE\_StructVectorGetBoxValues (C++ function),  
 81  
 HYPRE\_StructVectorGetBoxValues2 (C++ function),  
 81  
 HYPRE\_StructVectorGetValues (C++ function), 80  
 HYPRE\_StructVectorInitialize (C++ function), 79  
 HYPRE\_StructVectorInnerProd (C++ function), 81  
 HYPRE\_StructVectorPrint (C++ function), 81  
 HYPRE\_StructVectorRead (C++ function), 81  
 HYPRE\_StructVectorScale (C++ function), 81  
 HYPRE\_StructVectorSetBoxValues (C++ function),  
 80  
 HYPRE\_StructVectorSetBoxValues2 (C++ function),  
 80  
 HYPRE\_StructVectorSetConstantValues (C++  
 function), 79  
 HYPRE\_StructVectorSetRandomValues (C++ func-  
 tion), 80  
 HYPRE\_StructVectorSetValues (C++ function), 79  
 HYPRE\_TempParCSRSetupInterpreter (C++ func-  
 tion), 198  
 HYPRE\_VERSION (C macro), 216  
 HYPRE\_Version (C++ function), 216  
 HYPRE\_VersionNumber (C++ function), 216