



V. A. Dobrev, R. D. Falgout, Tz. V. Kolev, N. A. Petersson, J. B. Schroder, U. M. Yang
Center for Applied Scientific Computing (CASC)
Lawrence Livermore National Laboratory

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-SM-660398

Copyright (c) 2013, Lawrence Livermore National Security, LLC. Produced at the Lawrence Livermore National Laboratory. Written by the XBraid team. LLNL-CODE-660355. All rights reserved.

This file is part of XBraid. Please see the COPYRIGHT and LICENSE file for the copyright notice, disclaimer, and the GNU Lesser General Public License. Email xbraid-support@llnl.gov for support.

XBraid is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License (as published by the Free Software Foundation) version 2.1 dated February 1999.

XBraid is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the IMPLIED WARRANTY OF MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the terms and conditions of the GNU General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Contents

1	Abstract	1
2	Introduction	2
2.1	Meaning of the name	2
2.2	Advice to users	2
2.3	Overview of the XBraid Algorithm	2
2.3.1	Two-Grid Algorithm	6
2.3.2	Summary	7
2.4	Overview of the XBraid Code	7
2.4.1	Parallel decomposition and memory	8
2.4.2	Cycling and relaxation strategies	8
2.4.3	Overlapping communication and computation	9
2.4.4	Configuring the XBraid Hierarchy	10
2.4.5	Halting tolerance	10
2.5	Citing XBraid	11
2.6	Summary	11
3	Examples	12
3.1	The Simplest Example	12
3.2	Two-Dimensional Heat Equation	17
3.2.1	Scaling Study with this Example	21
4	Building XBraid	23
5	Examples: compiling and running	23
6	Coding Style	24
7	Using Doxygen	24
8	Regression Testing	25
9	Drivers: compiling and running	29
10	Module Index	30
10.1	Modules	30
11	Data Structure Index	30
11.1	Data Structures	30

12 File Index	30
12.1 File List	31
13 Module Documentation	31
13.1 User-written routines	31
13.1.1 Detailed Description	32
13.1.2 Typedef Documentation	32
13.2 User interface routines	35
13.2.1 Detailed Description	35
13.3 General Interface routines	36
13.3.1 Detailed Description	37
13.3.2 Typedef Documentation	37
13.3.3 Function Documentation	37
13.4 XBraid status routines	46
13.4.1 Detailed Description	47
13.4.2 Macro Definition Documentation	47
13.4.3 Typedef Documentation	47
13.4.4 Function Documentation	47
13.5 XBraid test routines	58
13.5.1 Detailed Description	58
13.5.2 Function Documentation	58
14 Data Structure Documentation	63
14.1 <code>_braid_AccessStatus</code> Struct Reference	63
14.1.1 Detailed Description	63
14.1.2 Field Documentation	63
14.2 <code>_braid_CoarsenRefStatus</code> Struct Reference	64
14.2.1 Detailed Description	64
14.2.2 Field Documentation	64
14.3 <code>_braid_CommHandle</code> Struct Reference	65
14.3.1 Detailed Description	65
14.3.2 Field Documentation	65
14.4 <code>_braid_Core</code> Struct Reference	66
14.4.1 Detailed Description	67
14.4.2 Field Documentation	67
14.5 <code>_braid_Grid</code> Struct Reference	71
14.5.1 Detailed Description	72
14.5.2 Field Documentation	72

14.6	_braid_StepStatus Struct Reference	73
14.6.1	Detailed Description	74
14.6.2	Field Documentation	74
15	File Documentation	75
15.1	_braid.h File Reference	75
15.1.1	Detailed Description	76
15.1.2	Macro Definition Documentation	77
15.1.3	Function Documentation	77
15.1.4	Variable Documentation	82
15.2	braid.h File Reference	82
15.2.1	Detailed Description	84
15.2.2	Macro Definition Documentation	84
15.3	braid_defs.h File Reference	84
15.3.1	Detailed Description	85
15.3.2	Macro Definition Documentation	85
15.3.3	Typedef Documentation	85
15.3.4	Function Documentation	86
15.3.5	Variable Documentation	86
15.4	braid_status.h File Reference	86
15.4.1	Detailed Description	87
15.5	braid_test.h File Reference	87
15.5.1	Detailed Description	88
Index		89

1 Abstract

This package implements an optimal-scaling multigrid solver for the (non)linear systems that arise from the discretization of problems with evolutionary behavior. Typically, solution algorithms for evolution equations are based on a time-marching approach, solving sequentially for one time step after the other. Parallelism in these traditional time-integration techniques is limited to spatial parallelism. However, current trends in computer architectures are leading towards systems with more, but not faster, processors, i.e., clock speeds are stagnate. Therefore, faster overall runtimes must come from greater parallelism. One approach to achieve parallelism in time is with multigrid, but extending classical multigrid methods for elliptic operators to this setting is a significant achievement. In this software, we implement a non-intrusive, optimal-scaling time-parallel method based on multigrid reduction techniques. The examples in the package demonstrate optimality of our multigrid-reduction-in-time algorithm (MGRIT) for solving a variety of equations in two and three spatial dimensions. These examples can also be used to show that MGRIT can achieve significant speedup in comparison to sequential time marching on modern architectures.

It is **strongly recommended** that you also read [Parallel Time Integration with Multigrid](#) after reading the [Overview of the XBraid Algorithm](#). It is a more in depth discussion of the algorithm and associated ex-

periments.

2 Introduction

2.1 Meaning of the name

We chose the package name XBraid to stand for *Time-Braid*, where *X* is the first letter in the Greek word for time, *Chronos*. The algorithm *braids* together time-grids of different granularity in order to create a multigrid method and achieve parallelism in the time dimension.

2.2 Advice to users

The field of parallel-in-time methods is in many ways under development, and success has been shown primarily for problems with some parabolic character. While there are ongoing projects (here and elsewhere) looking at varied applications such as hyperbolic problems, computational fluid dynamics, power grids, medical applications, and so on, expectations should take this fact into account. Please see our project [publications website](#) for our recent publications concerning some of these varied applications.

That being said, we strongly encourage new users to try our code for their application. Every new application has its own issues to address and this will help us to improve both the algorithm and the software.

For support, please email xbraid-support@llnl.gov. This email address automatically interfaces with our issue tracker and notifies all developers of the pending support request.

2.3 Overview of the XBraid Algorithm

The goal of XBraid is to solve a problem faster than a traditional time marching algorithm. Instead of sequential time marching, XBraid solves the problem iteratively by simultaneously updating a space-time solution guess over all time values. The initial solution guess can be anything, even a random function over space-time. The iterative updates to the solution guess are done by constructing a hierarchy of temporal grids, where the finest grid contains all of the time values for the simulation. Each subsequent grid is a coarser grid with fewer time values. The coarsest grid has a trivial number of time steps and can be quickly solved exactly. The effect is that solutions to the time marching problem on the coarser (i.e., cheaper) grids can be used to correct the original finest grid solution. Analogous to spatial multigrid, the coarse grid correction only *corrects* and *accelerates* convergence to the finest grid solution. The coarse grid does not need to represent an accurate time discretization in its own right. Thus, a problem with many time steps (thousands, tens of thousands or more) can be solved with 10 or 15 XBraid iterations, and the overall time to solution can be greatly sped up. However, this is achieved at the cost of more computational resources.

To understand how XBraid differs from traditional time marching, consider the simple linear advection equation, $u_t = -cu_x$. The next figure depicts how one would typically evolve a solution here with sequential time stepping. The initial condition is a wave, and this wave propagates sequentially across space as time increases.

XBraid instead begins with a solution guess over all of space-time, which for demonstration, we let be random. An XBraid iteration does

1. Relaxation on the fine grid, i.e., the grid that contains all of the desired time values. Relaxation is just a local application of the time stepping scheme, e.g., backward Euler.
2. Restriction to the first coarse grid, i.e., interpolate the problem to a grid that contains fewer time values, say every second or every third time value.
3. Relaxation on the first coarse grid
4. Restriction to the second coarse grid and so on...

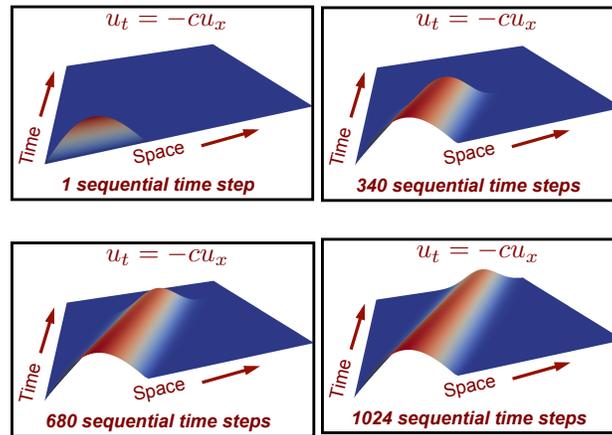


Figure 1: Sequential time stepping.

5. When a coarse grid of trivial size (say 2 time steps) is reached, it is solved exactly.
6. The solution is then interpolated from the coarsest grid to the finest grid

One XBraid iteration is called a *cycle* and these cycles continue until the solution is accurate enough. This is depicted in the next figure, where only a few iterations are required for this simple problem.

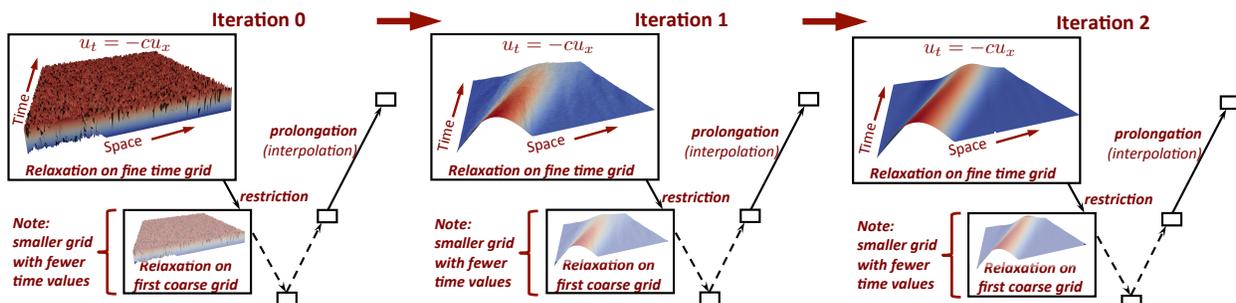


Figure 2: XBraid iterations.

There are a few important points to make.

- The coarse time grids allow for global propagation of information across space-time with only one XBraid iteration. This is visible in the above figure by observing how the solution is updated from iteration 0 to iteration 1.
- Using coarser (cheaper) grids to correct the fine grid is analogous to spatial multigrid.
- Only a few XBraid iterations are required to find the solution over 1024 time steps. Therefore if enough processors are available to parallelize XBraid, we can see a speedup over traditional time stepping (more on this later).
- This is a simple example, with evenly space time steps. XBraid is structured to handle variable time step sizes and adaptive time step sizes.

To firm up our understanding, let's do a little math. Assume that you have a general system of ordinary differential equations (ODEs),

$$u'(t) = f(t, u(t)), \quad u(0) = u_0, \quad t \in [0, T].$$

Next, let $t_i = i\delta t, i = 0, 1, \dots, N$ be a temporal mesh with spacing $\delta t = T/N$, and u_i be an approximation to $u(t_i)$. A

general one-step time discretization is now given by

$$\begin{aligned} u_0 &= g_0 \\ u_i &= \Phi_i(u_{i-1}) + g_i, \quad i = 1, 2, \dots, N. \end{aligned}$$

Traditional time marching would first solve for $i = 1$, then solve for $i = 2$, and so on. For linear time propagators $\{\Phi_i\}$, this can also be expressed as applying a direct solver (a forward solve) to the following system:

$$\mathbf{A}\mathbf{u} \equiv \begin{pmatrix} I & & & & \\ -\Phi_1 & I & & & \\ & \ddots & \ddots & & \\ & & & -\Phi_N & I \end{pmatrix} \begin{pmatrix} u_0 \\ u_1 \\ \vdots \\ u_N \end{pmatrix} = \begin{pmatrix} g_0 \\ g_1 \\ \vdots \\ g_N \end{pmatrix} \equiv \mathbf{g}$$

or

$$\mathbf{A}\mathbf{u} = \mathbf{g}.$$

This process is optimal and $O(N)$, but it is sequential. XBraid achieves parallelism in time by replacing this sequential solve with an optimal multigrid reduction iterative method¹ applied to only the time dimension. This approach is

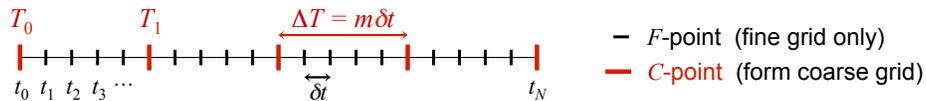
- nonintrusive, in that it coarsens only in time and the user defines Φ . Thus, users can continue using existing time stepping codes by wrapping them into our framework.
- optimal and $O(N)$, but $O(N)$ with a higher constant than time stepping. Thus with enough computational resources, XBraid will outperform sequential time stepping.
- highly parallel

We now describe the two-grid process in more detail, with the multilevel analogue being a recursive application of the process. We also assume that Φ is constant for notational simplicity. XBraid coarsens in the time dimension with factor $m > 1$ to yield a coarse time grid with $N_\Delta = N/m$ points and time step $\Delta T = m\delta t$. The corresponding coarse grid problem,

$$A_\Delta = \begin{pmatrix} I & & & & \\ -\Phi_\Delta & I & & & \\ & \ddots & \ddots & & \\ & & & -\Phi_\Delta & I \end{pmatrix},$$

is obtained by defining coarse grid propagators $\{\Phi_\Delta\}$ which are at least as cheap to apply as the fine scale propagators $\{\Phi\}$. The matrix A_Δ has fewer rows and columns than A , e.g., if we are coarsening in time by 2, A_Δ has one half as many rows and columns.

This coarse time grid induces a partition of the fine grid into C-points (associated with coarse grid points) and F-points, as visualized next. C-points exist on both the fine and coarse time grid, but F-points exist only on the fine time scale.



Every multigrid algorithm requires a relaxation method and an approach to transfer values between grids. Our relaxation scheme alternates between so-called F-relaxation and C-relaxation as illustrated next. F-relaxation updates the F-point values $\{u_j\}$ on interval (T_i, T_{i+1}) by simply propagating the C-point value u_{mi} across the interval using the time propagator $\{\Phi\}$. While this is a sequential process, each F-point interval update is independent from the others and can be computed in parallel. Similarly, C-relaxation updates the C-point value u_{mi} based on the F-point value u_{mi-1} and these updates can also be computed in parallel. This approach to relaxation can be thought of as line relaxation in space in that the residual is set to 0 for an entire time step.

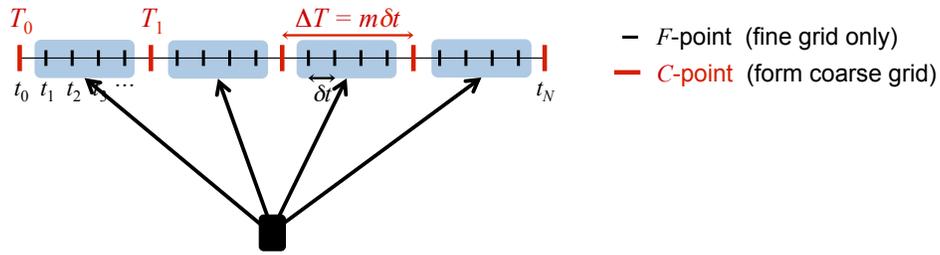


Figure 3: Update all F-point intervals in parallel, using the time propagator Φ .

The F updates are done simultaneously in parallel, as depicted next.

Following the F sweep, the C updates are also done simultaneously in parallel, as depicted next.

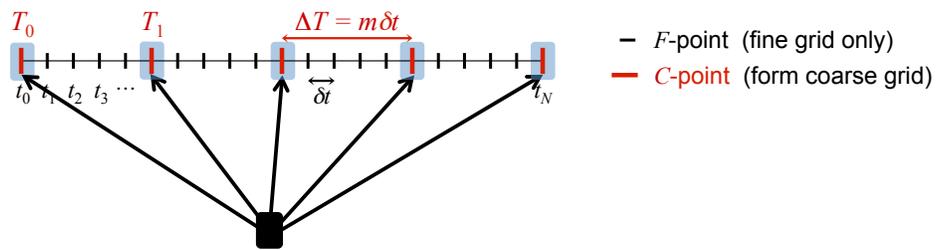
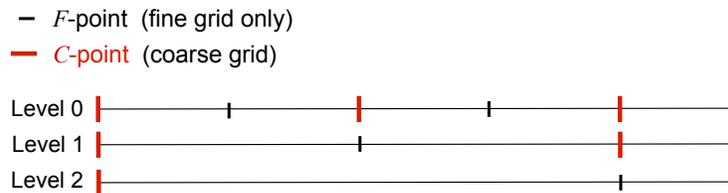


Figure 4: Update all C-points in parallel, using the time propagator Φ .

¹ Ries, Manfred, Ulrich Trottenberg, and Gerd Winter. "A note on MGR methods." Linear Algebra and its Applications 49 (1983): 1-26.

- The coarsest grid is solved exactly, i.e., sequentially, which can be a bottleneck for two-level methods like Parareal, ² but not for a multilevel scheme like XBraid where the coarsest grid is of trivial size.
- By forming the coarse grid to have the same sparsity structure and time stepper as the fine grid, the algorithm can recur easily and efficiently.
- Interpolation is ideal or exact, in that an application of interpolation leaves a zero residual at all F-points.
- The process is applied recursively until a trivially sized temporal grid is reached, e.g., 2 or 3 time points. Thus, the coarsening rate m determines how many levels there are in the hierarchy. For instance in this figure, a 3 level hierarchy is shown. Three levels are chosen because there are six time points, $m = 2$ and $m^2 < 6 \leq m^3$. If the coarsening rate had been $m = 4$ then there would only be two levels because there would be no more points to coarsen!



By default, XBraid will subdivide the time domain into evenly sized time steps. XBraid is structured to handle variable time step sizes and adaptive time step sizes.

2.4 Overview of the XBraid Code

XBraid is designed to run in conjunction with an existing application code that can be wrapped per our interface. This application code will implement some time marching simulation like fluid flow. Essentially, the user has to take their application code and extract a stand-alone time-stepping function Φ that can evolve a solution from one time value to another, regardless of time step size. After this is done, the XBraid code takes care of the parallelism in the time dimension.

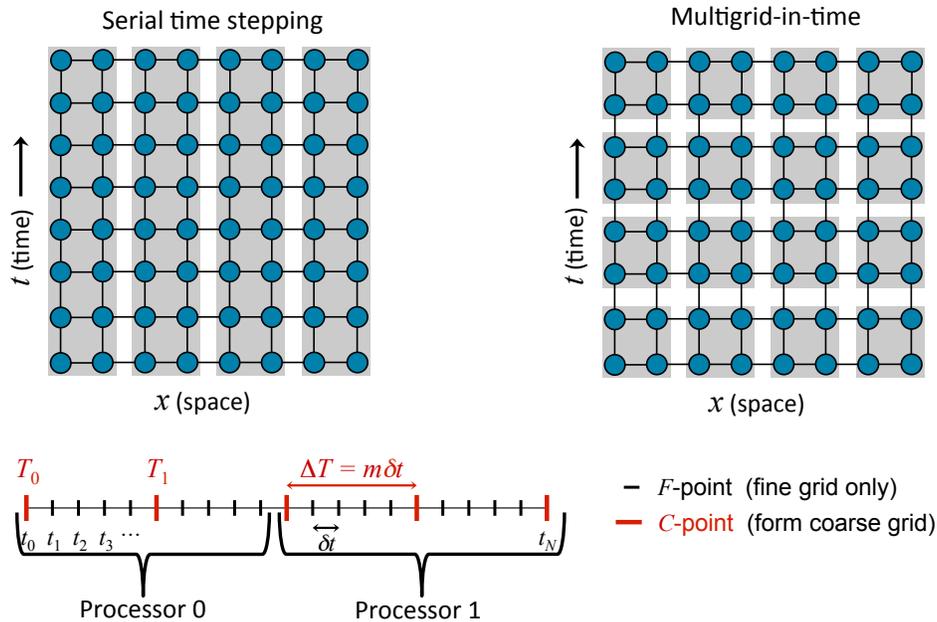
XBraid

- is written in C and can easily interface with Fortran and C++
- uses MPI for parallelism
- self documents through comments in the source code and through *.md files
- functions and structures are prefixed by *braid*
 - User routines are prefixed by `braid_`
 - Developer routines are prefixed by `_braid_`

2.4.1 Parallel decomposition and memory

- XBraid decomposes the problem in parallel as depicted next. As you can see, traditional time stepping only stores one time step at a time, but only enjoys a spatial data decomposition and spatial parallelism. On the other hand, XBraid stores multiple time steps simultaneously and each processor holds a space-time chunk reflecting both the spatial and temporal parallelism.
- XBraid only handles temporal parallelism and is agnostic to the spatial decomposition. See [braid_Split-Commworld](#). Each processor owns a certain number of CF intervals of points. In the following figure, processor 1 and processor 2 each own 2 CF intervals. XBraid distributes intervals evenly on the finest grid.

² Lions, J., Yvon Maday, and Gabriel Turinici. "A"parareal" in time discretization of PDE's." *Comptes Rendus de l'Academie des Sciences Series I Mathematics* 332.7 (2001): 661-668.

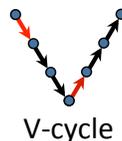


- XBraid increases the parallelism significantly, but now several time steps need to be stored, requiring more memory. XBraid employs two strategies to address the increased memory costs.
 - First, one need not solve the whole problem at once. Storing only one space-time slab is advisable. That is, solve for as many time steps (say k time steps) as you have available memory for. Then move on to the next k time steps.
 - Second, XBraid provides support for storing only C-points. Whenever an F-point is needed, it is generated by F-relaxation. More precisely, only the red C-point time values in the previous figure are stored. Coarsening is usually aggressive with $m = 8, 16, 32, \dots$, so the storage requirements of XBraid are significantly reduced when compared to storing all of the time values.

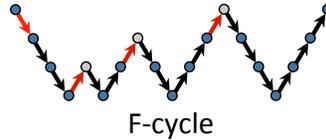
Overall, the memory multiplier per processor when using XBraid is $O(1)$ if space-time coarsening (see [The Simplest Example](#)) is used and $O(\log_m N)$ for time-only coarsening. The time-only coarsening option is the default and requires no user-written spatial interpolation/restriction routines (which is the case for space-time coarsening). We note that the base of the logarithm is m , which can be quite large.

2.4.2 Cycling and relaxation strategies

There are two main cycling strategies available in XBraid, F- and V-cycles. These two cycles differ in how often and the order in which coarse levels are visited. A V-cycle is depicted next, and is a simple recursive application of the [Two-Grid Algorithm](#).



An F-cycle visits coarse grids more frequently and in a different order. Essentially, an F-cycle uses a V-cycle as the post-smoother, which is an expensive choice for relaxation. But, this extra work gives you a closer approximation to a two-grid cycle, and a faster convergence rate at the extra expense of more work. The effectiveness of a V-cycle as a relaxation scheme can be seen in [Figure 2](#), where one V-cycle globally propagates and *smooths* the error. The cycling strategy of an F-cycle is depicted next.



Next, we make a few points about F- versus V-cycles.

- One V-cycle iteration is cheaper than one F-cycle iteration.
- But, F-cycles often converge more quickly. For some test cases, this difference can be quite large. The cycle choice for the best time to solution will be problem dependent. See [Scaling Study with this Example](#) for a case study of cycling strategies.
- For exceptionally strong F-cycles, the option `braid_SetNFMGVcyc` can be set to use multiple V-cycles as relaxation. This has proven useful for some problems with a strongly advective nature.

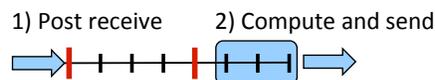
The number of FC relaxation sweeps is another important algorithmic setting. Note that at least one F-relaxation sweep is always done on a level. A few summary points about relaxation are as follows.

- Using FCF, FCFCF, or FCFCFCF relaxation corresponds to passing `braid_SetNRelax` a value of 1, 2 or 3 respectively, and will result in an XBraid cycle that converges more quickly as the number of relaxations grows.
- But as the number of relaxations grows, each XBraid cycle becomes more expensive. The optimal relaxation strategy for the best time to solution will be problem dependent.
- However, a good first step is to try FCF on all levels (i.e., `braid_SetNRelax(core, -1, 1)`).
- A common optimization is to first set FCF on all levels (i.e., `braid_setnrelax(core, -1, 1)`), but then overwrite the FCF option on level 0 so that only F-relaxation is done on level 0, (i.e., `braid_setnrelax(core, 0, 1)`). Another strategy is to use F-relaxation on all levels together with F-cycles.
- See [Scaling Study with this Example](#) for a case study of relaxation strategies.

Last, [Parallel Time Integration with Multigrid](#) has a more in depth case study of cycling and relaxation strategies

2.4.3 Overlapping communication and computation

XBraid effectively overlaps communication and computation. The main computational kernel of XBraid is one relaxation sweep touching all the CF intervals. At the start of a relaxation sweep, each process first posts a non-blocking receive at its left-most point. It then carries out F-relaxation in each interval, starting with the right-most interval to send the data to the neighboring process as soon as possible. If each process has multiple CF intervals at this XBraid level, the strategy allows for complete overlap.



2.4.4 Configuring the XBraid Hierarchy

Some of the more basic XBraid function calls allow you to control aspects discussed here.

- `braid_SetFMG`: switches between using F- and V-cycles.

- [braid_SetMaxIter](#): sets the maximum number of XBraid iterations
- [braid_SetCFactor](#): sets the coarsening factor for any (or all levels)
- [braid_SetNRelax](#): sets the number of CF-relaxation sweeps for any (or all levels)
- [braid_SetRelTol](#), [braid_SetAbsTol](#): sets the stopping tolerance
- [braid_SetMinCoarse](#): sets the minimum possible coarse grid size
- [braid_SetMaxLevels](#): sets the maximum number of levels in the XBraid hierarchy

2.4.5 Halting tolerance

Another important configuration aspect regards setting a residual halting tolerance. Setting a tolerance involves these three XBraid options:

1. [braid_PtFcnSpatialNorm](#)

This user-defined function carries out a spatial norm by taking the norm of a `braid_Vector`. A common choice is the standard Euclidean norm (2-norm), but many other choices are possible, such as an L2-norm based on a finite element space.

2. [braid_SetTemporalNorm](#)

This option determines how to obtain a global space-time residual norm. That is, this decides how to combine the spatial norms returned by [braid_PtFcnSpatialNorm](#) at each time step to obtain a global norm over space and time. It is this global norm that then controls halting.

There are three *tnorm* options supported by [braid_SetTemporalNorm](#). We let the summation index *i* be over all C-point values on the fine time grid, *k* refer to the current XBraid iteration, *r* be residual values, *space_time* norms be a norm over the entire space-time domain and *spatial_norm* be the user-defined spatial norm from [braid_PtFcnSpatialNorm](#). Thus, r_i is the residual at the *i*th C-point, and $r^{(k)}$ is the residual at the *k*th XBraid iteration. The three options are then defined as,

- *tnorm=1*: One-norm summation of spatial norms

$$\|r^{(k)}\|_{\text{space_time}} = \sum_i \|r_i^{(k)}\|_{\text{spatial_norm}}$$

If [braid_PtFcnSpatialNorm](#) is the one-norm over space, then this is equivalent to the one-norm of the global space-time residual vector.

- *tnorm=2*: Two-norm summation of spatial norms

$$\|r^{(k)}\|_{\text{space_time}} = \left(\sum_i \|r_i^{(k)}\|_{\text{spatial_norm}}^2 \right)^{1/2}$$

If [braid_PtFcnSpatialNorm](#) is the Euclidean norm (two-norm) over space, then this is equivalent to the Euclidean-norm of the global space-time residual vector.

- *tnorm=3*: Infinity-norm combination of spatial norms

$$\|r^{(k)}\|_{\text{space_time}} = \max_i \|r_i^{(k)}\|_{\text{spatial_norm}}$$

If [braid_PtFcnSpatialNorm](#) is the infinity-norm over space, then this is equivalent to the infinity-norm of the global space-time residual vector.

The default choice is *tnorm=2*

3. [braid_SetAbsTol](#), [braid_SetRelTol](#)

- If an absolute tolerance is used, then

$$\|r^{(k)}\|_{\text{space_time}} < \text{tol}$$

defines when to halt.

- If a relative tolerance is used, then

$$\frac{\|r^{(k)}\|_{\text{space_time}}}{\|r^{(0)}\|_{\text{space_time}}} < \text{tol}$$

defines when to halt. That is, the current k th residual is scaled by the initial residual before comparison to the halting tolerance. This is similar to typical relative residual halting tolerances used in spatial multigrid, but can be a dangerous choice in this setting.

Care should be practiced when choosing a halting tolerance. For instance, if a relative tolerance is used, then issues can arise when the initial guess is zero for large numbers of time steps. Taking the case where the initial guess (defined by `braided_PtFcnInit`) is 0 for all time values $t > 0$, the initial residual norm will essentially only be nonzero at the first time value,

$$\|r^{(0)}\|_{\text{space_time}} \approx \|r_1^{(k)}\|_{\text{spatial_norm}}$$

This will skew the relative halting tolerance, especially if the number of time steps increases, but the initial residual norm does not.

A better strategy is to choose an absolute tolerance that takes your space-time domain size into account, as in Section [Scaling Study with this Example](#), or to use an infinity-norm temporal norm option.

2.5 Citing XBraid

To cite XBraid, please state in your text the version number from the VERSION file, and please cite the project website in your bibliography as

[1] XBraid: Parallel multigrid in time. <http://llnl.gov/casc/xbraid>.

The corresponding BibTeX entry is

```
@misc{xbraid-package,
  title = {{XB}raid: Parallel multigrid in time},
  howpublished = {\url{http://llnl.gov/casc/xbraid}}
}
```

2.6 Summary

- XBraid applies multigrid to the time dimension.
 - This exposes concurrency in the time dimension.
 - The potential for speedup is large, 10x, 100x, ...
- This is a non-intrusive approach, with an unchanged time discretization defined by user.
- Parallel time integration is only useful beyond some scale. This is evidenced by the experimental results below. For smaller numbers of cores sequential time stepping is faster, but at larger core counts XBraid is much faster.
- The more time steps that you can parallelize over, the better your speedup will be.
- XBraid is optimal for a variety of parabolic problems (see the examples directory).

3 Examples

3.1 The Simplest Example

User Defined Structures and Wrappers

The user must wrap their existing time stepping routine per the XBraid interface. To do this, the user must define two data structures and some wrapper routines. To make the idea more concrete, we now give these function definitions from `examples/ex-01`, which implements a scalar ODE,

$$u_t = \lambda u.$$

The two data structures are:

1. **App:** This holds a wide variety of information and is *global* in that it is passed to every function. This structure holds everything that the user will need to carry out a simulation. Here, this is just the global MPI communicator and few values describing the temporal domain.

```
typedef struct _braid_App_struct
{
    MPI_Comm comm;
    double tstart;
    double tstop;
    int ntime;
} my_App;
```

2. **Vector:** this defines (roughly) a state vector at a certain time value. It could also contain any other information related to this vector which is needed to evolve the vector to the next time value, like mesh information. Here, the vector is just a scalar double.

```
typedef struct _braid_Vector_struct
{
    double value;
} my_Vector;
```

The user must also define a few wrapper routines. Note, that the *app* structure is the first argument to every function.

1. **Step:** This function tells XBraid how to take a time step, and is the core user routine. The user must advance the vector *u* from time *tstart* to time *tstop*. Note how the time values are given to the user through the *status* structure and associated *Get* routine. **Important note:** the *g_i* function from [Overview of the XBraid Algorithm](#) must be incorporated into *Step*, so that the following equation is solved by default.

$$\Phi(u_i) = 0.$$

The *ustop* parameter serves as an approximation to the solution at time *tstop* and is not needed here. It can be useful for implicit schemes that require an initial guess for a linear or nonlinear solver. The use of *fstop* is an advanced parameter (not required) and forms the the right-hand side of the nonlinear problem on the given time grid. This value is only nonzero when providing a residual with [braid_SetResidual](#). More information on how to use this optional feature is given below. Also see the full code listing for this example in `examples/ex-01`.

Here advancing the solution just involves the scalar λ .

```
int
my_Step(braid_App app,
        braid_Vector ustop,
        braid_Vector fstop,
        braid_Vector u,
        braid_StepStatus status)
{
    double tstart; /* current time */
    double tstop; /* evolve to this time*/
    braid_StepStatusGetTstartTstop(status, &tstart, &tstop);

    /* On the finest grid, each value is half the previous value */
```

```

(u->value) = pow(0.5, tstop-tstart)*(u->value);

if (fstop != NULL)
{
    /* Nonzero rhs */
    (u->value) += (fstop->value);
}

/* no refinement */
braid_StepStatusSetRFactor(status, 1);

return 0;
}

```

2. **Init:** This function tells XBraid how to initialize a vector at time t . Here that is just allocating and setting a scalar on the heap.

```

int
my_Init(braid_App app,
        double t,
        braid_Vector *u_ptr)
{
    my_Vector *u;

    u = (my_Vector *) malloc(sizeof(my_Vector));
    if (t == 0.0)
    {
        /* Initial condition */
        (u->value) = 1.0;
    }
    else
    {
        /* Initialize all other time points */
        (u->value) = 0.456;
    }
    *u_ptr = u;

    return 0;
}

```

3. **Clone:** This function tells XBraid how to clone a vector into a new vector.

```

int
my_Clone(braid_App app,
         braid_Vector u,
         braid_Vector *v_ptr)
{
    my_Vector *v;

    v = (my_Vector *) malloc(sizeof(my_Vector));
    (v->value) = (u->value);
    *v_ptr = v;

    return 0;
}

```

4. **Free:** This function tells XBraid how to free a vector.

```

int
my_Free(braid_App app,
        braid_Vector u)
{
    free(u);

    return 0;
}

```

5. **Sum:** This function tells XBraid how to sum two vectors (AXPY operation).

```

int
my_Sum(braid_App app,
        double alpha,
        braid_Vector x,
        double beta,

```

```

        braid_Vector y)
    {
        (y->value) = alpha*(x->value) + beta*(y->value);

        return 0;
    }

```

6. **SpatialNorm:** This function tells XBraid how to take the norm of a *braid_Vector* and is used for halting. This norm is only over space. A common norm choice is the standard Euclidean norm, but many other choices are possible, such as an L2-norm based on a finite element space. The norm choice should be based on what makes sense for your problem. How to accumulate spatial norm values to obtain a global space-time residual norm for halting decisions is controlled by [braid_SetTemporalNorm](#).

```

int
my_SpatialNorm(braid_App    app,
               braid_Vector u,
               double       *norm_ptr)
{
    double dot;

    dot = (u->value)*(u->value);
    *norm_ptr = sqrt(dot);

    return 0;
}

```

7. **Access:** This function allows the user access to XBraid and the current solution vector at time t . This is most commonly used to print solution(s) to screen, file, etc... The user defines what is appropriate output. Notice how you are told the time value t of the vector u and even more information in *astatus*. This lets you tailor the output to only certain time values at certain XBraid iterations. Querying *astatus* for such information is done through *braid_AccessStatusGet**(..)* routines.

The frequency of the calls to *access* is controlled through [braid_SetAccessLevel](#). For instance, if *access_level* is set to 2, then *access* is called every XBraid iteration and on every XBraid level. In this case, querying *astatus* to determine the current XBraid level and iteration will be useful. This scenario allows for even more detailed tracking of the simulation. The default *access_level* is 1 and gives the user access only after the simulation ends and only on the finest time-grid.

Eventually, this routine will allow for broader access to XBraid and computational steering.

See [examples/ex-02](#) and [drivers/drive-04](#) for more advanced uses of the *access* function. In [drive-04](#), *access* is used to write solution vectors to a GLVIS visualization port, and [ex-02](#) uses *access* to write to .vtu files.

```

int
my_Access(braid_App    app,
          braid_Vector u,
          braid_AccessStatus astatus)
{
    MPI_Comm comm = (app->comm);
    double   tstart = (app->tstart);
    double   tstop  = (app->tstop);
    int      ntime  = (app->ntime);
    int      index, myid;
    char     filename[255];
    FILE     *file;
    double   t;

    braid_AccessStatusGetT(astatus, &t);
    index = ((t-tstart) / ((tstop-tstart)/ntime) + 0.1);

    MPI_Comm_rank(comm, &myid);

    sprintf(filename, "%s.%07d.%05d", "ex-01.out", index, myid);
    file = fopen(filename, "w");
    fprintf(file, "%.14e\n", (u->value));
    fflush(file);
}

```

```

    fclose(file);

    return 0;
}

```

8. **BufSize, BufPack, BufUnpack:** These three routines tell XBraid how to communicate vectors between processors. *BufPack* packs a vector into a `void *` buffer for MPI and then *BufUnPack* unpacks it from `void *` to vector. Here doing that for a scalar is trivial. *BufSize* computes the upper bound for the size of an arbitrary vector.

Note how *BufPack* also returns a size pointer. This size pointer should be the exact number of bytes packed, while *BufSize* should provide only an upper-bound on a possible buffer size. This flexibility allows for variable spatial grid sizes to result in smaller messages sent when appropriate. **To avoid MPI issues, it is very important that BufSize be pessimistic, provide an upper bound, and return the same value across processors.**

In general, the buffer should be self-contained. The receiving processor should be able to pull all necessary information from the buffer in order to properly interpret and unpack the buffer.

```

int
my_BufSize(braid_App app,
           int *size_ptr)
{
    *size_ptr = sizeof(double);
    return 0;
}

int
my_BufPack(braid_App app,
           braid_Vector u,
           void *buffer,
           braid_Int *size_ptr)
{
    double *dbuffer = buffer;

    dbuffer[0] = (u->value);
    *size_ptr = sizeof(double);

    return 0;
}

int
my_BufUnpack(braid_App app,
             void *buffer,
             braid_Vector *u_ptr)
{
    double *dbuffer = buffer;
    my_Vector *u;

    u = (my_Vector *) malloc(sizeof(my_Vector));
    (u->value) = dbuffer[0];
    *u_ptr = u;

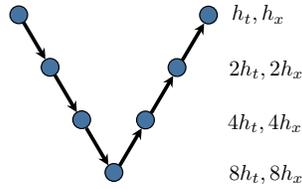
    return 0;
}

```

9. **Scoarsen, SRestrict** (optional): These are advanced options that allow for coarsening in space while you coarsen in time. This is useful for maintaining stable explicit schemes on coarse time scales and is not needed here. See for instance `drivers/drive-04` and `drivers/drive-02` which use these routines.

These functions allow you vary the spatial mesh size on XBraid levels as depicted here where the spatial and temporal grid sizes are halved every level.

10. **Residual** (optional): A user-defined residual can be provided with the function [braid_SetResidual](#) and can result in substantial computational savings, as explained below. *Residual* defines the nonlinear equation at each time step, i.e., it evaluates one block-row of the global space-time operator A with a right-hand side (a residual computation). It defines the equation which *Step* must solve. Because XBraid assumes a one-step method, the



equation to solve on each grid level has the form

$$A_i(u_i, u_{i-1}) = f_i,$$

where $A_i(\cdot)$ is the *Residual* function and $f_i = 0$ on the finest grid level. The nonzero right-hand-side on each coarse grid is needed to implement the FAS algorithm. The *Step* function provides an approximation to the solution u_i of this equation. That is,

$$u_i \approx \Phi(u_{i-1}, f_i),$$

where $\Phi(u, f)$ is the *Step* function augmented with a right-hand side term. When *Residual* is provided, *Step* does not need to produce an accurate time step, especially when all time points are stored (see [braid_SetStorage](#)). In this case, substantial computational savings are possible because one application of *Step* becomes much cheaper than a step for sequential time-stepping. For example if *Step* were backward Euler for a linear problem, *Residual* would essentially be a matrix-vector product, while *Step* would be an inaccurate solve of the corresponding linear system.

Users should write the *Residual* and *Step* routines such that the following holds:

$$A_i(\Phi(u_{i-1}, f_i), u_{i-1}) \approx f_i.$$

Note that when *Residual* is not provided, XBraid defines the residual in terms of the *Step* function as follows:

$$A_i(u_i, u_{i-1}) = u_i - \Phi(u_{i-1}) = 0.$$

In this case, we have exact equality above (i.e., \approx becomes $=$). In addition, the nonzero right-hand-side *fstop* needed on coarse grids to do FAS does not need to be provided to the user because the solution to the non-homogeneous equation is simply $u_i = \Phi(u_{i-1}) + f_i$, and this can easily be computed internally in XBraid. In other words, the *fstop* parameter for *Step* is always zero by default. Also note that in this default setting, *Step* must always be an accurate time step on the finest grid level so that residual computations are accurate.

11. Adaptive and variable time stepping are available by first calling the function [braid_SetRefine](#) in the main driver and then using [braid_StepStatusSetRRFactor](#) in the *Step* routine to set a refinement factor for interval [*tstart*, *tstop*]. In this way, user-defined criteria can subdivide intervals on the fly and adaptively refine in time. For instance, returning a refinement factor of 4 in *Step* will tell XBraid to subdivide that interval into 4 evenly spaced smaller intervals for the next iteration. Refinement can only be done on the finest XBraid level.

In this example, no refinement is being done (factor = 1). Currently, each refinement factor is constrained to be no larger than the coarsening factor. The final time grid is constructed adaptively in an FMG-like cycle by refining the initial grid according to the requested refinement factors. Refinement stops when the requested factors are all one or when various upper bounds are reached such as the max number of time points or max number of time grid refinement levels allowed.

Running XBraid for this Example

A typical flow of events in the *main* function is to first initialize the *app* structure.

```
/* set up app structure */
app = (my_App *) malloc(sizeof(my_App));
(app->comm) = comm;
```

```
(app->tstart) = tstart;
(app->tstop)  = tstop;
(app->ntime)  = ntime;
```

Then, the data structure definitions and wrapper routines are passed to XBraid. The core structure is used by XBraid for internal data structures.

```
braid_Core core;
braid_Init(MPI_COMM_WORLD, comm, tstart, tstop, ntime, app,
           my_Step, my_Init, my_Clone, my_Free, my_Sum, my_SpatialNorm,
           my_Access, my_BufSize, my_BufPack, my_BufUnpack, &core);
```

Then, XBraid options are set.

```
braid_SetPrintLevel( core, 1);
braid_SetMaxLevels(core, max_levels);
braid_SetNRelax(core, -1, nrelax);
braid_SetAbsTol(core, tol);
braid_SetCFactor(core, -1, cfactor);
braid_SetMaxIter(core, max_iter);
```

Then, the simulation is run.

```
braid_Drive(core);
```

Then, we clean up.

```
braid_Destroy(core);
```

Finally, to run ex-01, type

```
ex-01 -ml 5
```

This will run ex-01. See `examples/ex-0*` for more extensive examples.

3.2 Two-Dimensional Heat Equation

In this example, we assume familiarity with [The Simplest Example](#) and describe the major ways in which this example differs. This example is a full space-time parallel example, as opposed to [The Simplest Example](#), which implements only a scalar ODE for one degree-of-freedom in space. We solve the heat equation in 2D,

$$\delta/\delta_t u(x,y,t) = \Delta u(x,y,t) + g(x,y,t).$$

For spatial parallelism, we rely on the `hypre` package where the `SemiStruct` interface is used to define our spatial discretization stencil and form our time stepping scheme, the backward Euler method. The spatial discretization is just the standard 5-point finite difference stencil (`[-1; -1, 4, -1; -1]`), scaled by mesh widths, and the PFMG solver is used for the solves required by backward Euler. Please see the `hypre` manual and examples for more information on the `SemiStruct` interface and PFMG. Although, the `hypre` specific calls have mostly been abstracted away for this example, and so it is not necessary to be familiar with the `SemiStruct` interface for this example.

This example consists of three files and two executables.

- `examples/ex-02-serial.c`: This file compiles into its own executable `ex-02-serial` and represents a simple example user application. This file supports only parallelism in space and represents a basic approach to doing efficient sequential time stepping with the backward Euler scheme. Note that the `hypre` solver used (PFMG) to carry out the time stepping is highly efficient.

- `examples/ex-02.c`: This file compiles into its own executable `ex-02` and represents a basic example of wrapping the user application `ex-02-serial`. We will go over the wrappers below.
- `ex-02-lib.c`: This file contains shared functions used by the time-serial version and the time-parallel version. This is where most of the hypr specific calls reside. This file provides the basic functionality of this problem. For instance, `take_step(u, tstart, tstop, ...)` carries out a step, moving the vector `u` from time `tstart` to time `tstop` and `setUpImplicitMatrix(...)` constructs the matrix to be inverted by PFMG for the backward Euler method.

User Defined Structures and Wrappers

We now discuss in more detail the important data structures and wrapper routines in `examples/ex-02.c`. The actual code for this example is quite simple and it is recommended to read through it after this overview.

The two data structures are:

1. **App**: This holds a wide variety of information and is *global* in that it is passed to every user function. This structure holds everything that the user will need to carry out a simulation. One important structure contained in the *app* is the *simulation_manager*. This is a structure native to the user code `ex-02-lib.c`. This structure conveniently holds the information needed by the user code to carry out a time step. For instance,

```
app->man->A
is the time stepping matrix,
app->man->solver
is the hypr PFMG solver object,
app->man->dt
is the current time step size. The app is defined as
```

```
typedef struct _braid_App_struct {
    MPI_Comm      comm;           /* global communicator */
    MPI_Comm      comm_t;        /* communicator for parallelizing in time */
    MPI_Comm      comm_x;        /* communicator for parallelizing in space */
    int           pt;            /* number of processors in time */
    simulation_manager *man;      /* user's simulation manager structure */
    HYPRE_SStructVector e;        /* temporary vector used for error computations */
    int           nA;            /* number of spatial matrices created */
    HYPRE_SStructMatrix *A;       /* array of spatial matrices, size nA, one per level*/
    double        *dt_A;         /* array of time step sizes, size nA, one per level*/
    HYPRE_StructSolver *solver;   /* array of PFMG solvers, size nA, one per level*/
    int           use_rand;       /* binary value, use random or zero initial guess */
    int           *runtime_max_iter; /* runtime info for number of PFMG iterations*/
    int           *max_iter_x;    /* maximum iteration limits for PFMG */
} my_App;
```

The *app* contains all the information needed to take a time step with the user code for an arbitrary time step size. See the *Step* function below for more detail.

1. **Vector**: this defines a state vector at a certain time value. Here, the vector is a structure containing a native hypr data-type, the *SStructVector*, which describes a vector over the spatial grid. Note that *my_Vector* is used to define *braid_Vector*.

```
typedef struct _braid_Vector_struct {
    HYPRE_SStructVector x;
} my_Vector;
```

The user must also define a few wrapper routines. Note, that the *app* structure is the first argument to every function.

1. **Step**: This function tells XBraid how to take a time step, and is the core user routine. This function advances the vector `u` from time `tstart` to time `tstop`. A few important things to note are as follows.
 - The time values are given to the user through the *status* structure and associated *Get* routines.
 - The basic strategy is to see if a matrix and solver already exist for this `dt` value. If not, generate a new matrix and solver and store them in the *app* structure. If they do already exist, then re-use the data.

- To carry out a step, the user routines from `ex-02-lib.c` rely on a few crucial data members `man->dt`, `man->A` and `man->solver`. We overwrite these members with the correct information for the time step size in question. Then, we pass `man` and `u` to the user function `take_step(...)` which evolves `u`.
- The forcing term g_i is wrapped into the `take_step(...)` function. Thus, $\Phi(u_i) \rightarrow u_{i+1}$.

```

int my_Step(braid_App      app,
           braid_Vector   u,
           braid_StepStatus status)
{
    double tstart;          /* current time */
    double tstop;          /* evolve u to this time*/
    int i, A_idx;
    int iters_taken = -1;

    /* Grab status of current time step */
    braid_StepStatusGetTstartTstop(status, &tstart, &tstop);

    /* Check matrix lookup table to see if this matrix already exists*/
    A_idx = -1.0;
    for( i = 0; i < app->nA; i++ ){
        if( fabs( app->dt_A[i] - (tstop-tstart) )/(tstop-tstart) < 1e-10) {
            A_idx = i;
            break;
        }
    }

    /* We need to "trick" the user's manager with the new dt */
    app->man->dt = tstop - tstart;

    /* Set up a new matrix and solver and store in app */
    if( A_idx == -1.0 ){
        A_idx = i;
        app->nA++;
        app->dt_A[A_idx] = tstop-tstart;

        setUpImplicitMatrix( app->man );
        app->A[A_idx] = app->man->A;

        setUpStructSolver( app->man, u->x, u->x );
        app->solver[A_idx] = app->man->solver;
    }

    /* Time integration to next time point: Solve the system Ax = b.
     * First, "trick" the user's manager with the right matrix and solver */
    app->man->A = app->A[A_idx];
    app->man->solver = app->solver[A_idx];
    ...
    /* Take step */
    take_step(app->man, u->x, tstart, tstop);
    ...
    return 0;
}

```

2. There are other functions, **Init**, **Clone**, **Free**, **Sum**, **SpatialNorm**, **Access**, **BufSize**, **BufPack** and **BufUnpack**, which also must be written. These functions are all simple for this example, as for the case of [The Simplest Example](#). All we do here is standard operations on a spatial vector such as initialize, clone, take an inner-product, pack, etc... We refer the reader to `ex-02.c`.

Running XBraid for this Example

To initialize and run XBraid, the procedure is similar to [The Simplest Example](#). Only here, we have to both initialize the user code and XBraid. The code that is specific to the user's application comes directly from the existing serial simulation code. If you compare `ex-02-serial.c` and `ex-02.c`, you will see that most of the code setting up the user's data structures and defining the wrapper functions are simply lifted from the serial simulation.

Taking excerpts from the function *main()* in *ex-02.c*, we first initialize the user's simulation manager with code like

```
...
app->man->px    = 1; /* my processor number in the x-direction */
app->man->py    = 1; /* my processor number in the y-direction */
               /* px*py=num procs in space */
app->man->nx    = 17; /* number of points in the x-dim */
app->man->ny    = 17; /* number of points in the y-dim */
app->man->nt    = 32; /* number of time steps */
...
```

We also define default XBraid parameters with code like

```
...
max_levels    = 15; /* Max levels for XBraid solver */
min_coarse    = 3; /* Minimum possible coarse grid size */
nrelax        = 1; /* Number of CF relaxation sweeps on all levels */
...
```

The XBraid app must also be initialized with code like

```
...
app->comm      = comm;
app->tstart    = tstart;
app->tstop     = tstop;
app->ntime     = ntime;
```

Then, the data structure definitions and wrapper routines are passed to XBraid.

```
braid_Core core;
braid_Init(MPI_COMM_WORLD, comm, tstart, tstop, ntime, app,
           my_Step, my_Init, my_Clone, my_Free, my_Sum, my_SpatialNorm,
           my_Access, my_BufSize, my_BufPack, my_BufUnpack, &core);
```

Then, XBraid options are set with calls like

```
...
braid_SetPrintLevel( core, 1);
braid_SetMaxLevels(core, max_levels);
braid_SetNRelax(core, -1, nrelax);
...
```

Then, the simulation is run.

```
braid_Drive(core);
```

Then, we clean up.

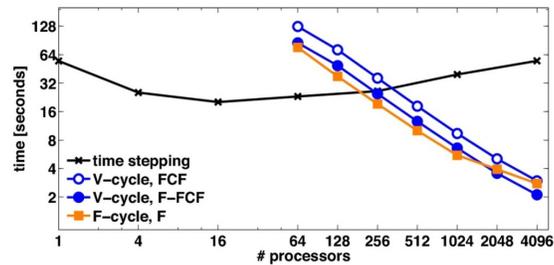
```
braid_Destroy(core);
```

Finally, to run *ex-02*, type

```
ex-02 -help
```

As a simple example, try the following.

```
mpirun -np 8 ex-02 -pgrid 2 2 2 -nt 256
```



3.2.1 Scaling Study with this Example

Here, we carry out a simple strong scaling study for this example. The "time stepping" data set represents sequential time stepping and was generated using `examples/ex-02-serial`. The time-parallel data set was generated using `examples/ex-02`. The problem setup is as follows.

- Backwards Euler is used as the time stepper. This is the only time stepper supported by `ex-02`.
- We used a Linux cluster with 4 cores per node, a Sandybridge Intel chipset, and a fast Infiniband interconnect.
- The space-time problem size was $129^2 \times 16,192$ over the unit cube $[0, 1] \times [0, 1] \times [0, 1]$.
- The coarsening factor was $m = 16$ on the finest level and $m = 2$ on coarser levels.
- Since 16 processors optimized the serial time stepping approach, 16 processors in space are also used for the XBraid experiments. So for instance 512 processors in the plot corresponds to 16 processors in space and 32 processors in time, $16 * 32 = 512$. Thus, each processor owns a space-time hypercube of $(129^2/16) \times (16, 192/32)$. See [Parallel decomposition and memory](#) for a depiction of how XBraid breaks the problem up.
- Various relaxation and V and F cycling strategies are experimented with.
 - *V-cycle, FCF* denotes V-cycles and FCF-relaxation on each level.
 - *V-cycle, F-FCF* denotes V-cycles and F-relaxation on the finest level and FCF-relaxation on all coarser levels.
 - *F-cycle, F* denotes F-cycles and F-relaxation on each level.
- The initial guess at time values for $t > 0$ is zero, which is typical.
- The halting tolerance corresponds to a discrete L2-norm and was

$$\text{tol} = \frac{10^{-8}}{\sqrt{(h_x)^2 h_t}},$$

where h_x and h_t are the spatial and temporal grid spacings, respectively.

This corresponds to passing `tol` to `braid_SetAbsTol`, passing 2 to `braid_SetTemporalNorm` and defining `braid_PtFcnSpatialNorm` to be the standard Euclidean 2-norm. All together, this appropriately scales the space-time residual in way that is relative to the number of space-time grid points (i.e., it approximates the L2-norm).

To re-run this scaling study, a sample run string for `ex-02` is

```
mpirun -np 64 ex-02 -pgrid 4 4 4 -nx 129 129 -nt 16129 -cf0 16 -cf 2 -nu 1 -use_rand 0
```

To re-run the baseline sequential time stepper, `ex-02-serial`, try

```
mpirun -np 64 ex-02-serial -pgrid 8 8 -nx 129 129 -nt 16129
```

For explanations of the command line parameters, type

```
ex-02-serial -help
ex-02 -help
```

Regarding the performance, we can say

- The best speedup is 10x and this would grow if more processors were available.
- Although not shown, the iteration counts here are about 10-15 XBraid iterations. See [Parallel Time Integration with Multigrid](#) for the exact iteration counts.
- At smaller core counts, serial time stepping is faster. But at about 256 processors, there is a crossover and XBraid is faster.
- You can see the impact of the cycling and relaxation strategies discussed in [Cycling and relaxation strategies](#). For instance, even though *V-cycle, F-FCF* is a weaker relaxation strategy than *V-cycle, FCF* (i.e., the XBraid convergence is slower), *V-cycle, F-FCF* has a faster time to solution than *V-cycle, FCF* because each cycle is cheaper.
- In general, one level of aggressive coarsening (here by a factor 16) followed by slower coarsening was found to be best on this machine.

Achieving the best speedup can require some tuning, and it is recommended to read [Parallel Time Integration with Multigrid](#) where this 2D heat equation example is explored in much more detail.

Running and Testing XBraid

The best overall test for XBraid, is to set the maximum number of levels to 1 (see [braid_SetMaxLevels](#)) which will carry out a sequential time stepping test. Take the output given to you by your *Access* function and compare it to output from a non-XBraid run. Is everything OK? Once this is complete, repeat for multilevel XBraid, and check that the solution is correct (that is, it matches a serial run to within tolerance).

At a lower level, to do sanity checks of your data structures and wrapper routines, there are also XBraid test functions, which can be easily run. The test routines also take as arguments the *app* structure, spatial communicator *comm_x*, a stream like *stdout* for test output and a time step size *dt* to test. After these arguments, function pointers to wrapper routines are the rest of the arguments. Some of the tests can return a boolean variable to indicate correctness.

```

/* Test init(), access(), free() */
braid_TestInitAccess( app, comm_x, stdout, dt, my_Init, my_Access, my_Free);

/* Test clone() */
braid_TestClone( app, comm_x, stdout, dt, my_Init, my_Access, my_Free, my_Clone);

/* Test sum() */
braid_TestSum( app, comm_x, stdout, dt, my_Init, my_Access, my_Free, my_Clone, my_Sum);

/* Test spatialnorm() */
correct = braid_TestSpatialNorm( app, comm_x, stdout, dt, my_Init, my_Free, my_Clone,
                                my_Sum, my_SpatialNorm);

/* Test bufsize(), bufpack(), bufunpack() */
correct = braid_TestBuf( app, comm_x, stdout, dt, my_Init, my_Free, my_Sum, my_SpatialNorm,
                        my_BufSize, my_BufPack, my_BufUnpack);

/* Test coarsen and refine */
correct = braid_TestCoarsenRefine(app, comm_x, stdout, 0.0, dt, 2*dt, my_Init,
                                my_Access, my_Free, my_Clone, my_Sum, my_SpatialNorm,
                                my_CoarsenInjection, my_Refine);
correct = braid_TestCoarsenRefine(app, comm_x, stdout, 0.0, dt, 2*dt, my_Init,
                                my_Access, my_Free, my_Clone, my_Sum, my_SpatialNorm,
                                my_CoarsenBilinear, my_Refine);

```

More Complicated Examples

We have Fortran90 and C++ interfaces. See `examples/ex-01f.f90`, `braid.hpp` and the various C++ examples in `drivers/drive-*.cpp`. For discussion of more complex problems please see our project [publications website](#) for our recent publications concerning some of these varied applications.

4 Building XBraid

- Copyright information and licensing restrictions can be found in the files COPYRIGHT and LICENSE.
- To specify the compilers, flags and options for your machine, edit `makefile.inc`. For now, we keep it simple and avoid using `configure` or `cmake`.
- To make the library, `libbraid.a`,

```
$ make
```
- To make the examples

```
$ make all
```
- The makefile lets you pass some parameters like `debug` with

```
$ make debug=yes
```

 or

```
$ make all debug=yes
```

 It would also be easy to add additional parameters, e.g., to compile with `insure`.
- To set compilers and library locations, look in `makefile.inc` where you can set up an option for your machine to define simple stuff like

```
CC = mpicc
MPICC = mpicc
MPICXX = mpiCC
LFLAGS = -lm
```

5 Examples: compiling and running

Type

```
ex-0* -help
```

for instructions on how to run any example.

To run the examples, type

```
mpirun -np 4 ex-* [args]
```

1. `ex-01` is the simplest example. It implements a scalar ODE and can be compiled and run with no outside dependencies. See Section ([The Simplest Example](#)) for more discussion of this example.
2. `ex-02` implements the 2D heat equation on a regular grid. You must have `hypre` installed and these variables in `examples/Makefile` set correctly

```
HYPRE_DIR = ../../linear_solvers/hypre
HYPRE_FLAGS = -I$(HYPRE_DIR)/include
HYPRE_LIB = -L$(HYPRE_DIR)/lib -lhypre
```

 Only implicit time stepping (backward Euler) is supported. See Section ([Two-Dimensional Heat Equation](#)) for more discussion of this example. The driver

```
drivers/drive-02
```

 is a more sophisticated version of this simple example that supports explicit time stepping and spatial coarsening.
3. `ex-burgers` implements Burger's equation in one spatial dimension, supports spatial coarsening and forward and backward Euler. It is under development, but a fun example. See also `viz-burgers.py` for visualizing the output.
4. `ex-lorenz` implements the Lorenz equation and is under development. See also `viz-lorenz.py` for visualizing the output.

6 Coding Style

Code should follow the `ellemte1` style. See `braid/misc/sample_c_code.c`, and for emacs and vim style files, see `braid/misc/sample.vimrc`, and `braid/misc/sample.emacs`.

7 Using Doxygen

To build the documentation, `doxygen` must be version 1.8 or greater. XBraid documentation uses a `markdown` syntax both in source file comments and in `*.md` files.

To make the documentation,

```
$ make user_manual
$ acread user_manual.pdf
```

or to make a more extensive reference manual for developers,

```
$ make developer_manual
$ acread developer_manual.pdf
```

Developers can run `doxygen` from a precompiled binary, which may or may not work for your machine,

```
/usr/casc/hypre/braid/share/doxygen/bin/doxygen
```

or build `doxygen` from

```
/usr/casc/hypre/braid/share/doxygen.tgz
```

- Compiling `doxygen` requires a number of dependencies like Bison, GraphViz and Flex. Configure will tell you what you're missing
- Unpack `doxygen.tgz`, then from the `doxygen` directory

```
./configure --prefix some_dir_in_your_path
make
make install
```

Documentation Strategy

- The `doxygen` comments are to be placed in the header files.
- A sample function declaration using the documentation approach using `markdown` (including typesetting equations) is in `braid.h` for the function `braid_Init()`
- A sample structure documentation is in `_braid.h` for `_braid_Core_struct`
- Descriptors for files can also be added, as at the top of `braid.h`
- The Doxygen manual is at <http://www.stack.nl/~dimitri/doxygen/manual/index.html>

XBraid Doxygen details

The user and developer manuals are ultimately produced by `Latex`. The formatting of the manuals is configured according to the following.

- `docs/local_doxygen.sty`
 - `Latex` style file used
- `docs/user_manual_header.tex`

- User manual title page and header info
- docs/developer_manual_header.tex
 - Developer manual title page and header info
- *.md
 - Any file ending in .md is extra documentation in markdown format, like [Introduction.md](#) or the various Readme.md files in each directory. This material can be read in plain-text or when it's compiled by Doxygen and Latex.
- docs/user_manual.conf
 - Doxygen configure file for the user manual
 - The FILE_NAMES tag is a filter to only include the user interface routines in [braid.h](#)
 - The INPUT tag orders the processing of the files and hence the section ordering
- docs/reference_manual.conf
 - Same as user_manual.conf, but the FILE_NAMES tag does not exclude any file from processing.
- docs/img
 - Contains the images
- To regenerate generic doxygen latex files, type


```
$ doxygen -w latex header.tex footer.tex doxygen.sty doxy.conf
```

 If this is done, then the .conf file must be changed to use the new header file and to copy the local_doxygen.sty file to the latex directory.

8 Regression Testing

Overview

- There are three levels in the testing framework. At each level, the fine-grain output from a `testscript.sh` is dumped into a directory `testscript.dir`, with the standard out and error stored in `testscript.out` and `testscript.err`. The test `testscript.sh` passes if `testscript.err` is empty (nothing is written to standard error).
- Basic instructions: run a test with a command like


```
$ ./test.sh diffusion2D.sh
```

 Then, see if `diffusion2D.err` is of size 0. If it is not, look at it's contents to see which test failed.
- To add a new regression test, create a new lowest level script like `diffusion2D.sh` and then call it from a machine script at level 2.
- Regression tests should be run before pushing code. It is recommended to run the basic (lowest level) tests like `diffusion2d.sh` or machine test like `machine-tux.sh`

Lowest Level Test Scripts

As an example, here we look at one of the lowest level tests, the `diffusion2d` test.

Files used:

- `test.sh`
- `diffusion2D.sh`
- `diffusion2D.saved`

Output:

- diffusion2D.dir
- diffusion2D.err
- diffusion2D.out

At this level, we execute

```
$ ./test.sh diffusion2D.sh
```

or just

```
$ ./diffusion2D.sh
```

The script `diffusion2D.sh` must create `diffusion2D.dir` and place all fine-grain test output in this directory. `test.sh` captures the standard out and error in `diffusion2D.out` and `diffusion2D.err`. The test `diffusion2D.sh` passes if `diffusion2D.err` is empty (nothing is written to standard error).

The strategy for low level scripts like `diffusion2D.sh` is to run a sequence of tests such as

```
$ mpirun -np 1 ../examples/ex-02 -pgrid 1 1 1 -nt 256
$ mpirun -np 4 ../examples/ex-02 -pgrid 1 1 4 -nt 256
```

The output from the first `mpirun` test must then be written to files named

```
diffusion2D.dir/unfiltered.std.out.0
diffusion2D.dir/std.out.0
diffusion2D.dir/std.err.0
```

and the second `mpirun` test similarly writes the files

```
diffusion2D.dir/unfiltered.std.out.1
diffusion2D.dir/std.out.1
diffusion2D.dir/std.err.1
```

Subsequent tests are written to higher numbered files. The `unfiltered.std.out.num` file contains all of the standard out for the test, while `std.out.num` contains filtered output (usually from a `grep` command) and could contain the output lines such as iteration numbers and number of levels. The file `std.err.num` contains the standard error output.

To see if a test ran correctly, `std.out.num` is compared to saved output in `diffusion2D.saved`. The file `diffusion2D.saved` contains the concatenated output from all the tests that `diffusion2D.sh` will run. For the above example, this file could look like

```
# Begin Test 1
number of levels      = 6
iterations            = 16
# Begin Test 2
number of levels      = 4
iterations            = 8
```

This saved output is split into an individual file for each test (using `# Begin Test` as a delimiter) and these new files are placed in `diffusion2D.dir`. So, after running these two regression tests, `diffusion2D.dir` will contain

```
diffusion2D.saved.0
diffusion2D.saved.1
unfiltered.std.out.0
std.out.0
std.err.0
unfiltered.std.out.1
std.out.1
std.err.1
```

An individual test has passed if `std.err.num` is empty. The file `std.err.num` contains a diff between `diffusion2D.save.num` and `std.out.num` (the diff ignores whitespace and the delimiter `# Begin Test`). Last in the directory where you ran `./test.sh diffusion2d.sh`, the files

```
diffusion2D.err
diffusion2D.out
```

will be created. If all the tests passed then `diffusion2D.err` will be empty. Otherwise, it will contain the filenames of the `std.err.num` files that are non-empty, representing failed tests.

Level 2 Scripts

As an example, here we look at one of the Level 2 tests, the `machine-tux` test that Jacob runs.

Files used:

- `machine-tux.sh`

Output:

- `machine-tux.dir`
- `machine-tux.err` (only generated if `autotest.sh` is used to run `machine-tux.sh`)
- `machine-tux.out` (only generated if `autotest.sh` is used to run `machine-tux.sh`)

At this level, we execute

```
./machine-tux.sh
```

The `autotest` framework (`autotest.sh`) calls machine scripts in this way. Each machine script should be short and call lower-level scripts like `diffusion2D.sh`. The output from lower-level scripts must be moved to `machine-tux.dir` like this:

```
$ ./test.sh diffusion2D.sh
$ mv -f diffusion2D.dir machine-tux.dir
$ mv -f diffusion2D.out machine-tux.dir
$ mv -f diffusion2D.err machine-tux.dir
```

All error files from `diffusion2D.sh` will be placed in `machine-tux.dir`, so if `machine-tux.dir` has all zero `*.err` files, then the `machine-tux` test has passed.

To begin testing on a new machine, like `vulcan`, add a new machine script similar to `machine-tux.sh` and change `autotest.sh` to recognize and run the new machine test. To then use `autotest.sh` with the machine script, you'll have to set up a passwordless connection from the new machine to

```
/usr/casc/hypre/braid/testing
```

Level 3 Script

Here we look at the highest level, where `autotest.sh` runs all of the level 2 machine tests and emails out the results.

Files used:

- `autotest.sh`

Output:

- `test/autotest_finished`

- `/usr/casc/hypre/braid/testing/AUTOTEST-20**.**.**-Day`
- Email to recipients listed in `autotest.sh`

At the highest level sits `autotest.sh` and is called automatically as a cronjob. If you just want to check to see if you've broken anything with a commit, just use lower level scripts.

There are four steps to running `autotest`.

- Step 1

```
$ ./autotest.sh -init
```

 will do a pull from master for the current working repository and recompile Braid. The `autotest` output files (`autotest.err` and `autotest.out`) and the output directory (`autotest_finished`) are initialized.
- Step 2

```
$ ./autotest.sh -tux343
```

 will run the autotests on `tux343`. This command will look for a `machine-tux.sh`, and execute it, moving the resulting


```
machine-tux.dir
machine-tux.err
machine-tux.out
```

 into `test/autotest_finished`.
- Step 3

```
$ ./autotest.sh -remote-copy
```

 will copy `/test/autotest_finished/*` to a time-stamped directory such as `/usr/casc/hypre/braid/testing/AUTOTEST-2013.11.18-Mon`
 Alternatively,

```
$ ./autotest.sh -remote-copy tux343
```

 will ssh through `tux343` to copy to `/usr/casc`. Multiple machines may independently be running regression tests and then copy to `AUTOTEST-2013.11.18-Mon`.
- Step 4

```
$ ./autotest.sh -summary-email
```

 will email everyone listed in the `$email_list` (an `autotest.sh` variable)

Cronfile

To add entries to your crontab, First, put your new cronjob lines into `cronfile`. Then see what you already have in your crontab file with

```
$ crontab -l
```

Next, append to `cronfile` whatever you already have

```
$ crontab -l >> cronfile
```

Finally, tell `crontab` to use your `cronfile`

```
$ crontab cronfile
```

Then make sure it took affect with

```
$ crontab -l
```

Crontab entry format uses `'*` to mean "every" and `'*/m'` to mean "every m-th". The first five entries on each line correspond respectively to:

- minute (0-56)
- hour (0-23)

- day of month (1-31)
- month (1-12)
- day of week (0-6)(0=Sunday)

Jacob's crontab (on tux343):

```
00 01 * * * source /etc/profile; source $HOME/.bashrc; cd $HOME/joint_repos/braid/test; ./autotest.sh -init
10 01 * * * source /etc/profile; source $HOME/.bashrc; cd $HOME/joint_repos/braid/test; ./autotest.sh -tux343
40 01 * * * source /etc/profile; source $HOME/.bashrc; cd $HOME/joint_repos/braid/test; ./autotest.sh -remote-c
50 01 * * * source /etc/profile; source $HOME/.bashrc; cd $HOME/joint_repos/braid/test; ./autotest.sh -summary-
00 02 * * * source /etc/profile; source $HOME/.bashrc; cd $HOME/joint_repos/braid/test; ./autotest.sh -create-t
```

9 Drivers: compiling and running

Type

```
drive-0* -help
```

for instructions on how to run any driver.

To run the examples, type

```
mpirun -np 4 drive-* [args]
```

1. drive-02 implements the 2D heat equation on a regular grid. You must have `hypre` installed and these variables in `examples/Makefile` set correctly

```
HYPRE_DIR = ../../linear_solvers/hypre
HYPRE_FLAGS = -I$(HYPRE_DIR)/include
HYPRE_LIB = -L$(HYPRE_DIR)/lib -lHYPRE
```

This driver also support spatial coarsening and explicit time stepping. This allows you to use explicit time stepping on each Braid level, regardless of time step size.

2. drive-03 implements the 3D heat equation on a regular grid, and assumes `hypre` is installed just like drive-02. This driver does not support spatial coarsening, and thus if explicit time stepping is used, the time stepping switches to implicit on coarse XBraid grids when the CFL condition is violated.
3. drive-04 is a sophisticated test bed for finite element discretizations of the heat equation. It relies on the `mfem` package to create general finite element discretizations for the spatial problem. Other packages must be installed in this order.

- Unpack and install `Metis`
- Unpack and install `hypre`
- Unpack `mfem`. Then make sure to set these variables correctly in the `mfem` Makefile:

```
USE_METIS_5 = YES
HYPRE_DIR = where_ever_linear_solvers_is/hypre
```

- Make the parallel version of `mfem` first by typing

```
make parallel
```

- Make `GLVIS`. Set these variables in the `glvis` makefile

```
MFEM_DIR = mfem_location
MFEM_LIB = -L$(MFEM_DIR) -lmfem
```

- Go to `braid/examples` and set these Makefile variables,

```
METIS_DIR = ../../metis-5.1.0/lib
MFEM_DIR = ../../mfem
MFEM_FLAGS = -I$(MFEM_DIR)
MFEM_LIB = -L$(MFEM_DIR) -lmfem -L$(METIS_DIR) -lmetis
```

then type

```
make drive-04
```

- To run drive-04 and glvis, open two windows. In one, start a glvis session

```
./glvis
```

Then, in the other window, run drive-04

```
mpirun -np ... drive-04 [args]
```

Glvis will listen on a port to which drive-04 will dump visualization information.

4. drive-05, 6, 7 and 10 are additional mfem examples compiled like drive-04.
 - drive-05 implements advection(-diffusion) with a discontinuous Galerkin discretization.
 - drive-06 implements nonlinear diffusion, i.e., the p -Laplacian.
 - drive-07 is under development.
 - drive-10 is under development.

10 Module Index

10.1 Modules

Here is a list of all modules:

User-written routines	31
User interface routines	35
General Interface routines	36
XBraid status routines	46
XBraid test routines	58

11 Data Structure Index

11.1 Data Structures

Here are the data structures with brief descriptions:

_braid_AccessStatus	63
_braid_CoarsenRefStatus	64
_braid_CommHandle	65
_braid_Core	66
_braid_Grid	71
_braid_StepStatus	73

12 File Index

12.1 File List

Here is a list of all files with brief descriptions:

_braid.h	Define headers for developer routines	75
braid.h	Define headers for user interface routines	82
braid_defs.h	Definitions of types, error flags, etc..	84
braid_status.h	Define headers for XBraid status structures, status get/set routines and status create/destroy routines	86
braid_test.h	Define headers for XBraid test routines	87

13 Module Documentation

13.1 User-written routines

Typedefs

- typedef struct [_braid_App_struct](#) * [braid_App](#)
- typedef struct [_braid_Vector_struct](#) * [braid_Vector](#)
- typedef [braid_Int](#)(* [braid_PtFcnStep](#))(braid_App app, [braid_Vector](#) ustop, [braid_Vector](#) fstop, [braid_Vector](#) u, [braid_StepStatus](#) status)
- typedef [braid_Int](#)(* [braid_PtFcnInit](#))(braid_App app, [braid_Real](#) t, [braid_Vector](#) *u_ptr)
- typedef [braid_Int](#)(* [braid_PtFcnClone](#))(braid_App app, [braid_Vector](#) u, [braid_Vector](#) *v_ptr)
- typedef [braid_Int](#)(* [braid_PtFcnFree](#))(braid_App app, [braid_Vector](#) u)
- typedef [braid_Int](#)(* [braid_PtFcnSum](#))(braid_App app, [braid_Real](#) alpha, [braid_Vector](#) x, [braid_Real](#) beta, [braid_Vector](#) y)
- typedef [braid_Int](#)(* [braid_PtFcnSpatialNorm](#))(braid_App app, [braid_Vector](#) u, [braid_Real](#) *norm_ptr)
- typedef [braid_Int](#)(* [braid_PtFcnAccess](#))(braid_App app, [braid_Vector](#) u, [braid_AccessStatus](#) status)
- typedef [braid_Int](#)(* [braid_PtFcnBufSize](#))(braid_App app, [braid_Int](#) *size_ptr)
- typedef [braid_Int](#)(* [braid_PtFcnBufPack](#))(braid_App app, [braid_Vector](#) u, void *buffer, [braid_Int](#) *size_ptr)
- typedef [braid_Int](#)(* [braid_PtFcnBufUnpack](#))(braid_App app, void *buffer, [braid_Vector](#) *u_ptr)
- typedef [braid_Int](#)(* [braid_PtFcnResidual](#))(braid_App app, [braid_Vector](#) ustop, [braid_Vector](#) r, [braid_StepStatus](#) status)
- typedef [braid_Int](#)(* [braid_PtFcnSCoarsen](#))(braid_App app, [braid_Vector](#) fu, [braid_Vector](#) *cu_ptr, [braid_CoarsenRefStatus](#) status)
- typedef [braid_Int](#)(* [braid_PtFcnSRefine](#))(braid_App app, [braid_Vector](#) cu, [braid_Vector](#) *fu_ptr, [braid_CoarsenRefStatus](#) status)

13.1.1 Detailed Description

These are all user-written data structures and routines. There are two data structures ([braid_App](#) and [braid_Vector](#)) for the user to define. And, there are a variety of function interfaces (defined through function pointer declarations) that the user must implement.

13.1.2 Typedef Documentation

13.1.2.1 typedef struct [_braid_App_struct](#)* [braid_App](#)

This holds a wide variety of information and is `global` in that it is passed to every function. This structure holds everything that the user will need to carry out a simulation. For a simple example, this could just hold the global MPI communicator and a few values describing the temporal domain.

13.1.2.2 typedef [braid_Int](#)(* [braid_PtFcnAccess](#))([braid_App](#) app,[braid_Vector](#) u,[braid_AccessStatus](#) status)

Gives user access to XBraid and to the current vector u at time t . Most commonly, this lets the user write the vector to screen, file, etc... The user decides what is appropriate. Note how you are told the time value t of the vector u and other information in *status*. This lets you tailor the output, e.g., for only certain time values at certain XBraid iterations. Querying status for such information is done through *[braid_AccessStatusGet](#)* routines.

The frequency of XBraid's calls to *access* is controlled through [braid_SetAccessLevel](#). For instance, if *access_level* is set to 2, then *access* is called every XBraid iteration and on every XBraid level. In this case, querying *status* to determine the current XBraid level and iteration will be useful. This scenario allows for even more detailed tracking of the simulation.

Eventually, access will be broadened to allow the user to steer XBraid.

13.1.2.3 typedef [braid_Int](#)(* [braid_PtFcnBufPack](#))([braid_App](#) app,[braid_Vector](#) u,void *buffer,[braid_Int](#) *size_ptr)

This allows XBraid to send messages containing [braid_Vectors](#). This routine packs a vector u into a *void * buffer* for MPI.

13.1.2.4 typedef [braid_Int](#)(* [braid_PtFcnBufSize](#))([braid_App](#) app,[braid_Int](#) *size_ptr)

This routine tells XBraid message sizes by computing an upper bound in bytes for an arbitrary [braid_Vector](#). This size must be an upper bound for what *BufPack* and *BufUnPack* will assume.

13.1.2.5 typedef [braid_Int](#)(* [braid_PtFcnBufUnpack](#))([braid_App](#) app,void *buffer,[braid_Vector](#) *u_ptr)

This allows XBraid to receive messages containing [braid_Vectors](#). This routine unpacks a *void * buffer* from MPI into a [braid_Vector](#).

13.1.2.6 typedef [braid_Int](#)(* [braid_PtFcnClone](#))([braid_App](#) app,[braid_Vector](#) u,[braid_Vector](#) *v_ptr)

Clone u into v_ptr

13.1.2.7 typedef [braid_Int](#)(* [braid_PtFcnFree](#))([braid_App](#) app,[braid_Vector](#) u)

Free and deallocate u

13.1.2.8 typedef [braid_Int](#)(* [braid_PtFcnInit](#))([braid_App](#) app,[braid_Real](#) t,[braid_Vector](#) *u_ptr)

Initializes a vector u_ptr at time t

13.1.2.9 **typedef braid_Int>(* braid_PtFcnResidual)(braid_App app,braid_Vector ustop,braid_Vector r,braid_StepStatus status)**

This function (optional) computes the residual r at time $tstop$. On input, r holds the value of u at $tstart$, and $ustop$ is the value of u at $tstop$. If used, set with [braid_SetResidual](#).

Query the status structure with [braid_StepStatusGetTstart\(status, &tstart\)](#) and [braid_StepStatusGetTstop\(status, &tstop\)](#) to get $tstart$ and $tstop$.

13.1.2.10 **typedef braid_Int>(* braid_PtFcnSCoarsen)(braid_App app,braid_Vector fu,braid_Vector *cu_ptr,braid_CoarsenRefStatus status)**

Spatial coarsening (optional). Allows the user to coarsen when going from a fine time grid to a coarse time grid. This function is called on every vector at each level, thus you can coarsen the entire space time domain. The action of this function should match the [braid_PtFcnSRefine](#) function.

The user should query the status structure at run time with [braid_CoarsenRefGet**\(\)](#) calls in order to determine how to coarsen. For instance, status tells you what the current time value is, and what the time step sizes on the fine and coarse levels are.

13.1.2.11 **typedef braid_Int>(* braid_PtFcnSpatialNorm)(braid_App app,braid_Vector u,braid_Real *norm_ptr)**

Carry out a spatial norm by taking the norm of a [braid_Vector](#) $norm_ptr = || u ||$. A common choice is the standard Euclidean norm, but many other choices are possible, such as an L2-norm based on a finite element space. See [braid_SetTemporalNorm](#) for information on how the spatial norm is combined over time for a global space-time residual norm. This global norm then controls halting.

13.1.2.12 **typedef braid_Int>(* braid_PtFcnSRefine)(braid_App app,braid_Vector cu,braid_Vector *fu_ptr,braid_CoarsenRefStatus status)**

Spatial refinement (optional). Allows the user to refine when going from a coarse time grid to a fine time grid. This function is called on every vector at each level, thus you can refine the entire space time domain. The action of this function should match the [braid_PtFcnSCoarsen](#) function.

The user should query the status structure at run time with [braid_CoarsenRefGet**\(\)](#) calls in order to determine how to coarsen. For instance, status tells you what the current time value is, and what the time step sizes on the fine and coarse levels are.

13.1.2.13 **typedef braid_Int>(* braid_PtFcnStep)(braid_App app,braid_Vector ustop,braid_Vector fstop,braid_Vector u,braid_StepStatus status)**

Defines the central time stepping function that the user must write.

The user must advance the vector u from time $tstart$ to $tstop$. The time step is taken assuming the right-hand-side vector $fstop$ at time $tstop$. The vector $ustop$ may be the same vector as u (in the case where not all unknowns are stored). The vector $fstop$ is set to NULL to indicate a zero right-hand-side.

Query the status structure with [braid_StepStatusGetTstart\(status, &tstart\)](#) and [braid_StepStatusGetTstop\(status, &tstop\)](#) to get $tstart$ and $tstop$. The status structure also allows for steering. For example, [braid_StepStatusSetR-Factor\(...\)](#) allows for setting a refinement factor, which tells XBraid to refine this time interval.

13.1.2.14 **typedef braid_Int>(* braid_PtFcnSum)(braid_App app,braid_Real alpha,braid_Vector x,braid_Real beta,braid_Vector y)**

AXPY, $alpha x + beta y \rightarrow y$

13.1.2.15 `typedef struct _braid_Vector_struct* braid_Vector`

This defines (roughly) a state vector at a certain time value. It could also contain any other information related to this vector which is needed to evolve the vector to the next time value, like mesh information.

13.2 User interface routines

Modules

- [General Interface routines](#)
- [XBraid status routines](#)

13.2.1 Detailed Description

These are all the user interface routines.

13.3 General Interface routines

Typedefs

- typedef struct `_braid_Core_struct` * `braid_Core`

Functions

- `braid_Int braid_Init` (`MPI_Comm comm_world`, `MPI_Comm comm`, `braid_Real tstart`, `braid_Real tstop`, `braid_Int ntime`, `braid_App app`, `braid_PtFcnStep step`, `braid_PtFcnInit init`, `braid_PtFcnClone clone`, `braid_PtFcnFree free`, `braid_PtFcnSum sum`, `braid_PtFcnSpatialNorm spatialnorm`, `braid_PtFcnAccess access`, `braid_PtFcnBufSize bufsize`, `braid_PtFcnBufPack bufpack`, `braid_PtFcnBufUnpack bufunpack`, `braid_Core *core_ptr`)
- `braid_Int braid_Drive` (`braid_Core core`)
- `braid_Int braid_Destroy` (`braid_Core core`)
- `braid_Int braid_PrintStats` (`braid_Core core`)
- `braid_Int braid_SetMaxLevels` (`braid_Core core`, `braid_Int max_levels`)
- `braid_Int braid_SetSkip` (`braid_Core core`, `braid_Int skip`)
- `braid_Int braid_SetRefine` (`braid_Core core`, `braid_Int refine`)
- `braid_Int braid_SetMaxRefinements` (`braid_Core core`, `braid_Int max_refinements`)
- `braid_Int braid_SetTPointsCutoff` (`braid_Core core`, `braid_Int tpoints_cutoff`)
- `braid_Int braid_SetMinCoarse` (`braid_Core core`, `braid_Int min_coarse`)
- `braid_Int braid_SetAbsTol` (`braid_Core core`, `braid_Real atol`)
- `braid_Int braid_SetRelTol` (`braid_Core core`, `braid_Real rtol`)
- `braid_Int braid_SetNRelax` (`braid_Core core`, `braid_Int level`, `braid_Int nrelax`)
- `braid_Int braid_SetCFactor` (`braid_Core core`, `braid_Int level`, `braid_Int cfactor`)
- `braid_Int braid_SetMaxIter` (`braid_Core core`, `braid_Int max_iter`)
- `braid_Int braid_SetFMG` (`braid_Core core`)
- `braid_Int braid_SetNFMG` (`braid_Core core`, `braid_Int k`)
- `braid_Int braid_SetNFMGVcyc` (`braid_Core core`, `braid_Int nfmng_Vcyc`)
- `braid_Int braid_SetStorage` (`braid_Core core`, `braid_Int storage`)
- `braid_Int braid_SetTemporalNorm` (`braid_Core core`, `braid_Int tnorm`)
- `braid_Int braid_SetResidual` (`braid_Core core`, `braid_PtFcnResidual residual`)
- `braid_Int braid_SetFullRNormRes` (`braid_Core core`, `braid_PtFcnResidual residual`)
- `braid_Int braid_SetSpatialCoarsen` (`braid_Core core`, `braid_PtFcnSCoarsen scoarsen`)
- `braid_Int braid_SetSpatialRefine` (`braid_Core core`, `braid_PtFcnSRefine srefine`)
- `braid_Int braid_SetPrintLevel` (`braid_Core core`, `braid_Int print_level`)
- `braid_Int braid_SetPrintFile` (`braid_Core core`, `const char *printfile_name`)
- `braid_Int braid_SetDefaultPrintFile` (`braid_Core core`)
- `braid_Int braid_SetAccessLevel` (`braid_Core core`, `braid_Int access_level`)
- `braid_Int braid_SplitCommworld` (`const MPI_Comm *comm_world`, `braid_Int px`, `MPI_Comm *comm_x`, `MPI_Comm *comm_t`)
- `braid_Int braid_GetNumIter` (`braid_Core core`, `braid_Int *niter_ptr`)
- `braid_Int braid_GetRNorms` (`braid_Core core`, `braid_Int *nrequest_ptr`, `braid_Real *rnorms`)
- `braid_Int braid_GetNLevels` (`braid_Core core`, `braid_Int *nlevels_ptr`)
- `braid_Int braid_GetSpatialAccuracy` (`braid_StepStatus status`, `braid_Real loose_tol`, `braid_Real tight_tol`, `braid_Real *tol_ptr`)
- `braid_Int braid_SetSeqSoln` (`braid_Core core`, `braid_Int seq_soln`)

13.3.1 Detailed Description

These are general interface routines, e.g., routines to initialize and run a XBraid solver, or to split a communicator into spatial and temporal components.

13.3.2 Typedef Documentation

13.3.2.1 typedef struct `_braid_Core_struct`* `braid_Core`

points to the core structure defined in [_braid.h](#)

13.3.3 Function Documentation

13.3.3.1 `braid_Int braid_Destroy (braid_Core core)`

Clean up and destroy core.

Parameters

<i>core</i>	<code>braid_Core</code> (_braid_Core) struct
-------------	--

13.3.3.2 `braid_Int braid_Drive (braid_Core core)`

Carry out a simulation with XBraid. Integrate in time.

Parameters

<i>core</i>	<code>braid_Core</code> (_braid_Core) struct
-------------	--

13.3.3.3 `braid_Int braid_GetNLevels (braid_Core core, braid_Int * nlevels_ptr)`

After Drive() finishes, this returns the number of XBraid levels

Parameters

<i>core</i>	<code>braid_Core</code> (_braid_Core) struct
<i>nlevels_ptr</i>	output, holds the number of XBraid levels

13.3.3.4 `braid_Int braid_GetNumIter (braid_Core core, braid_Int * niter_ptr)`

After Drive() finishes, this returns the number of iterations taken.

Parameters

<i>core</i>	<code>braid_Core</code> (_braid_Core) struct
<i>niter_ptr</i>	output, holds number of iterations taken

13.3.3.5 `braid_Int braid_GetRNorms (braid_Core core, braid_Int * nrequest_ptr, braid_Real * rnorms)`

After Drive() finishes, this returns XBraid residual history. If *nrequest_ptr* is negative, return the last *nrequest_ptr* residual norms. If positive, return the first *nrequest_ptr* residual norms. Upon exit, *nrequest_ptr* holds the number of residuals actually returned.

Parameters

<i>core</i>	braid_Core (<i>_braid_Core</i>) struct
<i>nrequest_ptr</i>	input/output, input: num requested resid norms, output: num actually returned
<i>rnorms</i>	output, holds residual norm history array

13.3.3.6 `braid_Int braid_GetSpatialAccuracy (braid_StepStatus status, braid_Real loose_tol, braid_Real tight_tol, braid_Real * tol_ptr)`

Example function to compute a tapered stopping tolerance for implicit time stepping routines, i.e., a tolerance *tol_ptr* for the spatial solves. This tapering only occurs on the fine grid.

This rule must be followed. The same tolerance must be returned over all processors, for a given XBraid and XBraid level. Different levels may have different tolerances and the same level may vary its tolerance from iteration to iteration, but for the same iteration and level, the tolerance must be constant.

This additional rule must be followed. The fine grid tolerance is never reduced (this is important for convergence)

On the fine level, the spatial stopping tolerance *tol_ptr* is interpolated from *loose_tol* to *tight_tol* based on the relationship between $rnorm / rnorm0$ and *tol*. Remember when $rnorm / rnorm0 < tol$, XBraid halts. Thus, this function lets us have a loose stopping tolerance while the Braid residual is still relatively large, and then we transition to a tight stopping tolerance as the Braid residual is reduced.

If the user has not defined a residual function, *tight_tol* is always returned.

The *loose_tol* is always used on coarse grids, excepting the above mentioned residual computations.

This function will normally be called from the user's step routine.

This function is also meant as a guide for users to develop their own routine.

Parameters

<i>status</i>	Current XBraid step status
<i>loose_tol</i>	Loosest allowed spatial solve stopping tol on fine grid
<i>tight_tol</i>	Tightest allowed spatial solve stopping tol on fine grid
<i>tol_ptr</i>	output, holds the computed spatial solve stopping tol

13.3.3.7 `braid_Int braid_Init (MPI_Comm comm_world, MPI_Comm comm, braid_Real tstart, braid_Real tstop, braid_Int ntime, braid_App app, braid_PtFcnStep step, braid_PtFcnInit init, braid_PtFcnClone clone, braid_PtFcnFree free, braid_PtFcnSum sum, braid_PtFcnSpatialNorm spatialnorm, braid_PtFcnAccess access, braid_PtFcnBufSize bufsize, braid_PtFcnBufPack bufpack, braid_PtFcnBufUnpack bufunpack, braid_Core * core_ptr)`

Create a core object with the required initial data.

This core is used by XBraid for internal data structures. The output is *core_ptr* which points to the newly created *braid_Core* structure.

Parameters

<i>comm_world</i>	Global communicator for space and time
<i>comm</i>	Communicator for temporal dimension
<i>tstart</i>	start time
<i>tstop</i>	End time

<i>ntime</i>	Initial number of temporal grid values
<i>app</i>	User-defined <code>_braid_App</code> structure
<i>step</i>	User time stepping routine to advance a <code>braid_Vector</code> forward one step
<i>init</i>	Initialize a <code>braid_Vector</code> on the finest temporal grid
<i>clone</i>	Clone a <code>braid_Vector</code>
<i>free</i>	Free a <code>braid_Vector</code>
<i>sum</i>	Compute vector sum of two <code>braid_Vectors</code>
<i>spatialnorm</i>	Compute norm of a <code>braid_Vector</code> , this is a norm only over space
<i>access</i>	Allows access to XBraid and current <code>braid_Vector</code>
<i>bufsize</i>	Computes size for MPI buffer for one <code>braid_Vector</code>
<i>bufpack</i>	Packs MPI buffer to contain one <code>braid_Vector</code>
<i>bufunpack</i>	Unpacks MPI buffer into a <code>braid_Vector</code>
<i>core_ptr</i>	Pointer to <code>braid_Core</code> (<code>_braid_Core</code>) struct

13.3.3.8 `braid_Int braid_PrintStats (braid_Core core)`

Print statistics after a XBraid run.

Parameters

<i>core</i>	<code>braid_Core</code> (<code>_braid_Core</code>) struct
-------------	---

13.3.3.9 `braid_Int braid_SetAbsTol (braid_Core core, braid_Real atol)`

Set absolute stopping tolerance.

Recommended option over relative tolerance

Parameters

<i>core</i>	<code>braid_Core</code> (<code>_braid_Core</code>) struct
<i>atol</i>	absolute stopping tolerance

13.3.3.10 `braid_Int braid_SetAccessLevel (braid_Core core, braid_Int access_level)`

Set access level for XBraid. This controls how often the user's access routine is called.

- Level 0: Never call the user's access routine
- Level 1: Only call the user's access routine after XBraid is finished
- Level 2: Call the user's access routine every iteration and on every level. This is during `_braid_FR restrict`, during the down-cycle part of a XBraid iteration.

Default is level 1.

Parameters

<i>core</i>	<code>braid_Core</code> (<code>_braid_Core</code>) struct
<i>access_level</i>	desired <code>access_level</code>

13.3.3.11 `braid_Int braid_SetCFactor (braid_Core core, braid_Int level, braid_Int cfactor)`

Set the coarsening factor `cfactor` on grid `level` (level 0 is the finest grid). The default factor is 2 on all levels. To change the default factor, use `level = -1`.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>level</i>	<i>level</i> to set coarsening factor on
<i>cfactor</i>	desired coarsening factor

13.3.3.12 **braid_Int** braid_SetDefaultPrintFile (**braid_Core** *core*)

Use default filename, *braid_runtime.out* for runtime print messages. This function is particularly useful for Fortran codes, where passing filename strings between C and Fortran is troublesome. Level of printing is controlled by [braid_SetPrintLevel](#).

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
-------------	---

13.3.3.13 **braid_Int** braid_SetFMG (**braid_Core** *core*)

Once called, XBraid will use FMG (i.e., F-cycles).

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
-------------	---

13.3.3.14 **braid_Int** braid_SetFullRNormRes (**braid_Core** *core*, **braid_PtFcnResidual** *residual*)

Set user-defined residual routine for computing full residual norm (all C/F points).

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>residual</i>	function pointer to residual routine

13.3.3.15 **braid_Int** braid_SetMaxIter (**braid_Core** *core*, **braid_Int** *max_iter*)

Set max number of multigrid iterations.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>max_iter</i>	maximum iterations to allow

13.3.3.16 **braid_Int** braid_SetMaxLevels (**braid_Core** *core*, **braid_Int** *max_levels*)

Set max number of multigrid levels.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>max_levels</i>	maximum levels to allow

13.3.3.17 **braid_Int** braid_SetMaxRefinements (**braid_Core** *core*, **braid_Int** *max_refinements*)

Set the max number of time grid refinement levels allowed.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>max_refinements</i>	maximum refinement levels allowed

13.3.3.18 **braid_Int** `braid_SetMinCoarse (braid_Core core, braid_Int min_coarse)`

Set minimum allowed coarse grid size. XBraid stops coarsening whenever creating the next coarser grid will result in a grid smaller than `min_coarse`. The maximum possible coarse grid size will be `min_coarse*coarsening_factor`.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>min_coarse</i>	minimum coarse grid size

13.3.3.19 **braid_Int** `braid_SetNFMG (braid_Core core, braid_Int k)`

Once called, XBraid will use FMG (i.e., F-cycles).

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>k</i>	number of initial F-cycles to do before switching to V-cycles

13.3.3.20 **braid_Int** `braid_SetNFMGVcyc (braid_Core core, braid_Int nfmg_Vcyc)`

Set number of V-cycles to use at each FMG level (standard is 1)

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>nfmg_Vcyc</i>	number of V-cycles to do each FMG level

13.3.3.21 **braid_Int** `braid_SetNRelax (braid_Core core, braid_Int level, braid_Int nrelax)`

Set the number of relaxation sweeps `nrelax` on grid `level` (level 0 is the finest grid). The default is 1 on all levels. To change the default factor, use `level = -1`. One sweep is a CF relaxation sweep.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>level</i>	<code>level</code> to set <code>nrelax</code> on
<i>nrelax</i>	number of relaxations to do on <code>level</code>

13.3.3.22 **braid_Int** `braid_SetPrintFile (braid_Core core, const char * printfile_name)`

Set output file for runtime print messages. Level of printing is controlled by `braid_SetPrintLevel`. Default is stdout.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>printfile_name</i>	output file for XBraid runtime output

13.3.3.23 **braid_Int** `braid_SetPrintLevel (braid_Core core, braid_Int print_level)`

Set print level for XBraid. This controls how much information is printed to the XBraid print file (`braid_SetPrintFile`).

- Level 0: no output
- Level 1: print typical information like a residual history, number of levels in the XBraid hierarchy, and so on.
- Level 2: level 1 output, plus debug level output.

Default is level 1.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>print_level</i>	desired print level

13.3.3.24 `braid_Int braid_SetRefine (braid_Core core, braid_Int refine)`

Turn time refinement on (`refine = 1`) or off (`refine = 0`).

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>refine</i>	boolean, refine in time or not

13.3.3.25 `braid_Int braid_SetRelTol (braid_Core core, braid_Real rtol)`

Set relative stopping tolerance, relative to the initial residual. Be careful. If your initial guess is all zero, then the initial residual may only be nonzero over one or two time values, and this will skew the relative tolerance. Absolute tolerances are recommended.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>rtol</i>	relative stopping tolerance

13.3.3.26 `braid_Int braid_SetResidual (braid_Core core, braid_PtFcnResidual residual)`

Set user-defined residual routine.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>residual</i>	function pointer to residual routine

13.3.3.27 `braid_Int braid_SetSeqSoln (braid_Core core, braid_Int seq_soln)`

Set the initial guess to XBraid as the sequential time stepping solution. This is primarily for debugging. When used with `storage=-2`, the initial residual should evaluate to exactly 0. The residual can also be 0 for other storage options if the time stepping is *exact*, e.g., the implicit solve in Step is done to full precision.

The value `seq_soln` is a Boolean

- 0: The user's `Init()` function initializes the state vector (default)
- 1: Sequential time stepping, with the user's initial condition from `Init(t=0)` initializes the state vector

Default is 0.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>seq_soln</i>	1: Init with sequential time stepping soln, 0: Use user's Init()

13.3.3.28 **braid_Int** braid_SetSkip (**braid_Core** *core*, **braid_Int** *skip*)

Set whether to skip all work on the first down cycle (skip = 1). On by default.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>skip</i>	boolean, whether to skip all work on first down-cycle

13.3.3.29 **braid_Int** braid_SetSpatialCoarsen (**braid_Core** *core*, **braid_PtFcnSCoarsen** *scoarsen*)

Set spatial coarsening routine with user-defined routine. Default is no spatial refinement or coarsening.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>scoarsen</i>	function pointer to spatial coarsening routine

13.3.3.30 **braid_Int** braid_SetSpatialRefine (**braid_Core** *core*, **braid_PtFcnSRefine** *srefine*)

Set spatial refinement routine with user-defined routine. Default is no spatial refinement or coarsening.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>srefine</i>	function pointer to spatial refinement routine

13.3.3.31 **braid_Int** braid_SetStorage (**braid_Core** *core*, **braid_Int** *storage*)

Sets the storage properties of the code.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>storage</i>	store C-points (0), all points (1)

13.3.3.32 **braid_Int** braid_SetTemporalNorm (**braid_Core** *core*, **braid_Int** *tnorm*)

Sets XBraid temporal norm.

This option determines how to obtain a global space-time residual norm. That is, this decides how to combine the spatial norms returned by [braid_PtFcnSpatialNorm](#) at each time step to obtain a global norm over space and time. It is this global norm that then controls halting.

There are three options for setting *tnorm*. See section [Halting tolerance](#) for a more detailed discussion (in [Introduction.-md](#)).

- *tnorm=1*: One-norm summation of spatial norms
- *tnorm=2*: Two-norm summation of spatial norms
- *tnorm=3*: Infinity-norm combination of spatial norms

The default choice is *tnorm=2*

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>tnorm</i>	choice of temporal norm

13.3.3.33 **braid_Int** braid_SetTPointsCutoff (**braid_Core** *core*, **braid_Int** *tpoints_cutoff*)

Set the number of time steps, beyond with refinements stop. If num(tpoints) > tpoints_cutoff, then stop doing refinements.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>tpoints_cutoff</i>	cutoff for stopping refinements

13.3.3.34 **braid_Int** braid_SplitCommworld (**const MPI_Comm** * *comm_world*, **braid_Int** *px*, **MPI_Comm** * *comm_x*, **MPI_Comm** * *comm_t*)

Split MPI commworld into *comm_x* and *comm_t*, the spatial and temporal communicators. The total number of processors will equal $P_x * P_t$, there P_x is the number of procs in space, and P_t is the number of procs in time.

Parameters

<i>comm_world</i>	Global communicator to split
<i>px</i>	Number of processors parallelizing space for a single time step
<i>comm_x</i>	Spatial communicator (written as output)
<i>comm_t</i>	Temporal communicator (written as output)

13.4 XBraid status routines

Data Structures

- struct [_braid_AccessStatus](#)
- struct [_braid_CoarsenRefStatus](#)
- struct [_braid_StepStatus](#)

Macros

- `#define _braid_StatusElt(status, elt) ((status) -> elt)`

Typedefs

- typedef struct
 [_braid_AccessStatus_struct](#) * [braid_AccessStatus](#)
- typedef struct
 [_braid_CoarsenRefStatus_struct](#) * [braid_CoarsenRefStatus](#)
- typedef struct
 [_braid_StepStatus_struct](#) * [braid_StepStatus](#)

Functions

- [braid_Int_braid_AccessStatusInit](#) ([braid_Real](#) t, [braid_Real](#) rnorm, [braid_Int](#) iter, [braid_Int](#) level, [braid_Int](#) nrefine, [braid_Int](#) gupper, [braid_Int](#) done, [braid_Int](#) wrapper_test, [braid_AccessStatus](#) status)
- [braid_Int_braid_AccessStatusDestroy](#) ([braid_AccessStatus](#) status)
- [braid_Int_braid_AccessStatusGetT](#) ([braid_AccessStatus](#) status, [braid_Real](#) *t_ptr)
- [braid_Int_braid_AccessStatusGetResidual](#) ([braid_AccessStatus](#) status, [braid_Real](#) *rnorm_ptr)
- [braid_Int_braid_AccessStatusGetIter](#) ([braid_AccessStatus](#) status, [braid_Int](#) *iter_ptr)
- [braid_Int_braid_AccessStatusGetLevel](#) ([braid_AccessStatus](#) status, [braid_Int](#) *level_ptr)
- [braid_Int_braid_AccessStatusGetNRefine](#) ([braid_AccessStatus](#) status, [braid_Int](#) *nrefine_ptr)
- [braid_Int_braid_AccessStatusGetNTPoints](#) ([braid_AccessStatus](#) status, [braid_Int](#) *ntpoints_ptr)
- [braid_Int_braid_AccessStatusGetDone](#) ([braid_AccessStatus](#) status, [braid_Int](#) *done_ptr)
- [braid_Int_braid_AccessStatusGetWrapperTest](#) ([braid_AccessStatus](#) status, [braid_Int](#) *wtest_ptr)
- [braid_Int_braid_AccessStatusGetTILD](#) ([braid_AccessStatus](#) status, [braid_Real](#) *t_ptr, [braid_Int](#) *iter_ptr, [braid_Int](#) *level_ptr, [braid_Int](#) *done_ptr)
- [braid_Int_braid_CoarsenRefStatusInit](#) ([braid_Real](#) tstart, [braid_Real](#) f_tprior, [braid_Real](#) f_tstop, [braid_Real](#) c_tprior, [braid_Real](#) c_tstop, [braid_Int](#) level, [braid_Int](#) nrefine, [braid_Int](#) gupper, [braid_CoarsenRefStatus](#) status)
- [braid_Int_braid_CoarsenRefStatusDestroy](#) ([braid_CoarsenRefStatus](#) status)
- [braid_Int_braid_CoarsenRefStatusGetTstart](#) ([braid_CoarsenRefStatus](#) status, [braid_Real](#) *tstart_ptr)
- [braid_Int_braid_CoarsenRefStatusGetFTstop](#) ([braid_CoarsenRefStatus](#) status, [braid_Real](#) *f_tstop_ptr)
- [braid_Int_braid_CoarsenRefStatusGetFTprior](#) ([braid_CoarsenRefStatus](#) status, [braid_Real](#) *f_tprior_ptr)
- [braid_Int_braid_CoarsenRefStatusGetCTstop](#) ([braid_CoarsenRefStatus](#) status, [braid_Real](#) *c_tstop_ptr)
- [braid_Int_braid_CoarsenRefStatusGetCTprior](#) ([braid_CoarsenRefStatus](#) status, [braid_Real](#) *c_tprior_ptr)
- [braid_Int_braid_CoarsenRefStatusGetTpriorTstop](#) ([braid_CoarsenRefStatus](#) status, [braid_Real](#) *tstart_ptr, [braid_Real](#) *f_tprior_ptr, [braid_Real](#) *f_tstop_ptr, [braid_Real](#) *c_tprior_ptr, [braid_Real](#) *c_tstop_ptr)
- [braid_Int_braid_CoarsenRefStatusGetLevel](#) ([braid_CoarsenRefStatus](#) status, [braid_Int](#) *level_ptr)
- [braid_Int_braid_CoarsenRefStatusGetNRefine](#) ([braid_CoarsenRefStatus](#) status, [braid_Int](#) *nrefine_ptr)
- [braid_Int_braid_CoarsenRefStatusGetNTPoints](#) ([braid_CoarsenRefStatus](#) status, [braid_Int](#) *ntpoints_ptr)

- `braid_Int_braid_StepStatusInit` (`braid_Real tstart`, `braid_Real tstop`, `braid_Real tol`, `braid_Int iter`, `braid_Int level`, `braid_Int nrefine`, `braid_Int gupper`, `braid_StepStatus status`)
- `braid_Int_braid_StepStatusDestroy` (`braid_StepStatus status`)
- `braid_Int_braid_StepStatusGetTstart` (`braid_StepStatus status`, `braid_Real *tstart_ptr`)
- `braid_Int_braid_StepStatusGetTstop` (`braid_StepStatus status`, `braid_Real *tstop_ptr`)
- `braid_Int_braid_StepStatusGetLevel` (`braid_StepStatus status`, `braid_Int *level_ptr`)
- `braid_Int_braid_StepStatusGetNRefine` (`braid_StepStatus status`, `braid_Int *nrefine_ptr`)
- `braid_Int_braid_StepStatusGetNTPoints` (`braid_StepStatus status`, `braid_Int *ntpoints_ptr`)
- `braid_Int_braid_StepStatusSetRFactor` (`braid_StepStatus status`, `braid_Real rfactor`)
- `braid_Int_braid_StepStatusGetTstartTstop` (`braid_StepStatus status`, `braid_Real *tstart_ptr`, `braid_Real *tstop_ptr`)
- `braid_Int_braid_StepStatusGetTol` (`braid_StepStatus status`, `braid_Real *tol_ptr`)
- `braid_Int_braid_StepStatusGetIter` (`braid_StepStatus status`, `braid_Int *iter_ptr`)
- `braid_Int_braid_StepStatusGetRNorms` (`braid_StepStatus status`, `braid_Int *nrequest_ptr`, `braid_Real *rnorms`)
- `braid_Int_braid_StepStatusGetOldFineTolx` (`braid_StepStatus status`, `braid_Real *old_fine_tolx_ptr`)
- `braid_Int_braid_StepStatusSetOldFineTolx` (`braid_StepStatus status`, `braid_Real old_fine_tolx_ptr`)
- `braid_Int_braid_StepStatusSetTightFineTolx` (`braid_StepStatus status`, `braid_Int tight_fine_tolx`)

13.4.1 Detailed Description

XBraid status structures are what tell the user the status of the simulation when their routines (step, coarsen/refine, access) are called.

13.4.2 Macro Definition Documentation

13.4.2.1 `#define_braid_StatusElt(status, elt) ((status)->elt)`

Accessor for all `braid**Status` attributes

13.4.3 Typedef Documentation

13.4.3.1 `typedef_struct_braid_AccessStatus_struct* braid_AccessStatus`

The user access routine will receive `braid_AccessStatus`, which will be a pointer to the actual `_braid_AccessStatus_struct`

13.4.3.2 `typedef_struct_braid_CoarsenRefStatus_struct* braid_CoarsenRefStatus`

The user coarsen and refine routines will receive `braid_CoarsenRefStatus`, which will be a pointer to the actual `_braid_CoarsenRefStatus_struct`

13.4.3.3 `typedef_struct_braid_StepStatus_struct* braid_StepStatus`

The user's step routine will receive `braid_StepStatus`, which will be a pointer to the actual `_braid_StepStatus_struct`

13.4.4 Function Documentation

13.4.4.1 `braid_Int_braid_AccessStatusDestroy(braid_AccessStatus status)`

Destroy a `braid_AccessStatus` structure to be destroyed

13.4.4.2 `braid_Int_braid_AccessStatusInit (braid_Real t, braid_Real rnorm, braid_Int iter, braid_Int level, braid_Int nrefine, braid_Int gupper, braid_Int done, braid_Int wrapper_test, braid_AccessStatus status)`

Initialize a `braid_AccessStatus` structure

Parameters

<i>t</i>	current time
<i>rnorm</i>	current residual norm in XBraid
<i>iter</i>	current iteration in XBraid
<i>level</i>	current level in XBraid
<i>nrefine</i>	number of refinements done
<i>gupper</i>	global size of the fine grid
<i>done</i>	boolean describing whether XBraid has finished
<i>wrapper_test</i>	boolean describing whether this call is only a wrapper test
<i>status</i>	structure to initialize

13.4.4.3 **braid_Int_braid_CoarsenRefStatusDestroy (braid_CoarsenRefStatus status)**

Destroy a `braid_CoarsenRefStatus` structure to be destroyed

13.4.4.4 **braid_Int_braid_CoarsenRefStatusInit (braid_Real tstart, braid_Real f_tprior, braid_Real f_tstop, braid_Real c_tprior, braid_Real c_tstop, braid_Int level, braid_Int nrefine, braid_Int gupper, braid_CoarsenRefStatus status)**

Initialize a `braid_CoarsenRefStatus` structure

Parameters

<i>tstart</i>	time value for current vector
<i>f_tprior</i>	time value to the left of tstart on fine grid
<i>f_tstop</i>	time value to the right of tstart on fine grid
<i>c_tprior</i>	time value to the left of tstart on coarse grid
<i>c_tstop</i>	time value to the right of tstart on coarse grid
<i>level</i>	current fine level in XBraid
<i>nrefine</i>	number of refinements done
<i>gupper</i>	global size of the fine grid
<i>status</i>	structure to initialize

13.4.4.5 **braid_Int_braid_StepStatusDestroy (braid_StepStatus status)**

Destroy a `braid_StepStatus` structure to be destroyed

13.4.4.6 **braid_Int_braid_StepStatusInit (braid_Real tstart, braid_Real tstop, braid_Real tol, braid_Int iter, braid_Int level, braid_Int nrefine, braid_Int gupper, braid_StepStatus status)**

Initialize a `braid_StepStatus` structure

Parameters

<i>tstart</i>	current time value
<i>tstop</i>	time value to evolve towards, time value to the right of tstart
<i>tol</i>	Current XBraid stopping tolerance
<i>iter</i>	Current XBraid iteration (also equal to length of rnorms)
<i>level</i>	current level in XBraid

<i>nrefine</i>	number of refinements done
<i>gupper</i>	global size of the fine grid
<i>status</i>	structure to initialize

13.4.4.7 `braid_Int braid_AccessStatusGetDone (braid_AccessStatus status, braid_Int * done_ptr)`

Return whether XBraid is done for the current simulation.

done_ptr = 1 indicates that XBraid has finished iterating, (either maxiter has been reached, or the tolerance has been met).

Parameters

<i>status</i>	structure containing current simulation info
<i>done_ptr</i>	output, =1 if XBraid has finished, else =0

13.4.4.8 `braid_Int braid_AccessStatusGetIter (braid_AccessStatus status, braid_Int * iter_ptr)`

Return the current iteration from the AccessStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>iter_ptr</i>	output, current XBraid iteration number

13.4.4.9 `braid_Int braid_AccessStatusGetLevel (braid_AccessStatus status, braid_Int * level_ptr)`

Return the current XBraid level from the AccessStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>level_ptr</i>	output, current level in XBraid

13.4.4.10 `braid_Int braid_AccessStatusGetNRefine (braid_AccessStatus status, braid_Int * nrefine_ptr)`

Return the number of refinements done.

Parameters

<i>status</i>	structure containing current simulation info
<i>nrefine_ptr</i>	output, number of refinements done

13.4.4.11 `braid_Int braid_AccessStatusGetNTPoints (braid_AccessStatus status, braid_Int * ntpoints_ptr)`

Return the global number of time points on the fine grid.

Parameters

<i>status</i>	structure containing current simulation info
<i>ntpoints_ptr</i>	output, number of time points on the fine grid

13.4.4.12 `braid_Int braid_AccessStatusGetResidual (braid_AccessStatus status, braid_Real * rnorm_ptr)`

Return the current residual norm from the AccessStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>rnorm_ptr</i>	output, current residual norm

13.4.4.13 **braided_Int** braided_AccessStatusGetT (**braided_AccessStatus** *status*, **braided_Real** * *t_ptr*)

Return the current time from the AccessStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>t_ptr</i>	output, current time

13.4.4.14 **braided_Int** braided_AccessStatusGetTILD (**braided_AccessStatus** *status*, **braided_Real** * *t_ptr*, **braided_Int** * *iter_ptr*, **braided_Int** * *level_ptr*, **braided_Int** * *done_ptr*)

Return XBraid status for the current simulation. Four values are returned.

TILD : time, iteration, level, done

These values are also available through individual Get routines. These individual routines are the location of detailed documentation on each parameter, e.g., see *braided_AccessStatusGetDone* for more information on the *done* value.

Parameters

<i>status</i>	structure containing current simulation info
<i>t_ptr</i>	output, current time
<i>iter_ptr</i>	output, current iteration in XBraid
<i>level_ptr</i>	output, current level in XBraid
<i>done_ptr</i>	output, boolean describing whether XBraid has finished

13.4.4.15 **braided_Int** braided_AccessStatusGetWrapperTest (**braided_AccessStatus** *status*, **braided_Int** * *wtest_ptr*)

Return whether this is a wrapper test or an XBraid run

Parameters

<i>status</i>	structure containing current simulation info
<i>wtest_ptr</i>	output, =1 if this is a wrapper test, =0 if XBraid run

13.4.4.16 **braided_Int** braided_CoarsenRefStatusGetCTprior (**braided_CoarsenRefStatus** *status*, **braided_Real** * *c_tprior_ptr*)

Return the **coarse grid** time value to the left of the current time value from the CoarsenRefStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>c_tprior_ptr</i>	output, time value to the left of current time value on coarse grid

13.4.4.17 **braided_Int** braided_CoarsenRefStatusGetCTstop (**braided_CoarsenRefStatus** *status*, **braided_Real** * *c_tstop_ptr*)

Return the **coarse grid** time value to the right of the current time value from the CoarsenRefStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>c_tstop_ptr</i>	output, time value to the right of current time value on coarse grid

13.4.4.18 **braid_Int** **braid_CoarsenRefStatusGetFTprior** (**braid_CoarsenRefStatus** *status*, **braid_Real** * *f_tprior_ptr*)

Return the **fine grid** time value to the left of the current time value from the CoarsenRefStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>f_tprior_ptr</i>	output, time value to the left of current time value on fine grid

13.4.4.19 **braid_Int** **braid_CoarsenRefStatusGetFTstop** (**braid_CoarsenRefStatus** *status*, **braid_Real** * *f_tstop_ptr*)

Return the **fine grid** time value to the right of the current time value from the CoarsenRefStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>f_tstop_ptr</i>	output, time value to the right of current time value on fine grid

13.4.4.20 **braid_Int** **braid_CoarsenRefStatusGetLevel** (**braid_CoarsenRefStatus** *status*, **braid_Int** * *level_ptr*)

Return the current XBraid level from the CoarsenRefStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>level_ptr</i>	output, current fine level in XBraid

13.4.4.21 **braid_Int** **braid_CoarsenRefStatusGetNRefine** (**braid_CoarsenRefStatus** *status*, **braid_Int** * *nrefine_ptr*)

Return the number of refinements done.

Parameters

<i>status</i>	structure containing current simulation info
<i>nrefine_ptr</i>	output, number of refinements done

13.4.4.22 **braid_Int** **braid_CoarsenRefStatusGetNTPoints** (**braid_CoarsenRefStatus** *status*, **braid_Int** * *ntpoints_ptr*)

Return the global number of time points on the fine grid.

Parameters

<i>status</i>	structure containing current simulation info
<i>ntpoints_ptr</i>	output, number of time points on the fine grid

13.4.4.23 **braid_Int** **braid_CoarsenRefStatusGetTpriorTstop** (**braid_CoarsenRefStatus** *status*, **braid_Real** * *tstart_ptr*, **braid_Real** * *f_tprior_ptr*, **braid_Real** * *f_tstop_ptr*, **braid_Real** * *c_tprior_ptr*, **braid_Real** * *c_tstop_ptr*)

Return XBraid status for the current simulation. Five values are returned, *tstart*, *f_tprior*, *f_tstop*, *c_tprior*, *c_tstop*.

These values are also available through individual Get routines. These individual routines are the location of detailed documentation on each parameter, e.g., see *braid_CoarsenRefStatusGetCTprior* for more information on the *c_tprior*

value.

Parameters

<i>status</i>	structure containing current simulation info
<i>tstart_ptr</i>	output, time value current vector
<i>f_tprior_ptr</i>	output, time value to the left of tstart on fine grid
<i>f_tstop_ptr</i>	output, time value to the right of tstart on fine grid
<i>c_tprior_ptr</i>	output, time value to the left of tstart on coarse grid
<i>c_tstop_ptr</i>	output, time value to the right of tstart on coarse grid

13.4.4.24 **braid_Int** braid_CoarsenRefStatusGetTstart (**braid_CoarsenRefStatus** *status*, **braid_Real** * *tstart_ptr*)

Return the current time value from the CoarsenRefStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>tstart_ptr</i>	output, current time

13.4.4.25 **braid_Int** braid_StepStatusGetIter (**braid_StepStatus** *status*, **braid_Int** * *iter_ptr*)

Return the current XBraid iteration from the StepStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>iter_ptr</i>	output, current iteration in XBraid

13.4.4.26 **braid_Int** braid_StepStatusGetLevel (**braid_StepStatus** *status*, **braid_Int** * *level_ptr*)

Return the current XBraid level from the StepStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>level_ptr</i>	output, current level in XBraid

13.4.4.27 **braid_Int** braid_StepStatusGetNRefine (**braid_StepStatus** *status*, **braid_Int** * *nrefine_ptr*)

Return the number of refinements done.

Parameters

<i>status</i>	structure containing current simulation info
<i>nrefine_ptr</i>	output, number of refinements done

13.4.4.28 **braid_Int** braid_StepStatusGetNTPoints (**braid_StepStatus** *status*, **braid_Int** * *ntpoints_ptr*)

Return the global number of time points on the fine grid.

Parameters

<i>status</i>	structure containing current simulation info
<i>ntpoints_ptr</i>	output, number of time points on the fine grid

13.4.4.29 `braid_Int braid_StepStatusGetOldFineTolx (braid_StepStatus status, braid_Real * old_fine_tolx_ptr)`

Return the previous `old_fine_tolx` set through `braid_StepStatusSetOldFineTolx` This is used especially by `*braid_Get-SpatialAccuracy`

Parameters

<i>status</i>	structure containing current simulation info
<i>old_fine_tolx_ptr</i>	output, previous <i>old_fine_tolx</i> , set through <i>braid_StepStatusSetOldFineTolx</i>

13.4.4.30 braid_Int braid_StepStatusGetRNorms (braid_StepStatus status, braid_Int * nrequest_ptr, braid_Real * rnorms)

Return the current XBraid residual history. If *nrequest_ptr* is negative, return the last *nrequest_ptr* residual norms. If positive, return the first *nrequest_ptr* residual norms. Upon exit, *nrequest_ptr* holds the number of residuals actually returned.

Parameters

<i>status</i>	structure containing current simulation info
<i>nrequest_ptr</i>	input/output, input: number of requested residual norms, output: number actually copied
<i>rnorms</i>	output, XBraid residual norm history, of length <i>nrequest_ptr</i>

13.4.4.31 braid_Int braid_StepStatusGetTol (braid_StepStatus status, braid_Real * tol_ptr)

Return the current XBraid stopping tolerance

Parameters

<i>status</i>	structure containing current simulation info
<i>tol_ptr</i>	output, current XBraid stopping tolerance

13.4.4.32 braid_Int braid_StepStatusGetTstart (braid_StepStatus status, braid_Real * tstart_ptr)

Return the current time value from the StepStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>tstart_ptr</i>	output, current time

13.4.4.33 braid_Int braid_StepStatusGetTstartTstop (braid_StepStatus status, braid_Real * tstart_ptr, braid_Real * tstop_ptr)

Return XBraid status for the current simulation. Two values are returned, *tstart* and *tstop*.

These values are also available through individual Get routines. These individual routines are the location of detailed documentation on each parameter, e.g., see *braid_StepStatusGetTstart* for more information on the *tstart* value.

Parameters

<i>status</i>	structure containing current simulation info
<i>tstart_ptr</i>	output, current time
<i>tstop_ptr</i>	output, next time value to evolve towards

13.4.4.34 braid_Int braid_StepStatusGetTstop (braid_StepStatus status, braid_Real * tstop_ptr)

Return the time value to the right of the current time value from the StepStatus structure.

Parameters

<i>status</i>	structure containing current simulation info
<i>tstop_ptr</i>	output, next time value to evolve towards

13.4.4.35 **braid_Int braid_StepStatusSetOldFineTolx (braid_StepStatus status, braid_Real old_fine_tolx_ptr)**

Set *old_fine_tolx*, available for retrieval through *braid_StepStatusGetOldFineTolx* This is used especially by *braid_GetSpatialAccuracy

Parameters

<i>status</i>	structure containing current simulation info
<i>old_fine_tolx_ptr</i>	input, the last used fine_tolx

13.4.4.36 **braid_Int braid_StepStatusSetRfactor (braid_StepStatus status, braid_Real rfactor)**

Set the rfactor, a desired refinement factor for this interval. rfactor=1 indicates no refinement, otherwise, this interval is subdivided rfactor times.

Parameters

<i>status</i>	structure containing current simulation info
<i>rfactor</i>	user-determined desired rfactor

13.4.4.37 **braid_Int braid_StepStatusSetTightFineTolx (braid_StepStatus status, braid_Int tight_fine_tolx)**

Set *tight_fine_tolx*, boolean variable indicating whether the tightest tolerance has been used for spatial solves (implicit schemes). This value must be 1 in order for XBraid to halt (unless maxiter is reached)

Parameters

<i>status</i>	structure containing current simulation info
<i>tight_fine_tolx</i>	input, boolean indicating whether the tight tolx has been used

13.5 XBraid test routines

Functions

- `braid_Int braid_TestInitAccess` (`braid_App` app, `MPI_Comm` comm_x, `FILE *fp`, `braid_Real` t, `braid_PtFcnInit` init, `braid_PtFcnAccess` access, `braid_PtFcnFree` free)
- `braid_Int braid_TestClone` (`braid_App` app, `MPI_Comm` comm_x, `FILE *fp`, `braid_Real` t, `braid_PtFcnInit` init, `braid_PtFcnAccess` access, `braid_PtFcnFree` free, `braid_PtFcnClone` clone)
- `braid_Int braid_TestSum` (`braid_App` app, `MPI_Comm` comm_x, `FILE *fp`, `braid_Real` t, `braid_PtFcnInit` init, `braid_PtFcnAccess` access, `braid_PtFcnFree` free, `braid_PtFcnClone` clone, `braid_PtFcnSum` sum)
- `braid_Int braid_TestSpatialNorm` (`braid_App` app, `MPI_Comm` comm_x, `FILE *fp`, `braid_Real` t, `braid_PtFcnInit` init, `braid_PtFcnFree` free, `braid_PtFcnClone` clone, `braid_PtFcnSum` sum, `braid_PtFcnSpatialNorm` spatialnorm)
- `braid_Int braid_TestBuf` (`braid_App` app, `MPI_Comm` comm_x, `FILE *fp`, `braid_Real` t, `braid_PtFcnInit` init, `braid_PtFcnFree` free, `braid_PtFcnSum` sum, `braid_PtFcnSpatialNorm` spatialnorm, `braid_PtFcnBufSize` bufsize, `braid_PtFcnBufPack` bufpack, `braid_PtFcnBufUnpack` bufunpack)
- `braid_Int braid_TestCoarsenRefine` (`braid_App` app, `MPI_Comm` comm_x, `FILE *fp`, `braid_Real` t, `braid_Real` fdt, `braid_Real` cdt, `braid_PtFcnInit` init, `braid_PtFcnAccess` access, `braid_PtFcnFree` free, `braid_PtFcnClone` clone, `braid_PtFcnSum` sum, `braid_PtFcnSpatialNorm` spatialnorm, `braid_PtFcnSCoarsen` coarsen, `braid_PtFcnSRefine` refine)
- `braid_Int braid_TestAll` (`braid_App` app, `MPI_Comm` comm_x, `FILE *fp`, `braid_Real` t, `braid_Real` fdt, `braid_Real` cdt, `braid_PtFcnInit` init, `braid_PtFcnFree` free, `braid_PtFcnClone` clone, `braid_PtFcnSum` sum, `braid_PtFcnSpatialNorm` spatialnorm, `braid_PtFcnBufSize` bufsize, `braid_PtFcnBufPack` bufpack, `braid_PtFcnBufUnpack` bufunpack, `braid_PtFcnSCoarsen` coarsen, `braid_PtFcnSRefine` refine)

13.5.1 Detailed Description

These are sanity check routines to help a user test their XBraid code.

13.5.2 Function Documentation

13.5.2.1 `braid_Int braid_TestAll` (`braid_App` app, `MPI_Comm` comm_x, `FILE *fp`, `braid_Real` t, `braid_Real` fdt, `braid_Real` cdt, `braid_PtFcnInit` init, `braid_PtFcnFree` free, `braid_PtFcnClone` clone, `braid_PtFcnSum` sum, `braid_PtFcnSpatialNorm` spatialnorm, `braid_PtFcnBufSize` bufsize, `braid_PtFcnBufPack` bufpack, `braid_PtFcnBufUnpack` bufunpack, `braid_PtFcnSCoarsen` coarsen, `braid_PtFcnSRefine` refine)

Runs all of the individual `braid_Test*` routines

- Returns 0 if the tests fail
- Returns 1 if the tests pass
- Check the log messages to see details of which tests failed.

Parameters

<i>app</i>	User defined App structure
<i>comm_x</i>	Spatial communicator
<i>fp</i>	File pointer (could be stdout or stderr) for log messages

<i>t</i>	Time value to initialize test vectors with
<i>fdt</i>	Fine time step value that you spatially coarsen from
<i>cdt</i>	Coarse time step value that you coarsen to
<i>init</i>	Initialize a <code>braid_Vector</code> on finest temporal grid
<i>free</i>	Free a <code>braid_Vector</code>
<i>clone</i>	Clone a <code>braid_Vector</code>
<i>sum</i>	Compute vector sum of two <code>braid_Vectors</code>
<i>spatialnorm</i>	Compute norm of a <code>braid_Vector</code> , this is a norm only over space
<i>bufsize</i>	Computes size in bytes for one <code>braid_Vector</code> MPI buffer
<i>bufpack</i>	Packs MPI buffer to contain one <code>braid_Vector</code>
<i>bufunpack</i>	Unpacks MPI buffer into a <code>braid_Vector</code>
<i>coarsen</i>	Spatially coarsen a vector. If NULL, test is skipped.
<i>refine</i>	Spatially refine a vector. If NULL, test is skipped.

13.5.2.2 `braid_Int braid_TestBuf (braid_App app, MPI_Comm comm_x, FILE * fp, braid_Real t, braid_PtFcnInit init, braid_PtFcnFree free, braid_PtFcnSum sum, braid_PtFcnSpatialNorm spatialnorm, braid_PtFcnBufSize bufsize, braid_PtFcnBufPack bufpack, braid_PtFcnBufUnpack bufunpack)`

Test the `BufPack`, `BufUnpack` and `BufSize` functions.

A vector is initialized at time t , packed into a buffer, then unpacked from a buffer. The unpacked result must equal the original vector.

- Returns 0 if the tests fail
- Returns 1 if the tests pass
- Check the log messages to see details of which tests failed.

Parameters

<i>app</i>	User defined App structure
<i>comm_x</i>	Spatial communicator
<i>fp</i>	File pointer (could be stdout or stderr) for log messages
<i>t</i>	Time value to test Buffer routines (used to initialize the vectors)
<i>init</i>	Initialize a <code>braid_Vector</code> on finest temporal grid
<i>free</i>	Free a <code>braid_Vector</code>
<i>sum</i>	Compute vector sum of two <code>braid_Vectors</code>
<i>spatialnorm</i>	Compute norm of a <code>braid_Vector</code> , this is a norm only over space
<i>bufsize</i>	Computes size in bytes for one <code>braid_Vector</code> MPI buffer
<i>bufpack</i>	Packs MPI buffer to contain one <code>braid_Vector</code>
<i>bufunpack</i>	Unpacks MPI buffer containing one <code>braid_Vector</code>

13.5.2.3 `braid_Int braid_TestClone (braid_App app, MPI_Comm comm_x, FILE * fp, braid_Real t, braid_PtFcnInit init, braid_PtFcnAccess access, braid_PtFcnFree free, braid_PtFcnClone clone)`

Test the clone function.

A vector is initialized at time t , cloned, and both vectors are written. Then both vectors are free-d. The user is to check (via the access function) to see if it is identical.

Parameters

<i>app</i>	User defined App structure
<i>comm_x</i>	Spatial communicator
<i>fp</i>	File pointer (could be stdout or stderr) for log messages
<i>t</i>	Time value to test clone with (used to initialize the vectors)
<i>init</i>	Initialize a braid_Vector on finest temporal grid
<i>access</i>	Allows access to XBraid and current braid_Vector (can be NULL for no writing)
<i>free</i>	Free a braid_Vector
<i>clone</i>	Clone a braid_Vector

13.5.2.4 **braid_Int** braid_TestCoarsenRefine (**braid_App** *app*, **MPI_Comm** *comm_x*, **FILE *** *fp*, **braid_Real** *t*, **braid_Real** *fdt*, **braid_Real** *cdt*, **braid_PtFcnInit** *init*, **braid_PtFcnAccess** *access*, **braid_PtFcnFree** *free*, **braid_PtFcnClone** *clone*, **braid_PtFcnSum** *sum*, **braid_PtFcnSpatialNorm** *spatialnorm*, **braid_PtFcnSCoarsen** *coarsen*, **braid_PtFcnSRefine** *refine*)

Test the Coarsen and Refine functions.

A vector is initialized at time *t*, and various sanity checks on the spatial coarsening and refinement routines are run.

- Returns 0 if the tests fail
- Returns 1 if the tests pass
- Check the log messages to see details of which tests failed.

Parameters

<i>app</i>	User defined App structure
<i>comm_x</i>	Spatial communicator
<i>fp</i>	File pointer (could be stdout or stderr) for log messages
<i>t</i>	Time value to initialize test vectors
<i>fdt</i>	Fine time step value that you spatially coarsen from
<i>cdt</i>	Coarse time step value that you coarsen to
<i>init</i>	Initialize a braid_Vector on finest temporal grid
<i>access</i>	Allows access to XBraid and current braid_Vector (can be NULL for no writing)
<i>free</i>	Free a braid_Vector
<i>clone</i>	Clone a braid_Vector
<i>sum</i>	Compute vector sum of two braid_Vectors
<i>spatialnorm</i>	Compute norm of a braid_Vector, this is a norm only over space
<i>coarsen</i>	Spatially coarsen a vector
<i>refine</i>	Spatially refine a vector

13.5.2.5 **braid_Int** braid_TestInitAccess (**braid_App** *app*, **MPI_Comm** *comm_x*, **FILE *** *fp*, **braid_Real** *t*, **braid_PtFcnInit** *init*, **braid_PtFcnAccess** *access*, **braid_PtFcnFree** *free*)

Test the init, access and free functions.

A vector is initialized at time *t*, written, and then free-d

Parameters

<i>app</i>	User defined App structure
<i>comm_x</i>	Spatial communicator
<i>fp</i>	File pointer (could be stdout or stderr) for log messages
<i>t</i>	Time value to test init with (used to initialize the vectors)
<i>init</i>	Initialize a <code>braid_Vector</code> on finest temporal grid
<i>access</i>	Allows access to <code>XBraid</code> and current <code>braid_Vector</code> (can be NULL for no writing)
<i>free</i>	Free a <code>braid_Vector</code>

13.5.2.6 `braid_Int braid_TestSpatialNorm (braid_App app, MPI_Comm comm_x, FILE * fp, braid_Real t, braid_PtFcnInit init, braid_PtFcnFree free, braid_PtFcnClone clone, braid_PtFcnSum sum, braid_PtFcnSpatialNorm spatialnorm)`

Test the `spatialnorm` function.

A vector is initialized at time t and then cloned. Various norm evaluations like $\| 3 v \| / \| v \|$ with known output are then done.

- Returns 0 if the tests fail
- Returns 1 if the tests pass
- Check the log messages to see details of which tests failed.

Parameters

<i>app</i>	User defined App structure
<i>comm_x</i>	Spatial communicator
<i>fp</i>	File pointer (could be stdout or stderr) for log messages
<i>t</i>	Time value to test <code>SpatialNorm</code> with (used to initialize the vectors)
<i>init</i>	Initialize a <code>braid_Vector</code> on finest temporal grid
<i>free</i>	Free a <code>braid_Vector</code>
<i>clone</i>	Clone a <code>braid_Vector</code>
<i>sum</i>	Compute vector sum of two <code>braid_Vectors</code>
<i>spatialnorm</i>	Compute norm of a <code>braid_Vector</code> , this is a norm only over space

13.5.2.7 `braid_Int braid_TestSum (braid_App app, MPI_Comm comm_x, FILE * fp, braid_Real t, braid_PtFcnInit init, braid_PtFcnAccess access, braid_PtFcnFree free, braid_PtFcnClone clone, braid_PtFcnSum sum)`

Test the `sum` function.

A vector is initialized at time t , cloned, and then these two vectors are summed a few times, with the results written. The vectors are then free-d. The user is to check (via the `access` function) that the output matches the sum of the two original vectors.

Parameters

<i>app</i>	User defined App structure
<i>comm_x</i>	Spatial communicator
<i>fp</i>	File pointer (could be stdout or stderr) for log messages
<i>t</i>	Time value to test <code>Sum</code> with (used to initialize the vectors)
<i>init</i>	Initialize a <code>braid_Vector</code> on finest temporal grid

<i>access</i>	Allows access to XBraid and current braid_Vector (can be NULL for no writing)
<i>free</i>	Free a braid_Vector
<i>clone</i>	Clone a braid_Vector
<i>sum</i>	Compute vector sum of two braid_Vectors

14 Data Structure Documentation

14.1 `_braid_AccessStatus` Struct Reference

Data Fields

- [braid_Real t](#)
- [braid_Int iter](#)
- [braid_Int level](#)
- [braid_Int nrefine](#)
- [braid_Int gupper](#)
- [braid_Real rnorm](#)
- [braid_Int done](#)
- [braid_Int wrapper_test](#)

14.1.1 Detailed Description

`AccessStatus` structure which defines the status of XBraid at a given instant on some level during a run. The user accesses it through `braid_AccessStatusGet**()` functions.

14.1.2 Field Documentation

14.1.2.1 `braid_Int done`

boolean describing whether XBraid has finished

14.1.2.2 `braid_Int gupper`

global size of the fine grid

14.1.2.3 `braid_Int iter`

XBraid iteration number

14.1.2.4 `braid_Int level`

current level in XBraid

14.1.2.5 `braid_Int nrefine`

number of refinements done

14.1.2.6 `braid_Real rnorm`

residual norm

14.1.2.7 `braid_Real t`

current time

14.1.2.8 `braid_Int wrapper_test`

boolean describing whether this call is only a wrapper test

The documentation for this struct was generated from the following file:

- [braid_status.h](#)

14.2 `_braid_CoarsenRefStatus` Struct Reference

Data Fields

- [braid_Real tstart](#)
- [braid_Real f_tprior](#)
- [braid_Real f_tstop](#)
- [braid_Real c_tprior](#)
- [braid_Real c_tstop](#)
- [braid_Int level](#)
- [braid_Int nrefine](#)
- [braid_Int gupper](#)

14.2.1 Detailed Description

The user coarsen and refine routines will receive a `CoarsenRefStatus` structure, which defines the status of `XBraid` at a given instant of coarsening or refinement on some level during a run. The user accesses it through `braid_CoarsenRefStatusGet**()` functions.

14.2.2 Field Documentation

14.2.2.1 `braid_Real c_tprior`

time value to the left of `tstart` on coarse grid

14.2.2.2 `braid_Real c_tstop`

time value to the right of `tstart` on coarse grid

14.2.2.3 `braid_Real f_tprior`

time value to the left of `tstart` on fine grid

14.2.2.4 `braid_Real f_tstop`

time value to the right of `tstart` on fine grid

14.2.2.5 `braid_Int gupper`

global size of the fine grid

14.2.2.6 `braid_Int level`

current fine level in `XBraid`

14.2.2.7 `braid_Int nrefine`

number of refinements done

14.2.2.8 **braid_Real tstart**

current time value

The documentation for this struct was generated from the following file:

- [braid_status.h](#)

14.3 **_braid_CommHandle Struct Reference**

Data Fields

- [braid_Int request_type](#)
- [braid_Int num_requests](#)
- [MPI_Request * requests](#)
- [MPI_Status * status](#)
- [void * buffer](#)
- [braid_Vector * vector_ptr](#)

14.3.1 Detailed Description

XBraid comm handle structure

Used for initiating and completing nonblocking communication to pass `braid_Vectors` between processors.

14.3.2 Field Documentation

14.3.2.1 `void* buffer`

Buffer for message

14.3.2.2 `braid_Int num_requests`

number of active requests for this handle, usually 1

14.3.2.3 `braid_Int request_type`

two values: `recv type = 1`, and `send type = 0`

14.3.2.4 `MPI_Request* requests`

MPI request structure

14.3.2.5 `MPI_Status* status`

MPI status

14.3.2.6 `braid_Vector* vector_ptr`

`braid_vector` being sent/received

The documentation for this struct was generated from the following file:

- [_braid.h](#)

14.4 `_braid_Core` Struct Reference

Data Fields

- `MPI_Comm comm_world`
- `MPI_Comm comm`
- `braid_Int myid_world`
- `braid_Int myid`
- `braid_Real tstart`
- `braid_Real tstop`
- `braid_Int ntime`
- `braid_App app`
- `braid_PtFcnStep step`
- `braid_PtFcnInit init`
- `braid_PtFcnClone clone`
- `braid_PtFcnFree free`
- `braid_PtFcnSum sum`
- `braid_PtFcnSpatialNorm spatialnorm`
- `braid_PtFcnAccess access`
- `braid_PtFcnBufSize bufsize`
- `braid_PtFcnBufPack bufpack`
- `braid_PtFcnBufUnpack bufunpack`
- `braid_PtFcnResidual residual`
- `braid_PtFcnSCoarsen scoarsen`
- `braid_PtFcnSRefine srefine`
- `braid_Int access_level`
- `braid_Int print_level`
- `braid_Int seq_soln`
- `braid_Int max_levels`
- `braid_Int min_coarse`
- `braid_Real tol`
- `braid_Int rtol`
- `braid_Int * nrels`
- `braid_Int nrdefault`
- `braid_Int * cfactors`
- `braid_Int cfdefault`
- `braid_Int max_iter`
- `braid_Int niter`
- `braid_Int fmg`
- `braid_Int nfmfg`
- `braid_Int nfmfg_Vcyc`
- `braid_Int tnorm`
- `braid_Real * tnorm_a`
- `braid_Real rnorm0`
- `braid_Real * rnorms`
- `braid_PtFcnResidual full_rnorm_res`
- `braid_Real full_rnorm0`
- `braid_Real * full_rnorms`
- `braid_AccessStatus astatus`
- `braid_CoarsenRefStatus cstatus`
- `braid_StepStatus sstatus`

- [braid_Int storage](#)
- [braid_Int gupper](#)
- [braid_Int refine](#)
- [braid_Int * rfactors](#)
- [braid_Int rstopped](#)
- [braid_Int nrefine](#)
- [braid_Int max_refinements](#)
- [braid_Int tpoints_cutoff](#)
- [braid_Int skip](#)
- [braid_Int nlevels](#)
- [_braid_Grid ** grids](#)
- [braid_Real localtime](#)
- [braid_Real globaltime](#)

14.4.1 Detailed Description

The typedef [_braid_Core](#) struct is a **critical** part of XBraid and is passed to *each* routine in XBraid. It thus allows each routine access to XBraid attributes.

14.4.2 Field Documentation

14.4.2.1 [braid_PtFcnAccess](#) access

user access function to XBraid and current vector

14.4.2.2 [braid_Int access_level](#)

determines how often to call the user's access routine

14.4.2.3 [braid_App](#) app

application data for the user

14.4.2.4 [braid_AccessStatus](#) astatus

status structure passed to user-written Access routine

14.4.2.5 [braid_PtFcnBufPack](#) bufpack

pack a buffer

14.4.2.6 [braid_PtFcnBufSize](#) bufsize

return buffer size

14.4.2.7 [braid_PtFcnBufUnpack](#) bufunpack

unpack a buffer

14.4.2.8 [braid_Int* cfactors](#)

coarsening factors

14.4.2.9 braid_Int cfdefault

default coarsening factor

14.4.2.10 braid_PtFcnClone clone

clone a vector

14.4.2.11 MPI_Comm comm

communicator for the time dimension

14.4.2.12 MPI_Comm comm_world**14.4.2.13 braid_CoarsenRefStatus cstatus**

status structure passed to user-written coarsen/refine routines

14.4.2.14 braid_Int fmg

use FMG cycle

14.4.2.15 braid_PtFcnFree free

free up a vector

14.4.2.16 braid_Real full_rnorm0

(optional) initial full residual norm

14.4.2.17 braid_PtFcnResidual full_rnorm_res

(optional) used to compute full residual norm

14.4.2.18 braid_Real* full_rnorms

(optional) full residual norm history

14.4.2.19 braid_Real globaltime

global wall time for [braid_Drive\(\)](#)

14.4.2.20 _braid_Grid grids**

pointer to temporal grid structures for each level

14.4.2.21 braid_Int gupper

global size of the fine grid

14.4.2.22 braid_PtFcnInit init

return an initialized [braid_Vector](#)

14.4.2.23 braid_Real localtime

local wall time for [braid_Drive\(\)](#)

14.4.2.24 braid_Int max_iter

maximum number of multigrid in time iterations

14.4.2.25 braid_Int max_levels

maximum number of temporal grid levels

14.4.2.26 braid_Int max_refinements

maximum number of refinements

14.4.2.27 braid_Int min_coarse

minimum possible coarse grid size

14.4.2.28 braid_Int myid

my rank in the time communicator

14.4.2.29 braid_Int myid_world

my rank in the world communicator

14.4.2.30 braid_Int nfmfg

number of fmg cycles to do initially before switching to V-cycles

14.4.2.31 braid_Int nfmfg_Vcyc

number of V-cycle calls at each level in FMG

14.4.2.32 braid_Int niter

number of iterations

14.4.2.33 braid_Int nlevels

number of temporal grid levels

14.4.2.34 braid_Int nrdefault

default number of pre-relaxations

14.4.2.35 braid_Int nrefine

number of refinements done

14.4.2.36 braid_Int* nrels

number of pre-relaxations on each level

14.4.2.37 braid_Int ntime

initial number of time intervals

14.4.2.38 braid_Int print_level

determines amount of output printed to screen (0,1,2)

14.4.2.39 braid_Int refine

refine in time (refine = 1)

14.4.2.40 braid_PtFcnResidual residual

(optional) compute residual

14.4.2.41 braid_Int* ractors

refinement factors for finest grid (if any)

14.4.2.42 braid_Real rnorm0

initial residual norm

14.4.2.43 braid_Real* rnorms

residual norm history

14.4.2.44 braid_Int rstopped

refinement stopped at iteration rstopped

14.4.2.45 braid_Int rtol

use relative tolerance

14.4.2.46 braid_PtFcnSCoarsen scoarsen

(optional) return a spatially coarsened vector

14.4.2.47 braid_Int seq_soln

boolean, controls if the initial guess is from sequential time stepping

14.4.2.48 braid_Int skip

boolean, controls skipping any work on first down-cycle

14.4.2.49 braid_PtFcnSpatialNorm spatialnorm

Compute norm of a braid_Vector, this is a norm only over space

14.4.2.50 braid_PtFcnSRefine srefine

(optional) return a spatially refined vector

14.4.2.51 braid_StepStatus sstatus

status structure passed to user-written step routines

14.4.2.52 `braid_PtFcnStep` step

apply step function

14.4.2.53 `braid_Int` storage

storage = 0 (C-points), = 1 (all)

14.4.2.54 `braid_PtFcnSum` sum

vector sum

14.4.2.55 `braid_Int` tnorm

choice of temporal norm

14.4.2.56 `braid_Real*` tnorm_a

local array of residual norms on a proc's interval, used for inf-norm

14.4.2.57 `braid_Real` tol

stopping tolerance

14.4.2.58 `braid_Int` tpoints_cutoff

refinements halt after the number of time steps exceed this value

14.4.2.59 `braid_Real` tstart

start time

14.4.2.60 `braid_Real` tstop

stop time

The documentation for this struct was generated from the following file:

- [_braid.h](#)

14.5 `_braid_Grid` Struct Reference

Data Fields

- [braid_Int](#) level
- [braid_Int](#) ilower
- [braid_Int](#) iupper
- [braid_Int](#) clower
- [braid_Int](#) cupper
- [braid_Int](#) gupper
- [braid_Int](#) cfactor
- [braid_Int](#) ncpoints
- [braid_Int](#) nupoints
- [braid_Vector](#) * ua
- [braid_Real](#) * ta

- `braid_Vector * va`
- `braid_Vector * fa`
- `braid_Int recv_index`
- `braid_Int send_index`
- `_braid_CommHandle * recv_handle`
- `_braid_CommHandle * send_handle`
- `braid_Vector * ua_alloc`
- `braid_Real * ta_alloc`
- `braid_Vector * va_alloc`
- `braid_Vector * fa_alloc`

14.5.1 Detailed Description

XBraid Grid structure for a certain time level

Holds all the information for a processor related to the temporal grid at this level.

14.5.2 Field Documentation

14.5.2.1 `braid_Int cfactor`

coarsening factor

14.5.2.2 `braid_Int clower`

smallest C point index

14.5.2.3 `braid_Int cupper`

largest C point index

14.5.2.4 `braid_Vector* fa`

rhs vectors f (all points, NULL on level 0)

14.5.2.5 `braid_Vector* fa_alloc`

original memory allocation for fa

14.5.2.6 `braid_Int gupper`

global size of the grid

14.5.2.7 `braid_Int ilower`

smallest time index at this level

14.5.2.8 `braid_Int iupper`

largest time index at this level

14.5.2.9 `braid_Int level`

Level that grid is on

14.5.2.10 **braid_Int** ncpoints

number of C points

14.5.2.11 **braid_Int** nupoints

number of unknown vector points

14.5.2.12 **_braid_CommHandle*** recv_handle

Handle for nonblocking receives of **braid_Vectors**

14.5.2.13 **braid_Int** recv_index

-1 means no receive

14.5.2.14 **_braid_CommHandle*** send_handle

Handle for nonblocking sends of **braid_Vectors**

14.5.2.15 **braid_Int** send_index

-1 means no send

14.5.2.16 **braid_Real*** ta

time values (all points)

14.5.2.17 **braid_Real*** ta_alloc

original memory allocation for ta

14.5.2.18 **braid_Vector*** ua

unknown vectors (C-points at least)

14.5.2.19 **braid_Vector*** ua_alloc

original memory allocation for ua

14.5.2.20 **braid_Vector*** va

restricted unknown vectors (all points, NULL on level 0)

14.5.2.21 **braid_Vector*** va_alloc

original memory allocation for va

The documentation for this struct was generated from the following file:

- [_braid.h](#)

14.6 **_braid_StepStatus** Struct Reference

Data Fields

- [braid_Real](#) tstart

- `braid_Real` `tstop`
- `braid_Real` * `rnorms`
- `braid_Real` `old_fine_tolx`
- `braid_Int` `tight_fine_tolx`
- `braid_Real` `tol`
- `braid_Int` `iter`
- `braid_Int` `rfactor`
- `braid_Int` `level`
- `braid_Int` `nrefine`
- `braid_Int` `gupper`

14.6.1 Detailed Description

The user's step routine will receive a `StepStatus` structure, which defines the status of `XBraid` at the given instant for step evaluation on some level during a run. The user accesses it through `braid_StepStatusGet**()` functions.

14.6.2 Field Documentation

14.6.2.1 `braid_Int` `gupper`

global size of the fine grid

14.6.2.2 `braid_Int` `iter`

Current iteration (also equal to length of `rnorms`)

14.6.2.3 `braid_Int` `level`

current grid level

14.6.2.4 `braid_Int` `nrefine`

number of refinements done

14.6.2.5 `braid_Real` `old_fine_tolx`

Allows for storing the previously used fine tolerance from `GetSpatialAccuracy`

14.6.2.6 `braid_Int` `rfactor`

if set by user, allows for subdivision of this interval for better time accuracy

14.6.2.7 `braid_Real`* `rnorms`

residual norm history, (points to `Core->rnorms` object)

14.6.2.8 `braid_Int` `tight_fine_tolx`

Boolean, indicating whether the tightest fine `tolx` has been used, condition for halting

14.6.2.9 `braid_Real` `tol`

Current stopping tolerance

14.6.2.10 `braid_Real` `tstart`

current time value

14.6.2.11 `braid_Real` `tstop`

time value to evolve towards, time value to the right of `tstart`

The documentation for this struct was generated from the following file:

- [braid_status.h](#)

15 File Documentation

15.1 `_braid.h` File Reference

Data Structures

- [struct `_braid_CommHandle`](#)
- [struct `_braid_Grid`](#)
- [struct `_braid_Core`](#)

Macros

- `#define _braid_CommHandleElt(handle, elt) ((handle) -> elt)`
- `#define _braid_GridElt(grid, elt) ((grid) -> elt)`
- `#define _braid_CoreElt(core, elt) ((core) -> elt)`
- `#define _braid_CoreFcn(core, fcn) (*(core) -> fcn)`
- `#define _braid_MapFineToCoarse(findex, cfactor, cindex) (cindex = (findex)/(cfactor))`
- `#define _braid_MapCoarseToFine(cindex, cfactor, findex) (findex = (cindex)*(cfactor))`
- `#define _braid_IsFPoint(index, cfactor) ((index)%(cfactor))`
- `#define _braid_IsCPoint(index, cfactor) (!_braid_IsFPoint(index, cfactor))`
- `#define _braid_NextCPoint(index, cfactor) ((_braid_Int)((index)+(cfactor)-1)/(cfactor))*(cfactor))`

Functions

- [braid_Int `_braid_GetBlockDistInterval`](#) ([braid_Int](#) npoints, [braid_Int](#) nprocs, [braid_Int](#) proc, [braid_Int](#) *ilower_ptr, [braid_Int](#) *iupper_ptr)
- [braid_Int `_braid_GetBlockDistProc`](#) ([braid_Int](#) npoints, [braid_Int](#) nprocs, [braid_Int](#) index, [braid_Int](#) *proc_ptr)
- [braid_Int `_braid_GetDistribution`](#) ([braid_Core](#) core, [braid_Int](#) *ilower_ptr, [braid_Int](#) *iupper_ptr)
- [braid_Int `_braid_GetProc`](#) ([braid_Core](#) core, [braid_Int](#) level, [braid_Int](#) index, [braid_Int](#) *proc_ptr)
- [braid_Int `_braid_GetCFactor`](#) ([braid_Core](#) core, [braid_Int](#) level, [braid_Int](#) *cfactor_ptr)
- [braid_Int `_braid_CommRecvInit`](#) ([braid_Core](#) core, [braid_Int](#) level, [braid_Int](#) index, [braid_Vector](#) *vector_ptr, [_braid_CommHandle](#) **handle_ptr)
- [braid_Int `_braid_CommSendInit`](#) ([braid_Core](#) core, [braid_Int](#) level, [braid_Int](#) index, [braid_Vector](#) vector, [_braid_CommHandle](#) **handle_ptr)
- [braid_Int `_braid_CommWait`](#) ([braid_Core](#) core, [_braid_CommHandle](#) **handle_ptr)
- [braid_Int `_braid_UGetIndex`](#) ([braid_Core](#) core, [braid_Int](#) level, [braid_Int](#) index, [braid_Int](#) *uindex_ptr)
- [braid_Int `_braid_UGetVectorRef`](#) ([braid_Core](#) core, [braid_Int](#) level, [braid_Int](#) index, [braid_Vector](#) *u_ptr)
- [braid_Int `_braid_USetVectorRef`](#) ([braid_Core](#) core, [braid_Int](#) level, [braid_Int](#) index, [braid_Vector](#) u)

- `braid_Int_braid_UGetVector` (`braid_Core` core, `braid_Int` level, `braid_Int` index, `braid_Vector` *u_ptr)
- `braid_Int_braid_USetVector` (`braid_Core` core, `braid_Int` level, `braid_Int` index, `braid_Vector` u, `braid_Int` move)
- `braid_Int_braid_UCommInitBasic` (`braid_Core` core, `braid_Int` level, `braid_Int` recv_msg, `braid_Int` send_msg, `braid_Int` send_now)
- `braid_Int_braid_UCommInit` (`braid_Core` core, `braid_Int` level)
- `braid_Int_braid_UCommInitF` (`braid_Core` core, `braid_Int` level)
- `braid_Int_braid_UCommWait` (`braid_Core` core, `braid_Int` level)
- `braid_Int_braid_GetInterval` (`braid_Core` core, `braid_Int` level, `braid_Int` interval_index, `braid_Int` *flo_ptr, `braid_Int` *fhi_ptr, `braid_Int` *ci_ptr)
- `braid_Int_braid_AccessVector` (`braid_Core` core, `braid_AccessStatus` status, `braid_Vector` u)
- `braid_Int_braid_GetUInit` (`braid_Core` core, `braid_Int` level, `braid_Int` index, `braid_Vector` u, `braid_Vector` *ustop_ptr)
- `braid_Int_braid_Step` (`braid_Core` core, `braid_Int` level, `braid_Int` index, `braid_Vector` ustop, `braid_Vector` u)
- `braid_Int_braid_Residual` (`braid_Core` core, `braid_Int` level, `braid_Int` index, `braid_Vector` ustop, `braid_Vector` r)
- `braid_Int_braid_FASResidual` (`braid_Core` core, `braid_Int` level, `braid_Int` index, `braid_Vector` ustop, `braid_Vector` r)
- `braid_Int_braid_Coarsen` (`braid_Core` core, `braid_Int` level, `braid_Int` f_index, `braid_Int` c_index, `braid_Vector` fvector, `braid_Vector` *cvector)
- `braid_Int_braid_RefineBasic` (`braid_Core` core, `braid_Int` level, `braid_Real` *f_ta, `braid_Real` *c_ta, `braid_Vector` cvector, `braid_Vector` *fvector)
- `braid_Int_braid_Refine` (`braid_Core` core, `braid_Int` level, `braid_Int` f_index, `braid_Int` c_index, `braid_Vector` cvector, `braid_Vector` *fvector)
- `braid_Int_braid_GridInit` (`braid_Core` core, `braid_Int` level, `braid_Int` ilower, `braid_Int` iupper, `_braid_Grid` **grid_ptr)
- `braid_Int_braid_GridClean` (`braid_Core` core, `_braid_Grid` *grid)
- `braid_Int_braid_GridDestroy` (`braid_Core` core, `_braid_Grid` *grid)
- `braid_Int_braid_InitGuess` (`braid_Core` core, `braid_Int` level)
- `braid_Int_braid_ComputeFullRNorm` (`braid_Core` core, `braid_Int` level, `braid_Real` *return_rnorm)
- `braid_Int_braid_FCRRelax` (`braid_Core` core, `braid_Int` level)
- `braid_Int_braid_FRestrict` (`braid_Core` core, `braid_Int` level)
- `braid_Int_braid_FInterp` (`braid_Core` core, `braid_Int` level)
- `braid_Int_braid_FRefine` (`braid_Core` core, `braid_Int` *refined_ptr)
- `braid_Int_braid_FAccess` (`braid_Core` core, `braid_Int` level, `braid_Int` done)
- `braid_Int_braid_InitHierarchy` (`braid_Core` core, `_braid_Grid` *fine_grid, `braid_Int` refined)
- `braid_Int_braid_PrintSpatialNorms` (`braid_Core` core, `braid_Real` *rnorms, `braid_Int` n)
- `braid_Int_braid_CopyFineToCoarse` (`braid_Core` core)
- `braid_Int_braid_SetRNorm` (`braid_Core` core, `braid_Int` iter, `braid_Real` rnorm)
- `braid_Int_braid_GetRNorm` (`braid_Core` core, `braid_Int` iter, `braid_Real` *rnorm_ptr)
- `braid_Int_braid_SetFullRNorm` (`braid_Core` core, `braid_Int` iter, `braid_Real` rnorm)
- `braid_Int_braid_GetFullRNorm` (`braid_Core` core, `braid_Int` iter, `braid_Real` *rnorm_ptr)
- `braid_Int_braid_DeleteLastResidual` (`braid_Core` core)

Variables

- FILE * `_braid_printfile`

15.1.1 Detailed Description

Define headers for developer routines. This file contains the headers for developer routines.

15.1.2 Macro Definition Documentation

15.1.2.1 `#define _braid_CommHandleElt(handle, elt)((handle)-> elt)`

Accessor for `_braid_CommHandle` attributes

15.1.2.2 `#define _braid_CoreElt(core, elt)((core)-> elt)`

Accessor for `_braid_Core` attributes

15.1.2.3 `#define _braid_CoreFcn(core, fcn)(*(core)-> fcn)`

Accessor for `_braid_Core` functions

15.1.2.4 `#define _braid_GridElt(grid, elt)((grid)-> elt)`

Accessor for `_braid_Grid` attributes

15.1.2.5 `#define _braid_IsCPoint(index, cfactor)(!_braid_IsFPoint(index, cfactor))`

Boolean, returns whether a time index is an C-point

15.1.2.6 `#define _braid_IsFPoint(index, cfactor)((index)%(cfactor))`

Boolean, returns whether a time index is an F-point

15.1.2.7 `#define _braid_MapCoarseToFine(cindex, cfactor, findex) (findex = (cindex)*(cfactor))`

Map a coarse time index to a fine time index, assumes a uniform coarsening factor.

15.1.2.8 `#define _braid_MapFineToCoarse(findex, cfactor, cindex) (cindex = (findex)/(cfactor))`

Map a fine time index to a coarse time index, assumes a uniform coarsening factor.

15.1.2.9 `#define _braid_NextCPoint(index, cfactor)(((braid_Int)((index)+(cfactor)-1)/(cfactor))*(cfactor))`

Returns the index for the next C-point to the right of index (inclusive)

15.1.3 Function Documentation

15.1.3.1 `braid_Int _braid_AccessVector (braid_Core core, braid_AccessStatus status, braid_Vector u)`

Call user's access function in order to give access to XBraid and the current vector. Most commonly, this lets the user write u to screen, disk, etc... The vector u corresponds to time step $index$ on $level$. $status$ holds state information about the current XBraid iteration, time value, etc...

15.1.3.2 `braid_Int _braid_Coarsen (braid_Core core, braid_Int level, braid_Int f_index, braid_Int c_index, braid_Vector fvector, braid_Vector * cvector)`

Coarsen in space on $level$ by calling the user's coarsen function. The vector corresponding to the time step index f_index on the fine grid is coarsened to the time step index c_index on the coarse grid. The output goes in $cvector$ and the input vector is $fvector$.

15.1.3.3 `braid_Int_braid_CommRecvInit (braid_Core core, braid_Int level, braid_Int index, braid_Vector * vector_ptr, _braid_CommHandle ** handle_ptr)`

Initialize a receive to go into `vector_ptr` for the given time `index` on `level`. Also return a comm handle `handle_ptr` for querying later, to see if the receive has occurred.

15.1.3.4 `braid_Int_braid_CommSendInit (braid_Core core, braid_Int level, braid_Int index, braid_Vector vector, _braid_CommHandle ** handle_ptr)`

Initialize a send of `vector` for the given time `index` on `level`. Also return a comm handle `handle_ptr` for querying later, to see if the send has occurred.

15.1.3.5 `braid_Int_braid_CommWait (braid_Core core, _braid_CommHandle ** handle_ptr)`

Block on the comm handle `handle_ptr` until the MPI operation (send or recv) has completed

15.1.3.6 `braid_Int_braid_ComputeFullRNorm (braid_Core core, braid_Int level, braid_Real * return_rnorm)`

Compute full temporal residual norm with user-provided residual routine. Output goes in `*return_rnorm`.

15.1.3.7 `braid_Int_braid_CopyFineToCoarse (braid_Core core)`

Copy the initialized C-points on the fine grid, to all coarse levels. For instance, if a point `k` on level `m` corresponds to point `p` on level 0, then they are equivalent after this function. The only exception is any spatial coarsening the user decides to do. This function allows XBraid to skip all work on the first down cycle and start in FMG style on the coarsest level. Assumes level 0 C-points are initialized.

15.1.3.8 `braid_Int_braid_DeleteLastResidual (braid_Core core)`

Delete the last residual, for use if F-Refinement is done.

15.1.3.9 `braid_Int_braid_FAccess (braid_Core core, braid_Int level, braid_Int done)`

Call the user's access function in order to give access to XBraid and the current vector at grid `level` and iteration `*iter`. Most commonly, this lets the user write solutions to screen, disk, etc... The quantity `rnorm` denotes the last computed residual norm, and `done` is a boolean indicating whether XBraid has finished iterating and this is the last Access call.

15.1.3.10 `braid_Int_braid_FASResidual (braid_Core core, braid_Int level, braid_Int index, braid_Vector ustop, braid_Vector r)`

Compute FAS residual = $f - \text{residual}$

15.1.3.11 `braid_Int_braid_FCRelax (braid_Core core, braid_Int level)`

Do `nu` sweeps of F-then-C relaxation on `level`

15.1.3.12 `braid_Int_braid_FInterp (braid_Core core, braid_Int level)`

F-Relax on `level` and interpolate to `level-1`

The output is set in the `braid_Grid` in `core`, so that the vector `u` on `level` is created by interpolating from `level+1`.

If the user has set spatial refinement, then this user-defined routine is also called.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>level</i>	interp from level to level+1

15.1.3.13 **braid_Int_braid_FRefine (braid_Core core, braid_Int * refined_ptr)**

Create a new fine grid (level 0) and corresponding grid hierarchy by refining the current fine grid based on user-provided refinement factors. Return the boolean *refined_ptr* to indicate whether grid refinement was actually done. To simplify the algorithm, refinement factors are automatically capped to be no greater than the coarsening factor (for level 0). The grid data is also redistributed to achieve good load balance in the temporal dimension. If the refinement factor is 1 in each time interval, no refinement is done.

15.1.3.14 **braid_Int_braid_FRestrict (braid_Core core, braid_Int level)**

F-Relax on *level* and then restrict to *level+1*

The output is set in the *braid_Grid* in *core*, so that the restricted vectors *va* and *fa* will be created, representing *level+1* versions of the unknown and rhs vectors.

If the user has set spatial coarsening, then this user-defined routine is also called.

If *level==0*, then *rnorm_ptr* will contain the residual norm.

Parameters

<i>core</i>	braid_Core (_braid_Core) struct
<i>level</i>	restrict from level to level+1

15.1.3.15 **braid_Int_braid_GetBlockDistInterval (braid_Int npoints, braid_Int nprocs, braid_Int proc, braid_Int * ilower_ptr, braid_Int * iupper_ptr)**

Returns the index interval for *proc* in a blocked data distribution.

15.1.3.16 **braid_Int_braid_GetBlockDistProc (braid_Int npoints, braid_Int nprocs, braid_Int index, braid_Int * proc_ptr)**

Returns the processor that owns *index* in a blocked data distribution (returns -1 if *index* is out of range).

15.1.3.17 **braid_Int_braid_GetCFactor (braid_Core core, braid_Int level, braid_Int * cfactor_ptr)**

Returns the coarsening factor to use on grid *level*.

15.1.3.18 **braid_Int_braid_GetDistribution (braid_Core core, braid_Int * ilower_ptr, braid_Int * iupper_ptr)**

Returns the index interval for my processor on the finest grid level. For the processor rank calling this function, it returns the smallest and largest time indices (*ilower_ptr* and *iupper_ptr*) that belong to that processor (the indices may be F or C points).

15.1.3.19 **braid_Int_braid_GetFullRNorm (braid_Core core, braid_Int iter, braid_Real * rnorm_ptr)**

Same as *GetRNorm*, but gets full residual norm.

15.1.3.20 **braid_Int_braid_GetInterval (braid_Core core, braid_Int level, braid_Int interval_index, braid_Int * flo_ptr, braid_Int * fhi_ptr, braid_Int * ci_ptr)**

Retrieve the time step indices at this *level* corresponding to a local FC interval given by *interval_index*. Argument *ci_ptr* is the time step index for the C-pt and *flo_ptr* and *fhi_ptr* are the smallest and largest F-pt indices in this interval. The

C-pt is always to the right of the F-interval, but neither a C-pt or an F-interval are guaranteed. If the `ci_ptr` returns a -1, there is no C-pt. If the `flo_ptr` is greater than the `fhi_ptr`, there is no F-interval.

15.1.3.21 `braid_Int_braid_GetProc (braid_Core core, braid_Int level, braid_Int index, braid_Int * proc_ptr)`

Returns the processor number in `proc_ptr` on which the time step `index` lives for the given `level`. Returns -1 if `index` is out of range.

15.1.3.22 `braid_Int_braid_GetRNorm (braid_Core core, braid_Int iter, braid_Real * rnorm_ptr)`

Get the residual norm for iteration `iter`. If `iter < 0`, get the `rnorm` for the last iteration minus `|iter|-1`.

15.1.3.23 `braid_Int_braid_GetUlnit (braid_Core core, braid_Int level, braid_Int index, braid_Vector u, braid_Vector * ustop_ptr)`

Return an initial guess in `ustop_ptr` to use in the step routine for implicit schemes. The value returned depends on the storage options used. If the return value is NULL, no initial guess is available.

15.1.3.24 `braid_Int_braid_GridClean (braid_Core core, _braid_Grid * grid)`

Destroy the vectors on `grid`

15.1.3.25 `braid_Int_braid_GridDestroy (braid_Core core, _braid_Grid * grid)`

Destroy `grid`

15.1.3.26 `braid_Int_braid_GridInit (braid_Core core, braid_Int level, braid_Int ilower, braid_Int iupper, _braid_Grid ** grid_ptr)`

Create a new grid object `grid_ptr` with level indicator `level`. Arguments `ilower` and `iupper` correspond to the lower and upper time index values for this processor on this grid.

15.1.3.27 `braid_Int_braid_InitGuess (braid_Core core, braid_Int level)`

Set initial guess on `level`. Only C-pts are initialized on level 0, otherwise stored values are initialized based on restricted fine-grid values.

15.1.3.28 `braid_Int_braid_InitHierarchy (braid_Core core, _braid_Grid * fine_grid, braid_Int refined)`

Initialize grid hierarchy with `fine_grid` serving as the finest grid. Boolean `refined` indicates whether `fine_grid` was created by refining a coarser grid (in the `FRefine()` routine), which has implications on how to define CF-intervals.

15.1.3.29 `braid_Int_braid_PrintSpatialNorms (braid_Core core, braid_Real * rnorms, braid_Int n)`

Print out the residual norm for every C-point. Processor 0 gathers all the `rnorms` and prints them in order through a `gather` operator

15.1.3.30 `braid_Int_braid_Refine (braid_Core core, braid_Int level, braid_Int f_index, braid_Int c_index, braid_Vector cvector, braid_Vector * fvector)`

Refine in space on `level` by calling the user's refine function. The vector corresponding to the time step index `c_index` on the coarse grid is refined to the time step index `f_index` on the fine grid. The output goes in `fvector` and the input vector is `cvector`.

15.1.3.31 **braid_Int_braid_RefineBasic** (**braid_Core** *core*, **braid_Int** *level*, **braid_Real** * *f_ta*, **braid_Real** * *c_ta*, **braid_Vector** *cvector*, **braid_Vector** * *fvector*)

Refine in space (basic routine)

15.1.3.32 **braid_Int_braid_Residual** (**braid_Core** *core*, **braid_Int** *level*, **braid_Int** *index*, **braid_Vector** *ustop*, **braid_Vector** *r*)

Compute residual *r*

15.1.3.33 **braid_Int_braid_SetFullRNorm** (**braid_Core** *core*, **braid_Int** *iter*, **braid_Real** *rnorm*)

Same as SetRNorm, but sets full residual norm.

15.1.3.34 **braid_Int_braid_SetRNorm** (**braid_Core** *core*, **braid_Int** *iter*, **braid_Real** *rnorm*)

Set the residual norm for iteration *iter*. If *iter* < 0, set the *rnorm* for the last iteration minus $|\text{iter}|-1$. Also set the initial residual norm.

15.1.3.35 **braid_Int_braid_Step** (**braid_Core** *core*, **braid_Int** *level*, **braid_Int** *index*, **braid_Vector** *ustop*, **braid_Vector** *u*)

Integrate one time step at time step *index* to time step *index*+1.

15.1.3.36 **braid_Int_braid_UCommlnit** (**braid_Core** *core*, **braid_Int** *level*)

This routine initiates communication under the assumption that work will be done on all intervals (F or C) on *level*. It posts a receive for the point to the left of *ilower* (regardless whether *ilower* is F or C), and it posts a send of *iupper* if *iupper* is a C point.

15.1.3.37 **braid_Int_braid_UCommlnitBasic** (**braid_Core** *core*, **braid_Int** *level*, **braid_Int** *recv_msg*, **braid_Int** *send_msg*, **braid_Int** *send_now*)

Basic communication (from the left, to the right). Arguments *recv_msg* and *send_msg* are booleans that indicate whether or not to initiate a receive from the left and a send to the right respectively. Argument *send_now* indicates that the send should be initiated immediately.

15.1.3.38 **braid_Int_braid_UCommlnitF** (**braid_Core** *core*, **braid_Int** *level*)

This routine initiates communication under the assumption that work will be done on only F-pt intervals on *level*. It only posts a receive for the point to the left of *ilower* if *ilower* is an F point, and it posts a send of *iupper* if *iupper* is a C point.

15.1.3.39 **braid_Int_braid_UCommWait** (**braid_Core** *core*, **braid_Int** *level*)

Finish up communication. On *level*, wait on both the *recv* and *send* handles at this level.

15.1.3.40 **braid_Int_braid_UGetIndex** (**braid_Core** *core*, **braid_Int** *level*, **braid_Int** *index*, **braid_Int** * *uindex_ptr*)

Returns an index into the local u-vector for grid *level* at point *index*. If the u-vector is not stored, returns -1.

15.1.3.41 **braid_Int_braid_UGetVector** (**braid_Core** *core*, **braid_Int** *level*, **braid_Int** *index*, **braid_Vector** * *u_ptr*)

Returns a copy of the u-vector on grid *level* at point *index*. If *index* is my "receive index" (as set by UCommlnit()), for example), the u-vector will be received from a neighbor processor. If the u-vector is not stored, NULL is returned.

15.1.3.42 `braid_Int_braid_UGetVectorRef (braid_Core core, braid_Int level, braid_Int index, braid_Vector * u_ptr)`

Returns a reference to the local u-vector on grid *level* at point *index*. If the u-vector is not stored, returns NULL.

15.1.3.43 `braid_Int_braid_USetVector (braid_Core core, braid_Int level, braid_Int index, braid_Vector u, braid_Int move)`

Stores the u-vector on grid *level* at point *index*. If *index* is my "send index", a send is initiated to a neighbor processor. If *move* is true, the u-vector is moved into core storage instead of copied. If the u-vector is not stored, nothing is done.

15.1.3.44 `braid_Int_braid_USetVectorRef (braid_Core core, braid_Int level, braid_Int index, braid_Vector u)`

Stores a reference to the local u-vector on grid *level* at point *index*. If the u-vector is not stored, nothing is done.

15.1.4 Variable Documentation

15.1.4.1 FILE*_braid_printfile

This is the print file for redirecting stdout for all XBraid screen output

15.2 braid.h File Reference

Macros

- `#define braid_FMANGLE 1`
- `#define braid_Fortran_SpatialCoarsen 0`
- `#define braid_Fortran_Residual 0`
- `#define braid_INVALID_RNORM -1`
- `#define braid_ERROR_GENERIC 1 /* generic error */`
- `#define braid_ERROR_MEMORY 2 /* unable to allocate memory */`
- `#define braid_ERROR_ARG 4 /* argument error */`

Typedefs

- `typedef struct _braid_App_struct * braid_App`
- `typedef struct _braid_Vector_struct * braid_Vector`
- `typedef braid_Int(* braid_PtFcnStep)(braid_App app, braid_Vector ustop, braid_Vector fstop, braid_Vector u, braid_StepStatus status)`
- `typedef braid_Int(* braid_PtFcnInit)(braid_App app, braid_Real t, braid_Vector *u_ptr)`
- `typedef braid_Int(* braid_PtFcnClone)(braid_App app, braid_Vector u, braid_Vector *v_ptr)`
- `typedef braid_Int(* braid_PtFcnFree)(braid_App app, braid_Vector u)`
- `typedef braid_Int(* braid_PtFcnSum)(braid_App app, braid_Real alpha, braid_Vector x, braid_Real beta, braid_Vector y)`
- `typedef braid_Int(* braid_PtFcnSpatialNorm)(braid_App app, braid_Vector u, braid_Real *norm_ptr)`
- `typedef braid_Int(* braid_PtFcnAccess)(braid_App app, braid_Vector u, braid_AccessStatus status)`
- `typedef braid_Int(* braid_PtFcnBufSize)(braid_App app, braid_Int *size_ptr)`
- `typedef braid_Int(* braid_PtFcnBufPack)(braid_App app, braid_Vector u, void *buffer, braid_Int *size_ptr)`
- `typedef braid_Int(* braid_PtFcnBufUnpack)(braid_App app, void *buffer, braid_Vector *u_ptr)`
- `typedef braid_Int(* braid_PtFcnResidual)(braid_App app, braid_Vector ustop, braid_Vector r, braid_StepStatus status)`

- typedef `braid_Int`(* `braid_PtFcnSCoarsen`)(`braid_App` app, `braid_Vector` fu, `braid_Vector` *cu_ptr, `braid_CoarsenRefStatus` status)
- typedef `braid_Int`(* `braid_PtFcnSRefine`)(`braid_App` app, `braid_Vector` cu, `braid_Vector` *fu_ptr, `braid_CoarsenRefStatus` status)
- typedef struct `_braid_Core_struct` * `braid_Core`

Functions

- `braid_Int` `braid_Init` (`MPI_Comm` comm_world, `MPI_Comm` comm, `braid_Real` tstart, `braid_Real` tstop, `braid_Int` ntime, `braid_App` app, `braid_PtFcnStep` step, `braid_PtFcnInit` init, `braid_PtFcnClone` clone, `braid_PtFcnFree` free, `braid_PtFcnSum` sum, `braid_PtFcnSpatialNorm` spatialnorm, `braid_PtFcnAccess` access, `braid_PtFcnBufSize` bufsize, `braid_PtFcnBufPack` bufpack, `braid_PtFcnBufUnpack` bufunpack, `braid_Core` *core_ptr)
- `braid_Int` `braid_Drive` (`braid_Core` core)
- `braid_Int` `braid_Destroy` (`braid_Core` core)
- `braid_Int` `braid_PrintStats` (`braid_Core` core)
- `braid_Int` `braid_SetMaxLevels` (`braid_Core` core, `braid_Int` max_levels)
- `braid_Int` `braid_SetSkip` (`braid_Core` core, `braid_Int` skip)
- `braid_Int` `braid_SetRefine` (`braid_Core` core, `braid_Int` refine)
- `braid_Int` `braid_SetMaxRefinements` (`braid_Core` core, `braid_Int` max_refinements)
- `braid_Int` `braid_SetTPointsCutoff` (`braid_Core` core, `braid_Int` tpoints_cutoff)
- `braid_Int` `braid_SetMinCoarse` (`braid_Core` core, `braid_Int` min_coarse)
- `braid_Int` `braid_SetAbsTol` (`braid_Core` core, `braid_Real` atol)
- `braid_Int` `braid_SetRelTol` (`braid_Core` core, `braid_Real` rtol)
- `braid_Int` `braid_SetNRRelax` (`braid_Core` core, `braid_Int` level, `braid_Int` nrelax)
- `braid_Int` `braid_SetCFactor` (`braid_Core` core, `braid_Int` level, `braid_Int` cfactor)
- `braid_Int` `braid_SetMaxIter` (`braid_Core` core, `braid_Int` max_iter)
- `braid_Int` `braid_SetFMG` (`braid_Core` core)
- `braid_Int` `braid_SetNFMG` (`braid_Core` core, `braid_Int` k)
- `braid_Int` `braid_SetNFMGVcyc` (`braid_Core` core, `braid_Int` nfmvg_Vcyc)
- `braid_Int` `braid_SetStorage` (`braid_Core` core, `braid_Int` storage)
- `braid_Int` `braid_SetTemporalNorm` (`braid_Core` core, `braid_Int` tnorm)
- `braid_Int` `braid_SetResidual` (`braid_Core` core, `braid_PtFcnResidual` residual)
- `braid_Int` `braid_SetFullRNormRes` (`braid_Core` core, `braid_PtFcnResidual` residual)
- `braid_Int` `braid_SetSpatialCoarsen` (`braid_Core` core, `braid_PtFcnSCoarsen` scoarsen)
- `braid_Int` `braid_SetSpatialRefine` (`braid_Core` core, `braid_PtFcnSRefine` srefine)
- `braid_Int` `braid_SetPrintLevel` (`braid_Core` core, `braid_Int` print_level)
- `braid_Int` `braid_SetPrintFile` (`braid_Core` core, const char *printfile_name)
- `braid_Int` `braid_SetDefaultPrintFile` (`braid_Core` core)
- `braid_Int` `braid_SetAccessLevel` (`braid_Core` core, `braid_Int` access_level)
- `braid_Int` `braid_SplitCommworld` (const `MPI_Comm` *comm_world, `braid_Int` px, `MPI_Comm` *comm_x, `MPI_Comm` *comm_t)
- `braid_Int` `braid_GetNumIter` (`braid_Core` core, `braid_Int` *niter_ptr)
- `braid_Int` `braid_GetRNorms` (`braid_Core` core, `braid_Int` *nrequest_ptr, `braid_Real` *rnorms)
- `braid_Int` `braid_GetNLevels` (`braid_Core` core, `braid_Int` *nlevels_ptr)
- `braid_Int` `braid_GetSpatialAccuracy` (`braid_StepStatus` status, `braid_Real` loose_tol, `braid_Real` tight_tol, `braid_Real` *tol_ptr)
- `braid_Int` `braid_SetSeqSoln` (`braid_Core` core, `braid_Int` seq_soln)

15.2.1 Detailed Description

Define headers for user interface routines. This file contains routines used to allow the user to initialize, run and get and set a XBraid solver.

15.2.2 Macro Definition Documentation

15.2.2.1 `#define braid_ERROR_ARG 4 /* argument error */`

15.2.2.2 `#define braid_ERROR_GENERIC 1 /* generic error */`

15.2.2.3 `#define braid_ERROR_MEMORY 2 /* unable to allocate memory */`

15.2.2.4 `#define braid_FMANGLE 1`

Turn on Fortran 90 interface options manually

15.2.2.5 `#define braid_Fortran_Residual 0`

15.2.2.6 `#define braid_Fortran_SpatialCoarsen 0`

15.2.2.7 `#define braid_INVALID_RNORM -1`

Value used to represent an invalid residual norm

15.3 braid_defs.h File Reference

Macros

- `#define braid_Int_Max INT_MAX;`
- `#define braid_Int_Min INT_MIN;`
- `#define braid_MPI_REAL MPI_DOUBLE`
- `#define braid_MPI_INT MPI_INT`
- `#define _braid_Error(IERR, msg) _braid_ErrorHandler(__FILE__, __LINE__, IERR, msg)`
- `#define _braid_ErrorInArg(IARG, msg) _braid_Error(HYPRE_ERROR_ARG | IARG << 3, msg)`
- `#define _braid_TAlloc(type, count) ((type *)malloc((size_t)(sizeof(type) * (count))))`
- `#define _braid_CTAlloc(type, count) ((type *)calloc((size_t)(count), (size_t)sizeof(type)))`
- `#define _braid_TReAlloc(ptr, type, count) ((type *)realloc((char *)ptr, (size_t)(sizeof(type) * (count))))`
- `#define _braid_TFree(ptr) (free((char *)ptr), ptr = NULL)`
- `#define _braid_max(a, b) (((a) < (b)) ? (b) : (a))`
- `#define _braid_min(a, b) (((a) < (b)) ? (a) : (b))`
- `#define braid_isnan(a) (a != a)`

Typedefs

- typedef int [braid_Int](#)
- typedef double [braid_Real](#)

Functions

- void [_braid_ErrorHandler](#) (const char *filename, [braid_Int](#) line, [braid_Int](#) ierr, const char *msg)

Variables

- [braid_Int_braid_error_flag](#)

15.3.1 Detailed Description

Definitions of types, error flags, etc...

15.3.2 Macro Definition Documentation

15.3.2.1 `#define _braid_CTAlloc(type, count) ((type *)calloc((size_t)(count), (size_t)sizeof(type)))`

Allocation macro

15.3.2.2 `#define _braid_Error(IERR, msg) _braid_ErrorHandler(__FILE__, __LINE__, IERR, msg)`

15.3.2.3 `#define _braid_ErrorInArg(IARG, msg) _braid_Error(HYPRE_ERROR_ARG | IARG<<3, msg)`

15.3.2.4 `#define _braid_max(a, b) (((a)<(b)) ? (b) : (a))`

15.3.2.5 `#define _braid_min(a, b) (((a)<(b)) ? (a) : (b))`

15.3.2.6 `#define _braid_TAlloc(type, count) ((type *)malloc((size_t)(sizeof(type) * (count))))`

Allocation macro

15.3.2.7 `#define _braid_TFree(ptr) (free((char *)ptr), ptr = NULL)`

Free memory macro

15.3.2.8 `#define _braid_TReAlloc(ptr, type, count) ((type *)realloc((char *)ptr, (size_t)(sizeof(type) * (count))))`

Re-allocation macro

15.3.2.9 `#define braid_Int_Max INT_MAX;`

15.3.2.10 `#define braid_Int_Min INT_MIN;`

15.3.2.11 `#define braid_isnan(a) (a != a)`

15.3.2.12 `#define braid_MPI_INT MPI_INT`

15.3.2.13 `#define braid_MPI_REAL MPI_DOUBLE`

15.3.3 Typedef Documentation

15.3.3.1 `typedef int braid_Int`

Defines integer type

15.3.3.2 `typedef double braid_Real`

Defines floating point type

15.3.4 Function Documentation

15.3.4.1 void `braid_ErrorHandler` (const char * *filename*, braid_Int *line*, braid_Int *ierr*, const char * *msg*)

15.3.5 Variable Documentation

15.3.5.1 braid_Int `_braid_error_flag`

This is the global XBraid error flag. If it is ever nonzero, an error has occurred.

15.4 braid_status.h File Reference

Data Structures

- struct `_braid_AccessStatus`
- struct `_braid_CoarsenRefStatus`
- struct `_braid_StepStatus`

Macros

- #define `_braid_StatusElt(status, elt) ((status) -> elt)`

Typedefs

- typedef struct `_braid_AccessStatus_struct` * `braid_AccessStatus`
- typedef struct `_braid_CoarsenRefStatus_struct` * `braid_CoarsenRefStatus`
- typedef struct `_braid_StepStatus_struct` * `braid_StepStatus`

Functions

- `braid_Int braid_AccessStatusInit` (braid_Real *t*, braid_Real *rnorm*, braid_Int *iter*, braid_Int *level*, braid_Int *nrefine*, braid_Int *gupper*, braid_Int *done*, braid_Int *wrapper_test*, `braid_AccessStatus` *status*)
- `braid_Int braid_AccessStatusDestroy` (`braid_AccessStatus` *status*)
- `braid_Int braid_AccessStatusGetT` (`braid_AccessStatus` *status*, braid_Real **t_ptr*)
- `braid_Int braid_AccessStatusGetResidual` (`braid_AccessStatus` *status*, braid_Real **rnorm_ptr*)
- `braid_Int braid_AccessStatusGetIter` (`braid_AccessStatus` *status*, braid_Int **iter_ptr*)
- `braid_Int braid_AccessStatusGetLevel` (`braid_AccessStatus` *status*, braid_Int **level_ptr*)
- `braid_Int braid_AccessStatusGetNRefine` (`braid_AccessStatus` *status*, braid_Int **nrefine_ptr*)
- `braid_Int braid_AccessStatusGetNTPoints` (`braid_AccessStatus` *status*, braid_Int **ntpoints_ptr*)
- `braid_Int braid_AccessStatusGetDone` (`braid_AccessStatus` *status*, braid_Int **done_ptr*)
- `braid_Int braid_AccessStatusGetWrapperTest` (`braid_AccessStatus` *status*, braid_Int **wtest_ptr*)
- `braid_Int braid_AccessStatusGetTILD` (`braid_AccessStatus` *status*, braid_Real **t_ptr*, braid_Int **iter_ptr*, braid_Int **level_ptr*, braid_Int **done_ptr*)
- `braid_Int braid_CoarsenRefStatusInit` (braid_Real *tstart*, braid_Real *f_tprior*, braid_Real *f_tstop*, braid_Real *c_tprior*, braid_Real *c_tstop*, braid_Int *level*, braid_Int *nrefine*, braid_Int *gupper*, `braid_CoarsenRefStatus` *status*)
- `braid_Int braid_CoarsenRefStatusDestroy` (`braid_CoarsenRefStatus` *status*)
- `braid_Int braid_CoarsenRefStatusGetTstart` (`braid_CoarsenRefStatus` *status*, braid_Real **tstart_ptr*)

- [braid_Int braid_CoarsenRefStatusGetFTstop](#) ([braid_CoarsenRefStatus](#) status, [braid_Real](#) *f_tstop_ptr)
- [braid_Int braid_CoarsenRefStatusGetFTprior](#) ([braid_CoarsenRefStatus](#) status, [braid_Real](#) *f_tprior_ptr)
- [braid_Int braid_CoarsenRefStatusGetCTstop](#) ([braid_CoarsenRefStatus](#) status, [braid_Real](#) *c_tstop_ptr)
- [braid_Int braid_CoarsenRefStatusGetCTprior](#) ([braid_CoarsenRefStatus](#) status, [braid_Real](#) *c_tprior_ptr)
- [braid_Int braid_CoarsenRefStatusGetTpriorTstop](#) ([braid_CoarsenRefStatus](#) status, [braid_Real](#) *tstart_ptr, [braid_Real](#) *f_tprior_ptr, [braid_Real](#) *f_tstop_ptr, [braid_Real](#) *c_tprior_ptr, [braid_Real](#) *c_tstop_ptr)
- [braid_Int braid_CoarsenRefStatusGetLevel](#) ([braid_CoarsenRefStatus](#) status, [braid_Int](#) *level_ptr)
- [braid_Int braid_CoarsenRefStatusGetNRefine](#) ([braid_CoarsenRefStatus](#) status, [braid_Int](#) *nrefine_ptr)
- [braid_Int braid_CoarsenRefStatusGetNTPoints](#) ([braid_CoarsenRefStatus](#) status, [braid_Int](#) *ntpoints_ptr)
- [braid_Int _braid_StepStatusInit](#) ([braid_Real](#) tstart, [braid_Real](#) tstop, [braid_Real](#) tol, [braid_Int](#) iter, [braid_Int](#) level, [braid_Int](#) nrefine, [braid_Int](#) gupper, [braid_StepStatus](#) status)
- [braid_Int _braid_StepStatusDestroy](#) ([braid_StepStatus](#) status)
- [braid_Int braid_StepStatusGetTstart](#) ([braid_StepStatus](#) status, [braid_Real](#) *tstart_ptr)
- [braid_Int braid_StepStatusGetTstop](#) ([braid_StepStatus](#) status, [braid_Real](#) *tstop_ptr)
- [braid_Int braid_StepStatusGetLevel](#) ([braid_StepStatus](#) status, [braid_Int](#) *level_ptr)
- [braid_Int braid_StepStatusGetNRefine](#) ([braid_StepStatus](#) status, [braid_Int](#) *nrefine_ptr)
- [braid_Int braid_StepStatusGetNTPoints](#) ([braid_StepStatus](#) status, [braid_Int](#) *ntpoints_ptr)
- [braid_Int braid_StepStatusSetRFactor](#) ([braid_StepStatus](#) status, [braid_Real](#) rfactor)
- [braid_Int braid_StepStatusGetTstartTstop](#) ([braid_StepStatus](#) status, [braid_Real](#) *tstart_ptr, [braid_Real](#) *tstop_ptr)
- [braid_Int braid_StepStatusGetTol](#) ([braid_StepStatus](#) status, [braid_Real](#) *tol_ptr)
- [braid_Int braid_StepStatusGetIter](#) ([braid_StepStatus](#) status, [braid_Int](#) *iter_ptr)
- [braid_Int braid_StepStatusGetRNorms](#) ([braid_StepStatus](#) status, [braid_Int](#) *nrequest_ptr, [braid_Real](#) *rnorms)
- [braid_Int braid_StepStatusGetOldFineToIx](#) ([braid_StepStatus](#) status, [braid_Real](#) *old_fine_tolx_ptr)
- [braid_Int braid_StepStatusSetOldFineToIx](#) ([braid_StepStatus](#) status, [braid_Real](#) old_fine_tolx_ptr)
- [braid_Int braid_StepStatusSetTightFineToIx](#) ([braid_StepStatus](#) status, [braid_Int](#) tight_fine_tolx)

15.4.1 Detailed Description

Define headers for XBraid status structures, status get/set routines and status create/destroy routines.

15.5 braid_test.h File Reference

Functions

- [braid_Int braid_TestInitAccess](#) ([braid_App](#) app, [MPI_Comm](#) comm_x, [FILE](#) *fp, [braid_Real](#) t, [braid_PtFcnInit](#) init, [braid_PtFcnAccess](#) access, [braid_PtFcnFree](#) free)
- [braid_Int braid_TestClone](#) ([braid_App](#) app, [MPI_Comm](#) comm_x, [FILE](#) *fp, [braid_Real](#) t, [braid_PtFcnInit](#) init, [braid_PtFcnAccess](#) access, [braid_PtFcnFree](#) free, [braid_PtFcnClone](#) clone)
- [braid_Int braid_TestSum](#) ([braid_App](#) app, [MPI_Comm](#) comm_x, [FILE](#) *fp, [braid_Real](#) t, [braid_PtFcnInit](#) init, [braid_PtFcnAccess](#) access, [braid_PtFcnFree](#) free, [braid_PtFcnClone](#) clone, [braid_PtFcnSum](#) sum)
- [braid_Int braid_TestSpatialNorm](#) ([braid_App](#) app, [MPI_Comm](#) comm_x, [FILE](#) *fp, [braid_Real](#) t, [braid_PtFcnInit](#) init, [braid_PtFcnFree](#) free, [braid_PtFcnClone](#) clone, [braid_PtFcnSum](#) sum, [braid_PtFcnSpatialNorm](#) spatialnorm)
- [braid_Int braid_TestBuf](#) ([braid_App](#) app, [MPI_Comm](#) comm_x, [FILE](#) *fp, [braid_Real](#) t, [braid_PtFcnInit](#) init, [braid_PtFcnFree](#) free, [braid_PtFcnSum](#) sum, [braid_PtFcnSpatialNorm](#) spatialnorm, [braid_PtFcnBufSize](#) bufsize, [braid_PtFcnBufPack](#) bufpack, [braid_PtFcnBufUnpack](#) bufunpack)
- [braid_Int braid_TestCoarsenRefine](#) ([braid_App](#) app, [MPI_Comm](#) comm_x, [FILE](#) *fp, [braid_Real](#) t, [braid_Real](#) fdt, [braid_Real](#) cdt, [braid_PtFcnInit](#) init, [braid_PtFcnAccess](#) access, [braid_PtFcnFree](#) free, [braid_PtFcnClone](#) clone, [braid_PtFcnSum](#) sum, [braid_PtFcnSpatialNorm](#) spatialnorm, [braid_PtFcnSCoarsen](#) coarsen, [braid_PtFcnSRefine](#) refine)

- `braid_Int braid_TestAll` (`braid_App` app, `MPI_Comm` comm_x, `FILE` *fp, `braid_Real` t, `braid_Real` fdt, `braid_Real` cdt, `braid_PtFcnInit` init, `braid_PtFcnFree` free, `braid_PtFcnClone` clone, `braid_PtFcnSum` sum, `braid_PtFcnSpatialNorm` spatialnorm, `braid_PtFcnBufSize` bufsize, `braid_PtFcnBufPack` bufpack, `braid_PtFcnBufUnpack` bufunpack, `braid_PtFcnSCoarsen` coarsen, `braid_PtFcnSRefine` refine)

15.5.1 Detailed Description

Define headers for XBraid test routines. This file contains routines used to test a user's XBraid wrapper routines one-by-one.

Index

- [_braid.h, 75](#)
- [_braid_AccessVector, 77](#)
- [_braid_Coarsen, 77](#)
- [_braid_CommHandleElt, 77](#)
- [_braid_CommRecvInit, 77](#)
- [_braid_CommSendInit, 78](#)
- [_braid_CommWait, 78](#)
- [_braid_ComputeFullRNorm, 78](#)
- [_braid_CopyFineToCoarse, 78](#)
- [_braid_CoreElt, 77](#)
- [_braid_CoreFcn, 77](#)
- [_braid_DeleteLastResidual, 78](#)
- [_braid_FASResidual, 78](#)
- [_braid_FAccess, 78](#)
- [_braid_FCRelax, 78](#)
- [_braid_FInterp, 78](#)
- [_braid_FRefine, 79](#)
- [_braid_FRestrict, 79](#)
- [_braid_GetBlockDistInterval, 79](#)
- [_braid_GetBlockDistProc, 79](#)
- [_braid_GetCFactor, 79](#)
- [_braid_GetDistribution, 79](#)
- [_braid_GetFullRNorm, 79](#)
- [_braid_GetInterval, 79](#)
- [_braid_GetProc, 80](#)
- [_braid_GetRNorm, 80](#)
- [_braid_GetUInit, 80](#)
- [_braid_GridClean, 80](#)
- [_braid_GridDestroy, 80](#)
- [_braid_GridElt, 77](#)
- [_braid_GridInit, 80](#)
- [_braid_InitGuess, 80](#)
- [_braid_InitHierarchy, 80](#)
- [_braid_IsCPoint, 77](#)
- [_braid_IsFPoint, 77](#)
- [_braid_MapCoarseToFine, 77](#)
- [_braid_MapFineToCoarse, 77](#)
- [_braid_NextCPoint, 77](#)
- [_braid_PrintSpatialNorms, 80](#)
- [_braid_Refine, 80](#)
- [_braid_RefineBasic, 80](#)
- [_braid_Residual, 81](#)
- [_braid_SetFullRNorm, 81](#)
- [_braid_SetRNorm, 81](#)
- [_braid_Step, 81](#)
- [_braid_UCommInit, 81](#)
- [_braid_UCommInitBasic, 81](#)
- [_braid_UCommInitF, 81](#)
- [_braid_UCommWait, 81](#)
- [_braid_UGetIndex, 81](#)
- [_braid_UGetVector, 81](#)
- [_braid_UGetVectorRef, 81](#)
- [_braid_USetVector, 82](#)
- [_braid_USetVectorRef, 82](#)
- [_braid_printfile, 82](#)
- [_braid_AccessStatus, 63](#)
 - [done, 63](#)
 - [gupper, 63](#)
 - [iter, 63](#)
 - [level, 63](#)
 - [nrefine, 63](#)
 - [rnorm, 63](#)
 - [t, 63](#)
 - [wrapper_test, 63](#)
- [_braid_AccessStatusDestroy](#)
 - [XBraid status routines, 47](#)
- [_braid_AccessStatusInit](#)
 - [XBraid status routines, 47](#)
- [_braid_AccessVector](#)
 - [_braid.h, 77](#)
- [_braid_CTAlloc](#)
 - [braid_defs.h, 85](#)
- [_braid_Coarsen](#)
 - [_braid.h, 77](#)
- [_braid_CoarsenRefStatus, 64](#)
 - [c_tprior, 64](#)
 - [c_tstop, 64](#)
 - [f_tprior, 64](#)
 - [f_tstop, 64](#)
 - [gupper, 64](#)
 - [level, 64](#)
 - [nrefine, 64](#)
 - [tstart, 64](#)
- [_braid_CoarsenRefStatusDestroy](#)
 - [XBraid status routines, 49](#)
- [_braid_CoarsenRefStatusInit](#)
 - [XBraid status routines, 49](#)
- [_braid_CommHandle, 65](#)
 - [buffer, 65](#)
 - [num_requests, 65](#)
 - [request_type, 65](#)
 - [requests, 65](#)
 - [status, 65](#)
 - [vector_ptr, 65](#)
- [_braid_CommHandleElt](#)
 - [_braid.h, 77](#)
- [_braid_CommRecvInit](#)
 - [_braid.h, 77](#)
- [_braid_CommSendInit](#)
 - [_braid.h, 78](#)
- [_braid_CommWait](#)
 - [_braid.h, 78](#)

- [_braid_ComputeFullIRNorm](#)
 - [_braid.h, 78](#)
- [_braid_CopyFineToCoarse](#)
 - [_braid.h, 78](#)
- [_braid_Core, 66](#)
 - [access, 67](#)
 - [access_level, 67](#)
 - [app, 67](#)
 - [astatus, 67](#)
 - [bufpack, 67](#)
 - [bufsize, 67](#)
 - [bufunpack, 67](#)
 - [cfactors, 67](#)
 - [cfdefault, 67](#)
 - [clone, 68](#)
 - [comm, 68](#)
 - [comm_world, 68](#)
 - [cstatus, 68](#)
 - [fmg, 68](#)
 - [free, 68](#)
 - [full_rnorm0, 68](#)
 - [full_rnorm_res, 68](#)
 - [full_rnorms, 68](#)
 - [globaltime, 68](#)
 - [grids, 68](#)
 - [gupper, 68](#)
 - [init, 68](#)
 - [localtime, 68](#)
 - [max_iter, 68](#)
 - [max_levels, 69](#)
 - [max_refinements, 69](#)
 - [min_coarse, 69](#)
 - [myid, 69](#)
 - [myid_world, 69](#)
 - [nfmfg, 69](#)
 - [nfmfg_Vcyc, 69](#)
 - [niter, 69](#)
 - [nlevels, 69](#)
 - [nrdefault, 69](#)
 - [nrefine, 69](#)
 - [nrels, 69](#)
 - [ntime, 69](#)
 - [print_level, 69](#)
 - [refine, 70](#)
 - [residual, 70](#)
 - [rfactors, 70](#)
 - [rnorm0, 70](#)
 - [rnorms, 70](#)
 - [rstopped, 70](#)
 - [rtol, 70](#)
 - [scoarsen, 70](#)
 - [seq_soln, 70](#)
 - [skip, 70](#)
 - [spatialnorm, 70](#)
 - [srefine, 70](#)
 - [sstatus, 70](#)
 - [step, 70](#)
 - [storage, 71](#)
 - [sum, 71](#)
 - [tnorm, 71](#)
 - [tnorm_a, 71](#)
 - [tol, 71](#)
 - [tpoints_cutoff, 71](#)
 - [tstart, 71](#)
 - [tstop, 71](#)
- [_braid_CoreElt](#)
 - [_braid.h, 77](#)
- [_braid_CoreFcn](#)
 - [_braid.h, 77](#)
- [_braid_DeleteLastResidual](#)
 - [_braid.h, 78](#)
- [_braid_Error](#)
 - [braid_defs.h, 85](#)
- [_braid_ErrorHandler](#)
 - [braid_defs.h, 86](#)
- [_braid_ErrorInArg](#)
 - [braid_defs.h, 85](#)
- [_braid_FASResidual](#)
 - [_braid.h, 78](#)
- [_braid_FAccess](#)
 - [_braid.h, 78](#)
- [_braid_FCRelax](#)
 - [_braid.h, 78](#)
- [_braid_FInterp](#)
 - [_braid.h, 78](#)
- [_braid_FRefine](#)
 - [_braid.h, 79](#)
- [_braid_FRestrict](#)
 - [_braid.h, 79](#)
- [_braid_GetBlockDistInterval](#)
 - [_braid.h, 79](#)
- [_braid_GetBlockDistProc](#)
 - [_braid.h, 79](#)
- [_braid_GetCFactor](#)
 - [_braid.h, 79](#)
- [_braid_GetDistribution](#)
 - [_braid.h, 79](#)
- [_braid_GetFullIRNorm](#)
 - [_braid.h, 79](#)
- [_braid_GetInterval](#)
 - [_braid.h, 79](#)
- [_braid_GetProc](#)
 - [_braid.h, 80](#)
- [_braid_GetRNorm](#)
 - [_braid.h, 80](#)
- [_braid_GetUInit](#)
 - [_braid.h, 80](#)
- [_braid_Grid, 71](#)

cfactor, [72](#)
 clower, [72](#)
 cupper, [72](#)
 fa, [72](#)
 fa_alloc, [72](#)
 gupper, [72](#)
 ilower, [72](#)
 iupper, [72](#)
 level, [72](#)
 ncpoints, [72](#)
 nupoints, [73](#)
 recv_handle, [73](#)
 recv_index, [73](#)
 send_handle, [73](#)
 send_index, [73](#)
 ta, [73](#)
 ta_alloc, [73](#)
 ua, [73](#)
 ua_alloc, [73](#)
 va, [73](#)
 va_alloc, [73](#)
 _braid_GridClean
 _braid.h, [80](#)
 _braid_GridDestroy
 _braid.h, [80](#)
 _braid_GridElt
 _braid.h, [77](#)
 _braid_GridInit
 _braid.h, [80](#)
 _braid_InitGuess
 _braid.h, [80](#)
 _braid_InitHierarchy
 _braid.h, [80](#)
 _braid_IsCPoint
 _braid.h, [77](#)
 _braid_IsFPoint
 _braid.h, [77](#)
 _braid_MapCoarseToFine
 _braid.h, [77](#)
 _braid_MapFineToCoarse
 _braid.h, [77](#)
 _braid_NextCPoint
 _braid.h, [77](#)
 _braid_PrintSpatialNorms
 _braid.h, [80](#)
 _braid_Refine
 _braid.h, [80](#)
 _braid_RefineBasic
 _braid.h, [80](#)
 _braid_Residual
 _braid.h, [81](#)
 _braid_SetFullRNorm
 _braid.h, [81](#)
 _braid_SetRNorm
 _braid.h, [81](#)
 XBraid status routines, [47](#)
 _braid_Step
 _braid.h, [81](#)
 _braid_StepStatus, [73](#)
 gupper, [74](#)
 iter, [74](#)
 level, [74](#)
 nrefine, [74](#)
 old_fine_tolx, [74](#)
 rfactor, [74](#)
 rnorms, [74](#)
 tight_fine_tolx, [74](#)
 tol, [74](#)
 tstart, [74](#)
 tstop, [75](#)
 _braid_StepStatusDestroy
 XBraid status routines, [49](#)
 _braid_StepStatusInit
 XBraid status routines, [49](#)
 _braid_TAlloc
 braid_defs.h, [85](#)
 _braid_TFree
 braid_defs.h, [85](#)
 _braid_TReAlloc
 braid_defs.h, [85](#)
 _braid_UCommInit
 _braid.h, [81](#)
 _braid_UCommInitBasic
 _braid.h, [81](#)
 _braid_UCommInitF
 _braid.h, [81](#)
 _braid_UCommWait
 _braid.h, [81](#)
 _braid_UGetIndex
 _braid.h, [81](#)
 _braid_UGetVector
 _braid.h, [81](#)
 _braid_UGetVectorRef
 _braid.h, [81](#)
 _braid_USetVector
 _braid.h, [82](#)
 _braid_USetVectorRef
 _braid.h, [82](#)
 _braid_error_flag
 braid_defs.h, [86](#)
 _braid_max
 braid_defs.h, [85](#)
 _braid_min
 braid_defs.h, [85](#)
 _braid_printfile
 _braid.h, [82](#)

- access
 - _braid_Core, 67
- access_level
 - _braid_Core, 67
- app
 - _braid_Core, 67
- astatus
 - _braid_Core, 67
- braid.h, 82
 - braid_ERROR_ARG, 84
 - braid_ERROR_GENERIC, 84
 - braid_ERROR_MEMORY, 84
 - braid_FMANGLE, 84
 - braid_Fortran_Residual, 84
 - braid_Fortran_SpatialCoarsen, 84
 - braid_INVALID_RNORM, 84
- braid_AccessStatus
 - XBraid status routines, 47
- braid_AccessStatusGetDone
 - XBraid status routines, 50
- braid_AccessStatusGetIter
 - XBraid status routines, 50
- braid_AccessStatusGetLevel
 - XBraid status routines, 50
- braid_AccessStatusGetNRrefine
 - XBraid status routines, 50
- braid_AccessStatusGetNTPoints
 - XBraid status routines, 50
- braid_AccessStatusGetResidual
 - XBraid status routines, 50
- braid_AccessStatusGetT
 - XBraid status routines, 51
- braid_AccessStatusGetTILD
 - XBraid status routines, 51
- braid_AccessStatusGetWrapperTest
 - XBraid status routines, 51
- braid_App
 - User-written routines, 32
- braid_CoarsenRefStatus
 - XBraid status routines, 47
- braid_CoarsenRefStatusGetCTprior
 - XBraid status routines, 51
- braid_CoarsenRefStatusGetCTstop
 - XBraid status routines, 51
- braid_CoarsenRefStatusGetFTprior
 - XBraid status routines, 52
- braid_CoarsenRefStatusGetFTstop
 - XBraid status routines, 52
- braid_CoarsenRefStatusGetLevel
 - XBraid status routines, 52
- braid_CoarsenRefStatusGetNRrefine
 - XBraid status routines, 52
- braid_CoarsenRefStatusGetNTPoints
 - XBraid status routines, 52
- braid_Destroy
 - General Interface routines, 37
- braid_Drive
 - General Interface routines, 37
- braid_ERROR_ARG
 - braid.h, 84
- braid_ERROR_GENERIC
 - braid.h, 84
- braid_ERROR_MEMORY
 - braid.h, 84
- braid_FMANGLE
 - braid.h, 84
- braid_Fortran_Residual
 - braid.h, 84
- braid_Fortran_SpatialCoarsen
 - braid.h, 84
- braid_GetNLevels
 - General Interface routines, 37
- braid_GetNumIter
 - General Interface routines, 37
- braid_GetRNorms
 - General Interface routines, 37
- braid_GetSpatialAccuracy
 - General Interface routines, 38
- braid_INVALID_RNORM
 - braid.h, 84
- braid_Init
 - General Interface routines, 38
- braid_Int
 - braid_defs.h, 85
- braid_Int_Max
 - braid_defs.h, 85
- braid_Int_Min
 - braid_defs.h, 85
- braid_MPI_INT
 - braid_defs.h, 85
- braid_MPI_REAL
 - braid_defs.h, 85
- braid_PrintStats
 - General Interface routines, 39
- braid_PtFcnAccess
 - User-written routines, 32
- braid_PtFcnBufPack
 - User-written routines, 32
- braid_PtFcnBufSize
 - User-written routines, 32
- braid_PtFcnBufUnpack
 - User-written routines, 32

- User-written routines, [32](#)
- `braid_PtFcnClone`
 - User-written routines, [32](#)
- `braid_PtFcnFree`
 - User-written routines, [32](#)
- `braid_PtFcnInit`
 - User-written routines, [32](#)
- `braid_PtFcnResidual`
 - User-written routines, [32](#)
- `braid_PtFcnSCoarsen`
 - User-written routines, [33](#)
- `braid_PtFcnSRefine`
 - User-written routines, [33](#)
- `braid_PtFcnSpatialNorm`
 - User-written routines, [33](#)
- `braid_PtFcnStep`
 - User-written routines, [33](#)
- `braid_PtFcnSum`
 - User-written routines, [33](#)
- `braid_Real`
 - `braid_defs.h`, [85](#)
- `braid_SetAbsTol`
 - General Interface routines, [39](#)
- `braid_SetAccessLevel`
 - General Interface routines, [39](#)
- `braid_SetCFactor`
 - General Interface routines, [39](#)
- `braid_SetDefaultPrintFile`
 - General Interface routines, [40](#)
- `braid_SetFMG`
 - General Interface routines, [40](#)
- `braid_SetFullRNormRes`
 - General Interface routines, [40](#)
- `braid_SetMaxIter`
 - General Interface routines, [40](#)
- `braid_SetMaxLevels`
 - General Interface routines, [40](#)
- `braid_SetMaxRefinements`
 - General Interface routines, [40](#)
- `braid_SetMinCoarse`
 - General Interface routines, [41](#)
- `braid_SetNFMG`
 - General Interface routines, [41](#)
- `braid_SetNFMGVcyc`
 - General Interface routines, [41](#)
- `braid_SetNRelax`
 - General Interface routines, [41](#)
- `braid_SetPrintFile`
 - General Interface routines, [41](#)
- `braid_SetPrintLevel`
 - General Interface routines, [41](#)
- `braid_SetRefine`
 - General Interface routines, [42](#)
- `braid_SetRelTol`
 - General Interface routines, [42](#)
- `braid_SetResidual`
 - General Interface routines, [42](#)
- `braid_SetSeqSoln`
 - General Interface routines, [42](#)
- `braid_SetSkip`
 - General Interface routines, [43](#)
- `braid_SetSpatialCoarsen`
 - General Interface routines, [43](#)
- `braid_SetSpatialRefine`
 - General Interface routines, [43](#)
- `braid_SetStorage`
 - General Interface routines, [43](#)
- `braid_SetTPointsCutoff`
 - General Interface routines, [45](#)
- `braid_SetTemporalNorm`
 - General Interface routines, [43](#)
- `braid_SplitCommworld`
 - General Interface routines, [45](#)
- `braid_StepStatus`
 - XBraid status routines, [47](#)
- `braid_StepStatusGetIter`
 - XBraid status routines, [54](#)
- `braid_StepStatusGetLevel`
 - XBraid status routines, [54](#)
- `braid_StepStatusGetNRefine`
 - XBraid status routines, [54](#)
- `braid_StepStatusGetNTPoints`
 - XBraid status routines, [54](#)
- `braid_StepStatusGetOldFineTolx`
 - XBraid status routines, [54](#)
- `braid_StepStatusGetRNorms`
 - XBraid status routines, [56](#)
- `braid_StepStatusGetTol`
 - XBraid status routines, [56](#)
- `braid_StepStatusGetTstart`
 - XBraid status routines, [56](#)
- `braid_StepStatusGetTstartTstop`
 - XBraid status routines, [56](#)
- `braid_StepStatusGetTstop`
 - XBraid status routines, [56](#)
- `braid_StepStatusSetOldFineTolx`
 - XBraid status routines, [57](#)
- `braid_StepStatusSetRFactor`
 - XBraid status routines, [57](#)
- `braid_StepStatusSetTightFineTolx`
 - XBraid status routines, [57](#)
- `braid_TestAll`
 - XBraid test routines, [58](#)
- `braid_TestBuf`
 - XBraid test routines, [59](#)
- `braid_TestClone`
 - XBraid test routines, [59](#)
- `braid_TestCoarsenRefine`

- XBraid test routines, 60
- braid_TestInitAccess
 - XBraid test routines, 60
- braid_TestSpatialNorm
 - XBraid test routines, 61
- braid_TestSum
 - XBraid test routines, 61
- braid_Vector
 - User-written routines, 33
- braid_defs.h, 84
 - _braid_CTAlloc, 85
 - _braid_Error, 85
 - _braid_ErrorHandler, 86
 - _braid_ErrorInArg, 85
 - _braid_TAlloc, 85
 - _braid_TFree, 85
 - _braid_TReAlloc, 85
 - _braid_error_flag, 86
 - _braid_max, 85
 - _braid_min, 85
 - braid_Int, 85
 - braid_Int_Max, 85
 - braid_Int_Min, 85
 - braid_MPI_INT, 85
 - braid_MPI_REAL, 85
 - braid_Real, 85
 - braid_isnan, 85
- braid_isnan
 - braid_defs.h, 85
- braid_status.h, 86
- braid_test.h, 87
- buffer
 - _braid_CommHandle, 65
- bufpack
 - _braid_Core, 67
- bufsize
 - _braid_Core, 67
- bufunpack
 - _braid_Core, 67
- c_tprior
 - _braid_CoarsenRefStatus, 64
- c_tstop
 - _braid_CoarsenRefStatus, 64
- cfactor
 - _braid_Grid, 72
- cfactors
 - _braid_Core, 67
- cfdefault
 - _braid_Core, 67
- clone
 - _braid_Core, 68
- clower
 - _braid_Grid, 72
- comm
 - _braid_Core, 68
- comm_world
 - _braid_Core, 68
- cstatus
 - _braid_Core, 68
- cupper
 - _braid_Grid, 72
- done
 - _braid_AccessStatus, 63
- f_tprior
 - _braid_CoarsenRefStatus, 64
- f_tstop
 - _braid_CoarsenRefStatus, 64
- fa
 - _braid_Grid, 72
- fa_alloc
 - _braid_Grid, 72
- fmg
 - _braid_Core, 68
- free
 - _braid_Core, 68
- full_rnorm0
 - _braid_Core, 68
- full_rnorm_res
 - _braid_Core, 68
- full_rnorms
 - _braid_Core, 68
- General Interface routines, 36
 - braid_Core, 37
 - braid_Destroy, 37
 - braid_Drive, 37
 - braid_GetNLevels, 37
 - braid_GetNumIter, 37
 - braid_GetRNorms, 37
 - braid_GetSpatialAccuracy, 38
 - braid_Init, 38
 - braid_PrintStats, 39
 - braid_SetAbsTol, 39
 - braid_SetAccessLevel, 39
 - braid_SetCFactor, 39
 - braid_SetDefaultPrintFile, 40
 - braid_SetFMG, 40
 - braid_SetFullIRNormRes, 40
 - braid_SetMaxIter, 40
 - braid_SetMaxLevels, 40
 - braid_SetMaxRefinements, 40
 - braid_SetMinCoarse, 41
 - braid_SetNFMG, 41
 - braid_SetNFMGVcyc, 41
 - braid_SetNRelax, 41
 - braid_SetPrintFile, 41

- braid_SetPrintLevel, 41
- braid_SetRefine, 42
- braid_SetRelTol, 42
- braid_SetResidual, 42
- braid_SetSeqSoln, 42
- braid_SetSkip, 43
- braid_SetSpatialCoarsen, 43
- braid_SetSpatialRefine, 43
- braid_SetStorage, 43
- braid_SetTPointsCutoff, 45
- braid_SetTemporalNorm, 43
- braid_SplitCommworld, 45
- globaltime
 - _braid_Core, 68
- grids
 - _braid_Core, 68
- gupper
 - _braid_AccessStatus, 63
 - _braid_CoarsenRefStatus, 64
 - _braid_Core, 68
 - _braid_Grid, 72
 - _braid_StepStatus, 74
- ilower
 - _braid_Grid, 72
- init
 - _braid_Core, 68
- iter
 - _braid_AccessStatus, 63
 - _braid_StepStatus, 74
- iupper
 - _braid_Grid, 72
- level
 - _braid_AccessStatus, 63
 - _braid_CoarsenRefStatus, 64
 - _braid_Grid, 72
 - _braid_StepStatus, 74
- localtime
 - _braid_Core, 68
- max_iter
 - _braid_Core, 68
- max_levels
 - _braid_Core, 69
- max_refinements
 - _braid_Core, 69
- min_coarse
 - _braid_Core, 69
- myid
 - _braid_Core, 69
- myid_world
 - _braid_Core, 69
- ncpoints
 - _braid_Grid, 72
- nfmfg
 - _braid_Core, 69
- nfmfg_Vcyc
 - _braid_Core, 69
- niter
 - _braid_Core, 69
- nlevels
 - _braid_Core, 69
- nrdefault
 - _braid_Core, 69
- nrefine
 - _braid_AccessStatus, 63
 - _braid_CoarsenRefStatus, 64
 - _braid_Core, 69
 - _braid_StepStatus, 74
- nrels
 - _braid_Core, 69
- ntime
 - _braid_Core, 69
- num_requests
 - _braid_CommHandle, 65
- nupoints
 - _braid_Grid, 73
- old_fine_tolx
 - _braid_StepStatus, 74
- print_level
 - _braid_Core, 69
- recv_handle
 - _braid_Grid, 73
- recv_index
 - _braid_Grid, 73
- refine
 - _braid_Core, 70
- request_type
 - _braid_CommHandle, 65
- requests
 - _braid_CommHandle, 65
- residual
 - _braid_Core, 70
- rfactor
 - _braid_StepStatus, 74
- rfactors
 - _braid_Core, 70
- rnorm
 - _braid_AccessStatus, 63
- rnorm0
 - _braid_Core, 70
- rnorms
 - _braid_Core, 70
 - _braid_StepStatus, 74
- rstopped

- [_braid_Core, 70](#)
- rtol
 - [_braid_Core, 70](#)
- scoarsen
 - [_braid_Core, 70](#)
- send_handle
 - [_braid_Grid, 73](#)
- send_index
 - [_braid_Grid, 73](#)
- seq_soln
 - [_braid_Core, 70](#)
- skip
 - [_braid_Core, 70](#)
- spatialnorm
 - [_braid_Core, 70](#)
- srefine
 - [_braid_Core, 70](#)
- sstatus
 - [_braid_Core, 70](#)
- status
 - [_braid_CommHandle, 65](#)
- step
 - [_braid_Core, 70](#)
- storage
 - [_braid_Core, 71](#)
- sum
 - [_braid_Core, 71](#)
- t
 - [_braid_AccessStatus, 63](#)
- ta
 - [_braid_Grid, 73](#)
- ta_alloc
 - [_braid_Grid, 73](#)
- tight_fine_tolx
 - [_braid_StepStatus, 74](#)
- tnorm
 - [_braid_Core, 71](#)
- tnorm_a
 - [_braid_Core, 71](#)
- tol
 - [_braid_Core, 71](#)
 - [_braid_StepStatus, 74](#)
- tpoints_cutoff
 - [_braid_Core, 71](#)
- tstart
 - [_braid_CoarsenRefStatus, 64](#)
 - [_braid_Core, 71](#)
 - [_braid_StepStatus, 74](#)
- tstop
 - [_braid_Core, 71](#)
 - [_braid_StepStatus, 75](#)
- ua
 - [_braid_Grid, 73](#)
- ua_alloc
 - [_braid_Grid, 73](#)
- User interface routines, [35](#)
- User-written routines, [31](#)
 - [braid_App, 32](#)
 - [braid_PtFcnAccess, 32](#)
 - [braid_PtFcnBufPack, 32](#)
 - [braid_PtFcnBufSize, 32](#)
 - [braid_PtFcnBufUnpack, 32](#)
 - [braid_PtFcnClone, 32](#)
 - [braid_PtFcnFree, 32](#)
 - [braid_PtFcnInit, 32](#)
 - [braid_PtFcnResidual, 32](#)
 - [braid_PtFcnSCoarsen, 33](#)
 - [braid_PtFcnSRefine, 33](#)
 - [braid_PtFcnSpatialNorm, 33](#)
 - [braid_PtFcnStep, 33](#)
 - [braid_PtFcnSum, 33](#)
 - [braid_Vector, 33](#)
- va
 - [_braid_Grid, 73](#)
- va_alloc
 - [_braid_Grid, 73](#)
- vector_ptr
 - [_braid_CommHandle, 65](#)
- wrapper_test
 - [_braid_AccessStatus, 63](#)
- XBraid status routines, [46](#)
 - [_braid_AccessStatusDestroy, 47](#)
 - [_braid_AccessStatusInit, 47](#)
 - [_braid_CoarsenRefStatusDestroy, 49](#)
 - [_braid_CoarsenRefStatusInit, 49](#)
 - [_braid_StatusElt, 47](#)
 - [_braid_StepStatusDestroy, 49](#)
 - [_braid_StepStatusInit, 49](#)
 - [braid_AccessStatus, 47](#)
 - [braid_AccessStatusGetDone, 50](#)
 - [braid_AccessStatusGetIter, 50](#)
 - [braid_AccessStatusGetLevel, 50](#)
 - [braid_AccessStatusGetNRefine, 50](#)
 - [braid_AccessStatusGetNTPoints, 50](#)
 - [braid_AccessStatusGetResidual, 50](#)
 - [braid_AccessStatusGetT, 51](#)
 - [braid_AccessStatusGetTILD, 51](#)
 - [braid_AccessStatusGetWrapperTest, 51](#)
 - [braid_CoarsenRefStatus, 47](#)
 - [braid_CoarsenRefStatusGetCTprior, 51](#)
 - [braid_CoarsenRefStatusGetCTstop, 51](#)
 - [braid_CoarsenRefStatusGetFTPrior, 52](#)
 - [braid_CoarsenRefStatusGetFTstop, 52](#)
 - [braid_CoarsenRefStatusGetLevel, 52](#)

- braid_CoarsenRefStatusGetNRefine, [52](#)
- braid_CoarsenRefStatusGetNTPoints, [52](#)
- braid_CoarsenRefStatusGetTpriorTstop, [52](#)
- braid_CoarsenRefStatusGetTstart, [54](#)
- braid_StepStatus, [47](#)
- braid_StepStatusGetIter, [54](#)
- braid_StepStatusGetLevel, [54](#)
- braid_StepStatusGetNRefine, [54](#)
- braid_StepStatusGetNTPoints, [54](#)
- braid_StepStatusGetOldFineTolx, [54](#)
- braid_StepStatusGetRNorms, [56](#)
- braid_StepStatusGetTol, [56](#)
- braid_StepStatusGetTstart, [56](#)
- braid_StepStatusGetTstartTstop, [56](#)
- braid_StepStatusGetTstop, [56](#)
- braid_StepStatusSetOldFineTolx, [57](#)
- braid_StepStatusSetRFactor, [57](#)
- braid_StepStatusSetTightFineTolx, [57](#)
- XBraid test routines, [58](#)
 - braid_TestAll, [58](#)
 - braid_TestBuf, [59](#)
 - braid_TestClone, [59](#)
 - braid_TestCoarsenRefine, [60](#)
 - braid_TestInitAccess, [60](#)
 - braid_TestSpatialNorm, [61](#)
 - braid_TestSum, [61](#)