

GridPACK™ Overview

**Bruce Palmer, William Perkins, Kevin Glass, Yousu Chen, Shuangshuang Jin, Ruisheng
Diao, Mark Rice, Stephen Elbert, Zhenyu (Henry) Huang**

Table of Contents

Introduction	3
Building GridPACK™ Applications	4
GridPACK™ Framework Components	7
Network Module	8
Math Module	15
Network Components	24
Factories	33
Mapper Module	35
Parser Module	38
Serial IO Module	40
Configuration Module	43
Developing Applications	45
Advanced Functionality	62
Communicators	63
Task Manager	67
Timers	69
Exceptions	71
Hash Distribution Module	72
Generalize Matrix-Vector Interface	74
Fortran 2003 Interface	79

Introduction

The objective of the GridPACK™ toolkit project is to develop a framework to support the rapid development of power grid applications capable of running on high performance computing architectures (HPC) with high levels of performance and scalability. The toolkit will allow power system engineers to focus on developing working applications from their models without getting bogged down in the details of decomposing the computation across multiple processors, managing data transfers between processors, working out index transformations between power grid networks and the matrices generated by different power applications, and managing input and output. GridPACK™ is being designed to encapsulate as much of the book-keeping required to set up HPC applications as possible in high-level programming abstractions that allow developers to concentrate on the physics and mathematics of their problems.

This report will summarize the overall design of the GridPACK™ framework. The initial focus of the GridPACK™ design analysis was to target four power grid applications and to identify common features that span multiple applications as candidates for inclusion in a framework. This analysis included a breakdown of the application into phases and identification within each phase of the functionality required to complete them. The four applications originally targeted within this project were power flow simulations, contingency analysis, state estimation and dynamic simulation. The remainder of this document will describe the functionality incorporated into the GridPACK™ framework to support multiple power grid applications. The framework will continue to evolve as more real-world experience can be incorporated into the design process but many base classes that have already been identified that are capable of supporting a range of applications.

Four power grid applications were targeted for initial implementation within the GridPack framework. These consisted of

- 1) Powerflow simulations of the electric grid
- 2) Contingency analysis of the electric grid
- 3) State estimation based on electric grid measurements
- 4) Dynamic simulations of the electric grid

Based on these applications, several cross-cutting functionalities were identified that could be used to support multiple applications. These include modules to support

- 1) Network topology and behavior. The network topology is the starting point for any power grid analysis. The topology defines the initial network model and is the connection point between the physical problem definition in terms of buses and branches and the solution method, which is usually expressed in terms of matrices and vectors.
- 2) Network components and their properties (e.g. bus and branch models, measurements, etc.). Grid components are the objects associated with the buses and branches of the power grid network. Along with the network topology itself, these define the physical

system being modeled and in some cases the analysis that is to be performed. Bus and branch components can be differentiated into things like generators, loads, grounds, lines, transformers, measurements, etc. and depending on how they are defined and the level of detail incorporated into them, they define different power grid systems and analyses. The behavior of buses and branches can depend on the properties of branches or buses that are directly attached to them, e.g. figuring out the contribution of a particular bus to the solution procedure may require that properties of the attached branches are made available to the bus. The necessity for exchanging this data is built into the framework. Furthermore, these data exchanges must also be accounted for in a parallel computing context, since the grid component from which data is required may be located on a different processor.

- 3) Linear algebra and solvers. Basic algebraic objects, such as distributed matrices and vectors, are a core part of the solution algorithms required by power grid analyses. Most solution algorithms are dominated by sparse matrices but a few, such as Kalman filter analyses, require dense matrices. Vectors are typically dense. There exists a rich set of libraries for constructing distributed matrices and vectors and these are coupled to preconditioner and solver libraries. GridPACK™ can leverage this work heavily by creating wrappers within the framework to these libraries can be used in solution algorithms. Wrapping these libraries instead of using them directly will have the advantage that creating algebraic objects can be simplified somewhat for power grid applications but more importantly, it will allow framework developers to investigate new solver and algebraic libraries seamlessly, without disrupting other parts of the code.
- 4) Mapping between network and algebraic objects. The physical properties of power grid systems are defined by networks and the properties of the network components but the equations describing the networks are algebraic in nature. The mappings between the physical networks and the algebraic equations depend on the indexing scheme used to describe the network and the number of parameters in the network components that appear in the equations. Constructing a map between network parameters and their corresponding locations in a matrix or vector can be complicated and error prone. Fortunately, much of this work can be automated and developers can focus much more on developing code to evaluate individual matrix elements without worrying about where to locate them in the matrix. This can considerably simplify coding.

The elements described above have all been incorporated into the GridPACK™ modules described above. More details about these modules and their interactions are provided in the remainder of this document.

Building GridPACK™ Applications

GridPACK™ comes with several applications that are included in the main distribution. These currently include power flow, contingency analysis, dynamic simulation and state estimation applications as well as some non-power grid examples that illustrate features of the framework.

These applications are automatically built whenever the full GridPACK™ distribution is built. There is extensive documentation on how to build GridPACK™ and the libraries on which it depends on the website located at <https://gridpack.org>.

For applications developed outside the GridPACK™ distribution, the build process is fairly simple if you are using CMake (you will need to have CMake installed on your system to build GridPACK™ so using CMake for your application build should be a straightforward extension). For a CMake build, you need to create a CMakeLists.txt file in the same directory that includes your application files. A template for the CMakeLists.txt file is

```
1 cmake_minimum_required(VERSION 2.6.4)
2
3 if (NOT GRIDPACK_DIR)
4     set(GRIDPACK_DIR /HOME/gridpack-install
5         CACHE PATH "GridPACK installation directory")
6 endif()
7
8 include("${GRIDPACK_DIR}/lib/GridPACK.cmake")
9
10 project(MyProject)
11
12 enable_language(CXX)
13
14 gridpack_setup()
15
16 add_definitions(${GRIDPACK_DEFINITIONS})
17 include_directories(BEFORE ${CMAKE_CURRENT_SOURCE_DIR})
18 include_directories(BEFORE ${GRIDPACK_INCLUDE_DIRS})
19
20 add_executable(myapp.x
21     myapp_main.cpp
22     mayapp_driver.cpp
23     myapp_file1.cpp
24     myapp_file2.cpp
25 )
26 target_link_libraries(myapp.x ${GRIDPACK_LIBS})
27
28 add_custom_target(myapp.input
29
30     COMMAND ${CMAKE_COMMAND} -E copy
31     ${CMAKE_CURRENT_SOURCE_DIR}/input.xml
```

```

32  ${CMAKE_CURRENT_BINARY_DIR}
33
34  COMMAND ${CMAKE_COMMAND} -E copy
35  ${CMAKE_CURRENT_SOURCE_DIR}/myapp_test.raw
36  ${CMAKE_CURRENT_BINARY_DIR}
37
38  DEPENDS
39  ${CMAKE_CURRENT_SOURCE_DIR}/input.xml
40  ${CMAKE_CURRENT_SOURCE_DIR}/myapp_test.raw
41 )
42 add_dependencies(myapp.x myapp.input)

```

Lines 1-6 check to see if the CMake installation is recent enough and also make sure that the **GRIDPACK_DIR** variable has been defined in the configuration step. If it hasn't, then the CMake will try and use a default value in **/HOME/gridpack-install**. However, this is unlikely to be successful, so it is better to define **GRIDPACK_DIR** when configuring your application. Line 8 picks up a file that is used by the application build to link to libraries and header files in the GridPACK™ build and line 10 can be used to assign a name to your application. Lines 12-18 can be included as is, if all application files are in the same directory as the CMakeLists.txt file. If other directories are used source and header files, then they can be included using the directives in lines 17 and 18.

Lines 20-25 define the name of the executable and all the source code files that are used in the application. The **add_executable** command on line 26 adds the executable **myapp.x** to the build. The arguments to this command consist of the name of the executable followed by the executable source files. There can be an arbitrary number of source files associated with any one executable. Note that the source files just consist of the user application source files, the framework files are handled automatically.

The remaining lines 28-42 are optional and can be used to automatically copy files from the application source file directory to the build directory. These could include example input files or external configuration files that are called by the code to set internal parameters. The **add_custom_target** command on line 28 defines a list of files and what should be done with them. In this example, the two files **input.xml** and **myapp_test.raw** are the files to be copied. The **COMMAND** line specifies the action (copy) and the next two lines specify the location of the file to be copied and its destination. The **DEPENDS** keyword (line 38) indicates that any time the **input.xml** or **myapp_test.raw** files are modified, they should be recopied to the build directory if make is invoked and the **add_dependencies** command (line 42) binds the custom target to the build of the executable.

GridPACK™ Framework Components

This section will describe the components that have been implemented so far and the functionality they support. It will start off with two components that directly support the major underlying data objects, the power grid network and its associated network components and matrices and vectors. Additional components are then built on top of these (or at least in conjunction with them). These include grid components that describe the physics of the different network models or analyses, grid component factories that initialize the grid components, mappers that convert the current state of the grid components into matrices and vectors, solvers that supply the preconditioner and solver functionality necessary to implement solution algorithms, input and output modules that allow developers to import and export data, and other utility modules that support standard code develop operations like timing, event logging, and error handling.

Many of these modules are constructed using libraries developed elsewhere so as to minimize framework development time. However, by wrapping them in interfaces geared towards power grid applications these libraries can be made easier to use by power grid engineers. The interfaces also make it possible in the future to exchange libraries for new or improved implementations of specific functionality without requiring application developers to rewrite their codes. This can significantly reduce the cost of introducing new technology into the framework. The software layers in the GridPACK™ framework are shown schematically in Figure 1.

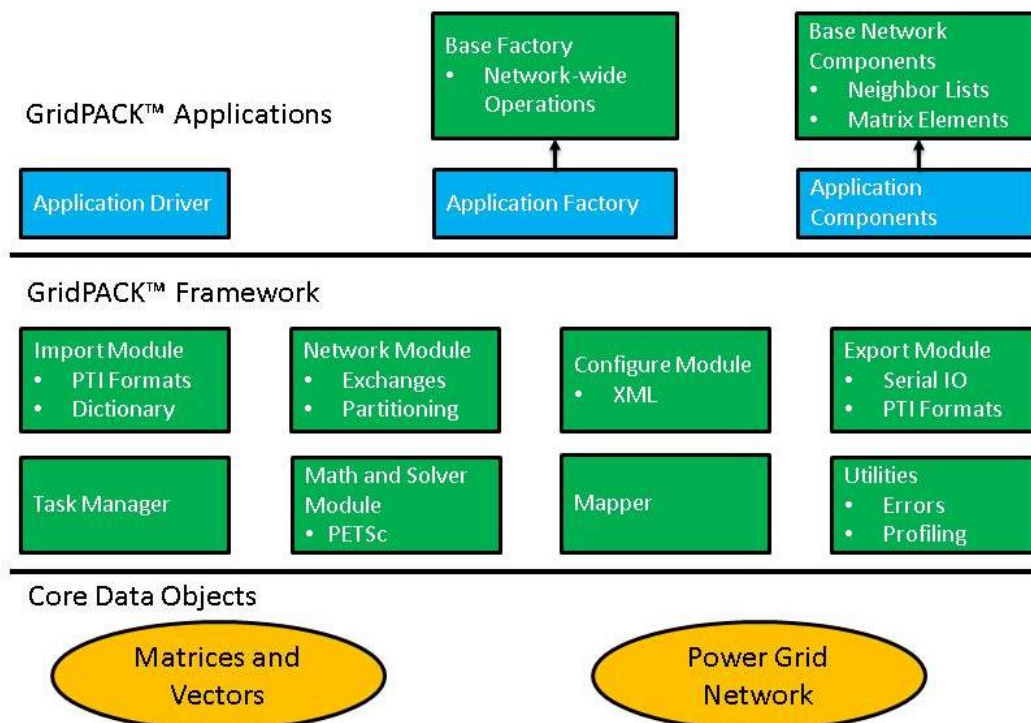


Figure 1. A schematic diagram of the GridPack framework software data stack.

Core framework components are described below. Before discussing the components themselves, some of the coding conventions and libraries used in GridPACK™ will be described.

Preliminaries: The GridPACK™ software uses a few coding conventions to help improve memory management and to minimize run-time errors. The first of these is to employ namespaces for all GridPACK modules. The entire GridPACK™ framework uses the **gridpack** namespace, individual modules within GridPACK™ are further delimited by their own namespaces. For example, the BaseNetwork class discussed in the next section resides in the **gridpack::network** namespace and other modules have similar delineations. The example applications included in the source code also have their own namespaces, but this is not a requirement for developing GridPACK™-based applications.

To help with memory management, many GridPACK™ functions return boost shared pointers instead of conventional C++ pointers. These can be converted to a conventional pointer using the **get()** command. We also recommend that pointers be converted using a **dynamic_cast** instead of conventional C-style cast.

Application files should include the **gridpack.hpp** header file. This can be done by adding the line

```
#include "gridpack/include/gridpack.hpp"
```

at the top of the application .hpp and/or .cpp files. This file contains definitions of all the GridPACK™ modules and their associated functions.

All matrices and vectors in GridPACK™ are currently complex. This is a consequence of building the underlying PETSc libraries to support complex numbers. Unfortunately, PETSc supports either complex numbers or real numbers but not both simultaneously. Work is underway to add support for both real and complex numbers at the GridPACK™ level. Complex numbers are represented in GridPACK™ as having type **ComplexType**. The real and imaginary parts of a complex number **x** of type **ComplexType** can be obtained using the functions **real(x)** and **imag(x)**.

Network Module

The network module is designed to represent the power grid has four major functions

- 1) The network is a container for the grid topology. The connectivity of the network is maintained by the network object and can be made available through requests to the network. The network also maintains the “ghost” status of locally held buses and

branches and determines whether a bus or branch is owned by a particular processor or represents a ghost image of a bus or branch owned by a neighboring processor.

- 2) The network topology can then be decorated with bus and branch objects that reflect the properties of the particular physical system under investigation. These bus and branch objects are written by the application developer and reflect the system being modeled and the analyses that need to be performed on it. Different applications will use different bus and branch implementations.
- 3) The network module is responsible for implementing update operations that can be used to fill in the value of ghost cell fields with current data from other processors. The update of ghost buses and ghost branches have been split into separate operations to give users flexibility in optimizing performance by minimizing the amount of data that needs to be communicated in the code.
- 4) The network contains the partitioner. The partitioner is embedded in the network module but represents a substantial technology in its own right. Partitioning is a key part of parallel application development. It represents the act of dividing up the problem so that each processor is left with approximately equal amounts of work and so that communication between processors (a major source of computational inefficiency in HPC programs) is minimized.

A network is illustrated schematically in Figure 2. Each bus and branch has an associated bus or branch object. The buses and branches are derived from base classes that specify certain functions that must be implemented by the application developer so that the network can interact with other GridPACK™ modules. In addition, the application can have functionality outside the base class that are unique to the particular application.






Figure 2. Schematic representation of a GridPACK™ network. The squares are branch objects and the circles are bus objects. Framework-specified interfaces are green and user supplied functionality is blue.

A major use of the partitioner is to rearrange the network in a form that is useful for computation immediately after it is read in from an external file. Typically, the information in the external file is not organized in a way that is necessarily optimal for computation, so the partitioner must redistribute data such that large connected blocks are all on the same processor. The partitioner is also responsible for adding the ghost buses and branches to the system.

Ghost buses and branches in a parallel program represent images of buses and branches that are owned by other processes. In order to carry out operations on buses and branches it is frequently necessary to gain access to data associated with attached buses and branches. The most efficient way to do this is to create copies of the buses and branches from other processors on each process so that all locally own objects are attached to these copies (ghosts). The ghost objects are then updated collectively with current information from their home processors at points in the computation. Updating all ghosts at once is almost always more efficient than accessing data from one remote bus or branch at a time.

The use of the partitioner to distribute the network between different processors and create ghost nodes and branches is illustrated in Figure 3. Figure 3(a) shows a simple network and Figures 3(b) and 3(c) show the result of distributing the network between two processors.






Figure 3. (a) a simple network (b) partition of network on processor 0 (b) partition of network on processor 1. Open circles indicate ghost buses and dotted lines indicate ghost branches.

Networks can be created using the templated base class **BaseNetwork<class Bus, class Branch>**, where **Bus** and **Branch** are application specific classes describing the properties of buses and branches in the network. The **BaseNetwork** class is defined within the **gridpack::network** namespace. In addition to the **Bus** and **Branch** classes, each bus and branch has an associated **DataCollection** object, which is described in more detail in the network components section. The **DataCollection** object is a collection of key-value pairs that acts as an intermediary between data that is read in from external configuration files and the bus and branch classes that define the network.

The **BaseNetwork** class contains a large number of methods, but only a relatively small number will be of interest to application developers. Most of the remaining methods are used primarily within other GridPACK™ modules to implement higher level capabilities. This document will focus on calls that are likely to be used by application developers.

The constructor for the network class is the function

BaseNetwork(const parallel::Communicator &comm)

The **Communicator** object can be used to define the set of processors over which the network is distributed. Communicators are discussed in more detail below. The network constructor creates an empty shell that does not contain any information about an actual network. The remainder of the network must be built up by adding buses and branches to it. Typically, buses

and branches are added by passing the network to a parser (see import module) which will create an initial version of the network. The constructor is paired with a corresponding destructor

~BaseNetwork()

that is called when the network object passes out of scope or is explicitly deleted by the user.

Two functions are available that return the number of buses or branches that are available on a process. This number includes both buses and branches that are held locally as well as any ghosts that may be located on the process.

int numBuses()

int numBranches()

There are also functions that will return the total number of buses or branches in the network. These numbers ignore ghost buses and ghost branches.

int totalBuses()

int totalBranches()

Buses and branches in the network can be identified using a local index that runs from 0 to the number of buses or branches on the process minus 1 (0-based indexing). For some calculations, it is necessary to identify one bus in the network as a reference bus. This bus is usually set when the network is created using an import parser. It can subsequently be identified using the function

int getReferenceBus()

If the reference bus is located on this processor (either as a local bus or a ghost) then this function returns the local index of the bus, otherwise it returns -1.

Ghost buses and branches are distinguished from locally owned buses and branches based on whether or not they are “active”. The two functions

bool getActiveBus(int idx)

bool getActiveBranch(int idx)

provide the active status of a bus or branch on a process. The index **idx** is a local index for the bus or branch.

Buses and branches are characterized by a number of different indices. One is the local index, already discussed above, but there are several others. Most of these are used internally by other parts of the framework but one index is of interest to application developers. This is the

“original” bus index. When the network is described in the input file, the buses are labeled with a (usually) positive integer. There are no requirements that these integers be consecutive, only that each bus has its own unique index. The value of this index can be recovered using the function

```
int getOriginalBusIndex(int idx)
```

The variable **idx** is the local index of the bus. Branches are usually described in terms of the original bus indices for the two buses at each end of the branch, so there is no corresponding function for branches. Instead, the procedure is to get the local indices of the two buses at each end of the branch and then get the corresponding original indices of the buses. This information is usually used for output.

It is frequently necessary to gain access to the objects associated with each bus or branch. The following four methods can be used to access these objects

```
boost::shared_ptr<Bus> getBus(int idx)
```

```
boost::shared_ptr<Branch> getBranch(int idx)
```

```
boost::shared_ptr<DataCollection> getBusData(int idx)
```

```
boost::shared_ptr<DataCollection> getBranchData(int idx)
```

The first two methods can be used to get Boost shared pointers to individual bus or branch objects indexed by local indices **idx**. The second two functions return pointers to the **DataCollection** objects associated with each bus or branch. These objects are usually used to initialize the bus and branch objects at the start of a calculation

```
void partition()
```

The partition function distributes the buses and branches across processors such that the connectivity to branches and buses on other processors is minimized. It is also responsible for adding ghost buses and branches to the network. This function should be called after the network is read in but before any other operations, such as setting up exchange buffers or creating neighbor lists have been performed.

Finally, two sets of functions are required in order to set up and execute data exchanges between buses and branches in a distributed network. These exchanges are used to move data from active components to ghost components residing on other processors. Before these functions can be called, the buffers in individual network components must be allocated. See the documentation below on network components and the network factory for more information on how to do this. Once the buffers are in place, bus and branch exchanges can be set up and executed with just a few calls. The functions

void initBusUpdate()

void initBranchUpdate()

are used to initialize the data structures inside the network object that manage data exchanges. Exchanges between buses and branches are handled separately, since not all applications will require exchanges between both sets of objects. The initialization routines are relatively complex and allocate several large internal data structures so they should not be called if there is no need to exchange data between buses or branches.

After the updates have been initialized, it is possible to execute a data exchange at any point in the code by calling the functions

void updateBuses()

void updateBranches()

These functions will cause data to be exchanged between active buses and branches and their corresponding ghosts buses and branches located on other processors.

The **BaseNetwork** methods described above are only a subset of the total functionality available but they represent most of the functionality that a typical developer would use. The remaining functions are primarily used to implement other parts of the GridPACK™ framework but are generally not required by people writing applications. More information on how these functions are used in practice can be found in the section on GridPACK™ factories.

Math Module

The math module is used to provide support for distributed matrices and vectors as well as linear solvers, non-linear solvers, and preconditioners. Once created, matrices can be treated as opaque objects and manipulated using a high level syntax that would be comparable to writing Matlab code. The distributed matrix and vector data structures themselves are based on existing solver libraries and represent relatively lightweight wrappers on existing code. The current math module is built on the PETSc library but other libraries, such as Hypre and Trilinos could be used instead to implement the math module.

The main functionality associated with the math module is the ability to instantiate new matrices and vectors, add individual matrix and vector elements (and their values) to the matrix/vector objects and invoke the assemble operation on the object. The assemble operation is designed to give the library a chance to set up internal data structures and repartition the matrix elements, etc. in a way that will optimize subsequent calculations. Inclusion of this operation follows the syntax of most solver libraries when they construct a matrix or vector. This module also includes some basic matrix and vector operations such as matrix-vector multiply and norms.

In addition to basic matrix operations, the math module contains linear and non-linear solvers and preconditioners. The math module provides a simple interface on top of the PETSc libraries that will allow users access to this functionality without having to be familiar with the libraries themselves. This should make it possible to construct solver routines that are comparable in complexity to Matlab scripts. The use of a wrapper instead of having users directly access the libraries will also make it simpler to switch the underlying library in an application. All that will be required will be for developers to link to an implementation of the math module interface that is built on a different library. There will not be a need to rewrite any application code. This has the advantage that if a different library is used for the math module in one application, it instantly becomes available for other applications.

The functionality in the math component is distributed between for classes, **Matrix**, **Vector**, **LinearSolver** and **NonlinearSolver**. Each of the classes is in the **gridpack::math** namespace and is described below. Like the **BaseNetwork** class, there are a lot of functions in **Matrix** and **Vector** that do not need to be used by users. Most of the functions related to matrix/vector instantiation and creation are actually located inside the mapper classes described below.

Matrices

The **Matrix** class is designed to create distributed matrices. It supports two types of matrix, **Dense** and **Sparse**. In most cases users will want to use the sparse matrix but some applications require dense matrices. The matrix constructor is

```
Matrix(const parallel::Communicator &dist,  
       const int &local_rows,  
       const int &cols,  
       const StorageType &storage_type=Sparse)
```

The communicator object **dist** specifies the set of processors that the matrix is defined on, the **local_rows** parameter corresponds to the number of rows contributed to the matrix by the processor, the **cols** parameter indicates what the second dimension of the matrix is and the **storage_type** parameter determines whether the matrix is sparse or dense. If the total dimension of the matrix is $M \times N$, then the sum of the **local_rows** parameters over all processors must equal M and the **cols** parameter is equal to N . The matrix destructor is

```
~Matrix()
```

Once a matrix has been created some inquiry functions can be used to probe the matrix size and distribution. The following functions return information about the matrix.

```
int rows() const
```



```
int localRows() const
```

```
void localRowRange(int &lo, int &hi) const
```

```
int cols()
```

The function **rows** will return the total number of rows in the matrix, **localRows** returns the number of rows associated with the calling processor, **localRowRange** returns the **lo** and **hi** index of the rows associated with the calling processor and **cols** returns the number of columns in the matrix. Note that matrices are partitioned into row blocks on each processor.

Addition functions can be used to add matrix elements to the matrix, either one at a time or in blocks. The following two calls can be used to reset existing elements or insert new ones.

```
void setElement(const int &i, const int &j,  
               const ComplexType &x)
```

```
void setElements(const int &n, const int *i, const int *j,  
                const ComplexType *x)
```

The first function will set the matrix element at the index location **(i, j)** to the value **x**. If the matrix element already exists, this function overwrites the value, if the element is not already part of the matrix, it gets added with the value **x**. Note that both **i** and **j** are zero-based indices. For the current PETSc based implementation of the math module, it is not required that the index **i** lie between the values of **lo** and **hi** obtained with **localRowRange** function, but for performance reasons it is desirable. Other implementations may require that **i** lie in this range. The second function can be used to add a collection of elements all at once. The variable **n** is the number elements to be added, the arrays **i** and **j** contain the row and column indices of the matrix elements and the array **x** contains their values. Again, it is preferable that all values in **i** lie within the range **[lo,hi]**.

Two functions that are similar to the set element functions above are the functions

```
void addElement(const int &i, const int &j,  
               const ComplexType &x)
```

```
void addElements(const int &n, const int *i, const int *j,  
                const ComplexType *x)
```

These differ from the set element functions only in that instead of overwriting the new values into the matrix, these functions will add the new values to whatever is already there. If no value is present in the matrix at that location the function inserts it.

In addition to setting or adding new elements, it is possible to retrieve matrix values using the functions

```
void getElement(const int &i, const int &j,  
                ComplexType &x) const
```

```
void getElements(const int &n, const int *i, const int *j,  
                 ComplexType *x) const
```

These functions can only access elements that are local to the processor. This means that the index **i** must lie in the range [**lo**,**hi**] returned by the function **localRowRange**.

Finally, before a matrix can be used in computations, it must be assembled and internal data structures must be set up. This can be accomplished by calling the function

```
void ready()
```

After this function has been invoked, the matrix is read for use and can be used in computations. In general, the procedure for building a matrix is 1) create the matrix object 2) determine local parameters such as **lo** and **hi** 3) set or add matrix elements and 4) assemble matrix using the **ready** function. Note that users can often avoid most of these operations by building matrices and vectors using the mapper functionality described below.

Some additional functions have been included in the matrix class that can be useful for creating matrices or writing out their values (e.g. for debugging purposes). It is often useful to create a copy of a matrix. This can be done using the clone method

```
Matrix* clone() const
```

The new matrix is an exact replica of the matrix that invokes this function.

Two functions that can be used to write the contents of a matrix, either to standard output or to a file are

```
void print (const char *filename=NULL) const  
void save(const char *filename) const
```

The first function will write the contents of the matrix to standard output if no filename is specified, otherwise it writes to the specified file, the second function will write a file in MatLAB format. These functions can be used for debugging or to create matrices that can be fed into other programs.

Once a matrix has been created, a variety of methods can be applied to it. Most of these are applied after the ready call has been made by the matrix, but some operations can be used to actually build a matrix. These functions are listed below.

void equate(const Matrix &A)

This function sets the calling matrix equal to matrix **A**.

void scale(const ComplexType &x)

Multiply all matrix elements by the value **x**.

void multiplyDiagonal(const Vector &x)

Multiply all elements on the diagonal of the calling matrix by the vector **x**. The **Vector** class is described below.

void addDiagonal(const Vector &x)

Add elements of **x** to the diagonal elements of the calling matrix.

void add(const Matrix &A)

Add the matrix **A** to the calling matrix. The two matrices must have the same number of rows and columns, but otherwise there are no restrictions on the data layout or the number and location of the non-zero entries.

void identity()

Create an identity matrix. This function assumes that the calling matrix has been created but no matrix elements have been assigned to it.

void zero()

Set all non-zero entries to zero.

void conjugate(void)

Set all entries to their complex conjugate value.

The following functions create a new matrix.

Matrix *multiply(const Matrix &A, const Matrix& B)

Multiply matrix **A** times matrix **B** to create a new matrix.

Vector *multiply(const Matrix &A, const Vector &x)

Multiply matrix **A** times vector **X** to get a new vector.

Matrix *transpose(const Matrix &A)

Take the transpose of matrix **A**.

Vectors

The vector class operates in much the same way as the matrix class. The vector functions will only be briefly reviewed here. The vector constructor is

Vector(const parallel::Communicator& comm, const int& local_length)

The parameter `local_length` is the number of contiguous elements in the vector that are held on the calling processor. The functions

int size(void) const
int localSize(void) const
void localIndexRange(int &lo, int &hi) const

can be used to get the global size of the vector or the size of the vector segment held locally on the calling processor. The **localIndexRange** function can be used to find the indices of the vector elements that are held locally.

Vector elements can be set and accessed using the functions

void setElement(const int &i, const ComplexType &x)
void setElementRange(const IdxType& lo, const int &hi, ComplexType *x)
void setElements(const int &n, const int *i, const ComplexType *x)
void addElement(const int &i, const ComplexType &x)
void addElements(const int& n, const int *i, const ComplexType *x)
void getElement(const int& i, ComplexType& x) const
void getElements(const int& n, const int *i, ComplexType *x) const
void getElementRange(const int& lo, const int& hi,
ComplexType *x) const
void ready(void)

These functions all operate in a similar way to the corresponding matrix operations. The **setElementRange** function, etc. are similar to the **setElements** function except that instead of specifying individual element indices in a separate vector, the low and high indices of the segment to which the values are assigned is specified (this assumes that the values in the array **x** represent a contiguous segment of the vector). The utility functions

Vector *clone(void) const
void print(const char* filename = NULL) const

```
void save(const char *filename) const
```

also have similar behaviors to their matrix counterparts.

Additional operations that can be performed on the entire vector include

```
void zero(void)
void equate(const Vector &x)
void fill(const ComplexType& v)
ComplexType norm1(void) const
ComplexType norm2(void) const
ComplexType normInfinity(void) const
void scale(const ComplexType& x)
void add(const ComplexType& x)
void add(const Vector& x, const ComplexType& scale = 1.0)
void elementMultiply(const Vector& x)
void elementDivide(const Vector& x)
```

The **zero** function sets all vector elements to zero, the **equate** function copies all values of the vector **x** to the corresponding elements of the calling vector, **fill** sets all elements to the value **v**, **norm1** returns the L_1 norm of the vector, **norm2** returns the L_2 norm and **normInfinity** returns the L_∞ norm. The **scale** function can be used to multiply all vector elements by the value **x**, the first **add** function can be used to add the constant **x** to all vector elements and the second **add** function can be used to add the vector **x** to the calling vector after first multiplying it by the value **scale**. The final two functions multiply or divide each element of the calling vector by the value in the vector **x**.

The following methods modify the values of the vector elements using some function of the element value.

```
void abs(void)
void real(void)
void imaginary(void)
void conjugate(void)
void exp(void)
void reciprocal(void)
```

The function **abs** replaces each element with its complex norm (absolute value), **real** and **imaginary** replace the elements with their real or imaginary values, **conjugate** replaces the vector elements with their conjugate values, **exp** replaces each vector element with the exponential of its original value and **reciprocal** replaces each element by its reciprocal.

Linear Solvers

The math module also contains solvers. The **LinearSolver** class contains a constructor

```
LinearSolver(const Matrix &A)
```

that creates an instance of the solver. The matrix **A** defines the set of linear equations $\mathbf{Ax}=\mathbf{b}$ that must be solved. The properties of the solver can be modified by calling the function

```
void configure(utility::Configuration::Cursor *props)
```

The **Configuration** module is described in more detail below. This function can be used to pass information from the input file to the solver to alter its properties.

Finally, the solver can be used to solve the set of linear equations by calling the method

```
void solve(const Vector &b, Vector &x) const
```

This function returns the solution **x** based on the right hand side vector **b**.

Non-linear Solvers

The math module also supports non-linear solvers for systems of the type $\mathbf{A}(\mathbf{x}) \cdot \mathbf{x} = \mathbf{b}(\mathbf{x})$ but the interface for these is more complicated than for the linear solvers. In order for the non-linear solver to work, two functions must be defined by the user. The first evaluates the Jacobian of the system for a given trial state of the system and the second computes the right hand side vector for a given trial state **x** of the system. The two functions are of type **JacobianBuilder** and **FunctionBuilder**. The **JacobianBuilder** function is a function with arguments

```
(const math::Vector &vec, math::Matrix &jacobian)
```

and **FunctionBuilder** is a function with arguments

```
(const math::Vector &xCurrent, math::Vector &newRHS)
```

These functions need to be added to the system somewhere. They can then be assigned to objects of type **JacobianBuilder** and **FunctionBuilder** and passed to the constructor of the non-linear solver. There are a number of ways to do this. In the following discussion, we will adopt the method used in the non-linear solver version of the power flow code that is distributed with GridPACK™.

The first step is to define a struct that can be used to implement the functions needed by the non-linear solver

```
struct SolverHelper : private utility::Uncopyable  
{  
    //Constructor
```

```

SolverHelper(// Arguments to initialize helper //)
{
    // Initialize non-linear calculation
}

:
boost::shared_ptr<math::Matrix> matrix; // Jacobian matrix
boost::shared_ptr<math::Vector> X; // Current state
:
void operator() (const math::Vector &xCurrent, math::vector &newRHS)
{
    // Evaluate RHS vector from current state xCurrent
}
void operator() (const math::Vector &xCurrent,
                  math::Matrix &Jacobian)
{
    // Evaluate Jacobian from current state xCurrent
}
}

```

Some additional data and utility functions might also be included in this struct to help with the remainder of the calculation, but the important ones for this discussion are the overloaded **operator ()** functions. In the application code, this helper struct can be initialized and used to create two functions of type **JacobianBuilder** and **FunctionBuilder** using the syntax

```

SolverHelper helper(//Arguments to initialize helper //);
math::JacobianBuilder jbuild = boost::ref(helper);
math::FunctionBuilder fbuild = boost::ref(helper);

```

At this point **jbuild** and **fbuild** are pointing to the overloaded functions in **helper** that have the appropriate arguments for a function of type **JacobianBuilder** and type **FunctionBuilder**. The **boost::ref** command provides a reference to the appropriate function in **helper** instead of making a copy, this preserves any state that might be present in **helper** between invocations of the functions **jbuild** and **fbuild** by the solver.

For the power flow application using a non-linear solver, the creation of the solver is a two-step process. First a pointer to a non-linear solver interface is created and then a particular solver instance is assigned to this interface. The power application can point to a hand-coded Newton-Raphson solver or a wrapper to the PETSc library of solvers. The code for this is the following

```

boost::scoped_ptr<math::NonlinearSolverInterface> solver;

```

```

if (useNewton) {
    math::NewtonRaphsonSolver *tmpsolver =
        new math::NewtonRaphsonSolver(*(helper.matrix), jbuild, fbuild);
    solver.reset(tmpsolver);
} else {
    solver.reset(new math::NonlinearSolver(*(helper.matrix), jbuild,
fbuild));
}

```

If you are only interested in using the **NonlinearSolver**, then it is possible to dispense with the **NonlinearSolverInterface** and just use the **NonlinearSolver** directly. The remaining call to invoke the solver is just

```

solver->solver(*helper.X);

```

Additional calls are likely to be added these to allow user-specified parameters from the input deck to be sent to the solver. In the case of the **NonlinearSolver**, these can be used to specify which PETSc solver should be used.

Network Components

Network component is a generic term for objects associated with buses and branches. These objects determine the behavior of the system and the type of analyses being done. Branch components can represent transmission lines and transformers while bus components could model loads, generators, or something else. Both kinds of components could represent measurements (e.g. for a state estimation analysis).

Network components cover a fairly broad range of behaviors and there is little that can be said about them outside the context of a specific problem. Each component inherits from a matrix-vector interface, which enables the framework to generate matrices and vectors from the network in a relatively straightforward way. In addition, buses inherit from a base bus interface and branches inherit from a base branch interface. The relationship between these interfaces is shown in Figure 4.




Figure 4. Schematic diagram showing the interface hierarchy for network components.

These base interfaces provide mechanisms for accessing the neighbors of a bus or branch and allow developers to specify what data is transferred in ghost exchanges. They do not define any physical properties of the bus or branch, it is up to application developers to do this.

Of these interfaces, the matrix-vector interface is the most important. It answers the question what block of data is contributed by a bus or network and what the dimensions of the block are. For example, if constructing the Y-matrix for a power flow problem using a real-valued formulation, the grid components on buses contribute a 2×2 block to the diagonal of the matrix. Similarly, the grid components on branches contribute a 2×2 block to the off-diagonal elements. (Note that if the Y-matrix is expressed as a complex matrix, then the blocks are of size 1×1 .) The location of these blocks in the matrix is determined by the location of the corresponding buses and branches in the network, but the indexing calculations required to determine this location can be made completely transparent to the user via the mapper module.

Because the matrix-vector interface focuses on small blocks, it is relatively easy for power grid engineers to write the corresponding methods. The full matrices and vectors can then be generated from the network using simple calls to the mapper interface (see the discussion below on the mapper module). All of the base network component classes reside in the **gridpack::component** namespace.

The **MatVecInterface** is probably the most important of the network component base classes, it is also the most difficult to understand. Its primary function is to enable developers to

build the matrices and vectors used in the solution algorithms from the network. It eliminates a large amount of tedious and error-prone index calculations that would otherwise need to be performed in order to determine where in a matrix a particular data element should be placed. The **MatVecInterface** includes basic constructors and destructors. The first set of non-trivial operations are generally implemented on buses and set the values of diagonal blocks in the matrix. Additional functions are usually implemented on branches and set values for off-diagonal elements. Vectors can be created by calling functions from buses. These functions are described in detail below.

The two functions that are used to create diagonal matrix blocks are

```
virtual bool matrixDiagSize(int *isize, int *jsize) const
```

```
virtual bool matrixDiagValues(ComplexType *values)
```

Both functions are virtual functions and are expected to be overwritten by application-specific bus and branch classes. The default behavior is to return 0 for **isize** and **jsize** and to return false for both functions. This means that these functions will not build a matrix unless overwritten by the application. Not all functions need to be overwritten by a given bus or branch class. Generally, only a subset of functions may be needed by an application.

The **matrixDiagSize** function returns the size of the matrix block that is contributed by the bus to a matrix. If a single complex number is contributed by the bus, the **matrixDiagSize** function returns 1 for both **isize** and **jsize**. If a real-valued formulation is being used so that the single returned value is expressed as a 2×2 block then both **isize** and **jsize** are set to 2. The return value is true if the bus contributes to the matrix, otherwise it is false. Returning false can occur, for example, if the bus is the reference bus in power flow calculation. For a more complicated calculation, such as a dynamic simulation with multiple generators on some buses, the size of the matrix blocks can differ from bus to bus. Note that the values returned by **matrixDiagSize** refer only to the particular bus that is invoking the function. It does not say anything about other buses in the system.

The **matrixDiagValues** function returns the actual values for the matrix block associated with the bus for which the function is invoked. The values are returned as a linear array with values returned in row-major order. For a 2×2 block, this means the first value is at the (0,0) position, the second value is at the (1,0) position, the third value is at the (0,1) position and the fourth value is at the (1,1) position. This function also returns true if the bus contributes to the matrix and false otherwise. This may seem redundant, since the **matrixDiagSize** function has already returned this information but it turns out there are certain applications where it is desirable for the **matrixDiagSize** function to return true and the **matrixDiagValues** function to return false. The buffer **values** is supplied by the calling program and is expected

to be big enough, based on the dimensions returned by the **matrixDiagSize** function, to contain all returned values.

The functions that are used to return values for off-diagonal matrix elements are listed below. These are usually only implemented for branches.

```
virtual bool matrixForwardSize(int *isize, int *jsize) const
```

```
virtual bool matrixForwardValues(ComplexType *values)
```

```
virtual bool matrixReverseSize(int *isize, int *jsize) const
```

```
virtual bool matrixReverseValues(ComplexType *values)
```

These functions work in a similar way to the functions for creating blocks along the diagonal, except that they split off-diagonal matrix calculations into forward elements and reverse elements. The initial approximate location of an off-diagonal matrix element in a matrix is based in some internal indices assigned to the buses at either end of the branch. Suppose that these indices are **i**, corresponding to the “from” bus and **j**, corresponding to the “to” bus. The “forward” functions assume that the request is for the **i****j** element while the “reverse” functions assume that the request is for the **j****i** element. Another way of looking at this is the following: as discussed below, branches contain pointers to two buses. The first is the “from” bus and the second is the “to” bus. The forward functions assume that the “from” bus corresponds to the first index of the element, the reverse functions assume that the “from” bus corresponds to the second index of the element. Note that if a bus does not contribute to a matrix, then the branches that are connected to the bus should also not contribute to the matrix.

The final set of functions in the **MatVecInterface** that are of interest to application developers are designed to set up vectors. These are usually implemented only for buses. These functions are analogous to the functions for creating matrix elements

```
virtual bool vectorSize(int *isize) const
```

```
virtual bool vectorValues(ComplexType *values)
```

The **vectorSize** function returns the number of elements contributed to the vector by a bus and the **vectorValues** returns the corresponding values. The **vectorValues** function expects the buffer values to be allocated by the calling program. In addition to functions that can be used to specify a vector, there is an additional function that can be used to push values from a vector back onto a bus. This function is

```
virtual void setValues(ComplexType *values)
```

The buffer contains values from the vector corresponding to internal variables in the bus and this function can be used to set the bus variables. The **setValues** function could be used to assign bus variables so that they can be used to recalculate matrices and vectors for an iterative loop in a non-linear solver or so that the results of a calculation can be exported to an output file.

The **BaseComponent** class contains additional functions that contribute to the base properties of a bus or branch. Again, most of the functions in this class are virtual and are expected to be overwritten by actual implementations. However, not all of them need to be overwritten by a particular bus or branch class. Many of these functions are used in conjunction with the **BaseFactory** class, which defines methods that run over all buses and branches in the network and invokes the functions defined below.

The **load** function

virtual void load(const boost::shared_ptr<DataCollection> &data)

is used to instantiate components based on data that is located in the network configuration file that is used to create the network. It is used in conjunction with the **DataCollection** object, which is described in more detail below. Networks are generally created by first instantiating a network parser and then using this to read in an external network file and create the network topology. The next step is to invoke the partition function on the network to get all network elements properly distributed between processors. At this point, the network, including ghost buses and branches, is complete and each bus and branch has a **DataCollection** object containing all the data in the network configuration file that pertains to that particular bus or branch. The data in the **DataCollection** object is stored as simple key-value pairs. This data is then used to initialize the corresponding bus or branch by invoking the load function on all buses and branches in the system. The bus and branch classes must implement the **load** function to extract the correct parameters from the **DataCollection** object and use them to assign internal bus and branch parameters.

Only one type of bus and one type of branch is associated with each network but many different types of equations can be generated by the network. To allow developers to embed many different behaviors into a single network and to control at what points in the simulation those behaviors can be manifested, the concept of modes is used. The function

virtual void setMode(int mode)

can be used to set an internal variable in the component that tells it how to behave. The variable “**mode**” usually corresponds to an enumerated constant that is part of the application definition. For example, in a power flow calculation it might be necessary to calculate both the Y-matrix and the equations for the power flow solution containing the Jacobian matrix and the right-hand

side vector. To control which matrix gets created, two modes are defined: “YBus” and “Jacobian”. Inside the matrix functions in the **MatVecInterface**, there is a condition

```
if (p_mode == YBus) {  
    // Return values for Y-matrix calculation  
} else if (p_mode == Jacobian) {  
    // Return values for power flow calculation  
}
```

The variable “**p_mode**” is an internal variable in the bus or branch that is set using the **setMode** function.

The function

```
virtual bool serialWrite(char *string, const int bufsize,  
                        const char *signal = NULL)
```

is used in the serial IO modules described below to write out properties of buses or branches to standard output. The character buffer “**string**” contains a formatted line of text representing the properties of the bus or branch that is written to standard output, the variable “**bufsize**” gives the number of characters that “**string**” can hold, and the variable “**signal**” can be used to control what data is written out and the return value is true if the bus or branch is writing out data and false otherwise. For example, if the application is writing out the properties of all buses with generators, then the signal “**generator**” might be passed to this subroutine. If a bus has generators, then a string is copied into the buffer “**string**” and the function returns true, otherwise it returns false. The buffer “**string**” is allocated by the calling program. The variable “**bufsize**” is provided so that the bus or branch can determine it is overwriting the buffer. Returning to the generator example, if this call returns a separate line for each generator, then it is possible that a bus with too many generators might exceed the buffer size. This could be detected by the implementation if the buffer size is known.

The **BaseComponent** class also contains two functions that must be implemented if buses and/or branches need to exchange data with other processors. Data that must be exchanged needs to be placed in buffers that have been allocated by the network. The bus and branch objects specify how large the buffers need to be by implementing the function

```
virtual int getXCBufSize()
```

This function must return the same value for all buses and all branches in the same bus or branch classes. Buses can return a different value than branches. For example, a power flow calculation, it is necessary that ghost buses get new values of the phase angle and voltage magnitude increments. These are both real numbers so the **getXCBusSize** routine needs to return the

value `2*sizeof(double)`. Note that all buses must return this value even if the bus is a reference bus and does not participate in the calculation.

This function is queried by the network and used to allocate a buffer of the appropriate size. The network then informs the bus and branch objects where the location of the buffer is by invoking the function

virtual void setXCBuf(void *buf)

The bus or branch can use this function to set internal pointers to this buffer that can be used to assign values to the buffer (which is done before a ghost exchange) or to collect values from the buffer (which is done after a ghost exchange). Continuing with the powerflow example, the bus implementation of the `setXCBuf` function would look like

```
setXCBuf(void *buf)
{
    p_Ang_ptr = (double*)buf;
    p_Mag_ptr = p_Ang_ptr;
}
```

The pointers `p_Ang_ptr` and `p_Mag_ptr` of type `double` are internal variables of the bus implementation and can be used elsewhere in the bus whenever the voltage angle and voltage magnitude variables are needed. After a network update operation, ghost buses will contain values for these variables that were calculated on the home processor that owns the corresponding bus.

The `BaseBusComponent` and `BaseBranchComponent` classes contain a few additional functions that are specific to whether or not a component is a bus or a branch. The `BaseBusComponent` class contains functions that can be used to identify attached buses or branches, determine if the bus is a reference bus, and recover the original indices of the bus. Other functions are included in the `BaseBusClass` but these are not usually required by application developers.

To get a list of pointers to all branches connected to a bus, the function

**void getNeighborBranches(
 std::vector<boost::shared_ptr<BaseComponent> > &nghbrs) const**

can be called. This provides a list of pointers to all branches that have the calling bus as one of its endpoints. This function can be used inside a bus method to loop over attached branches, which is a common motif in matrix calculations. For example, to evaluate the contribution to a diagonal element of the Y-matrix coming from transmission lines, it is necessary to perform the sum

$$Y_{ii} = -\sum_{j \neq i} Y_{ij}$$

where the Y_{ij} are the contribution due to transmission lines from the branch connecting i and j. The code inside a bus component that evaluates this sum can be written as

```
std::vector<boost::shared_ptr<BaseComponent> > branches;
getNeighborBranches(branches);
ComplexType y_diag(0.0,0.0);
for (int i=0; i<branches.size(); i++) {
    YBranch *branch = dynamic_cast<YBranch*>(branches[i].get());
    Y_diag += branch->getYContribution();
}
```

The function **getYContribution** evaluates the quantity Y_{ij} using parameters that are local to the branch. The return value is then accumulated into the bus variable **y_diag**, which is eventually returned through the **matrixDiagValues** function. The **dynamic_cast** is necessary to convert the pointer from a **BaseComponent** object to the application class **YBranch**. The **BaseComponent** class has no knowledge of the **getYContribution** function, this is only implemented in **YBranch**.

A function that is similar to **getNeighborBranches** is

```
void getNeighborBuses(
    std::vector<boost::shared_ptr<BaseComponent> > &nghbrs) const
```

which can be used to get a list of the buses that are connected to the calling bus via a single branch.

Many power grid problems require the specification of a special bus as a reference bus. This designation can be handled by the two functions

```
void setReferenceBus(bool status)
```

```
bool getReferenceBus() const
```

The first function can be used (if called with the argument true) to designate a bus as the reference bus and the second function can be called to inquire whether a bus is the reference bus.

Finally, it is often useful for exporting results if the original index of the bus is available. This can be recovered using the function

```
int getOriginalIndex() const
```

This function only works correctly after a call to the base factory method **setComponents**. This function is described below. Other functions in the **BaseBusComponent** class are needed within the framework but are not usually required by application developers.

The **BaseBranchComponent** class is similar to the **BaseBusComponent** class and provides basic information about branches and the buses at either end of the branch. To retrieve pointers to the buses at the ends of the branch, the following two functions are available

```
boost::shared_ptr<BaseComponent> getBus1() const
```

```
boost::shared_ptr<BaseComponent> getBus2() const
```

The **getBus1** function returns a pointer to the “from” bus, the **getBus2** function returns a pointer to the “to” bus.

Two other functions in the **BaseBranchComponent** class that are useful for writing output are

```
int getBus1OriginalIndex() const
```

```
int getBus2OriginalIndex() const
```

These functions get the original index of “from” and “to” buses. Unlike buses, the branches are not characterized by a single index. Similar to the **getOriginalIndex** function for the **BaseBusComponent** class, these functions will not work correctly until the **setComponents** method has been called in the base factory class.

Finally, a separate network component class that is associated with all buses and branches (including ghost buses and branches) is the **DataCollection** class. This class is a simple container that can be used to store key-value pairs. It also resides in the **gridpack::component** namespace. When the network is created using a standard parser to read a network configuration file (see more on parsers below), each bus and branch in the network, including the ghosts, has an associated **DataCollection** object that contains all parameters from the configuration file that are associated with that particular bus or branch. These can be retrieved from the **DataCollection** object using some simple accessors. Data can be stored in two ways inside the **DataCollection** object. The first method assumes that there is only a single instance of the key-value pair, the second assumes there are multiple instances. This second case can occur, for example, if there are multiple generators on a bus. Generators are characterized by a collection of parameters and each generator has its own set of parameters. The generator parameters can be indexed so that they can be matched with a specific generator.

Assuming that a parameter only appears once in the data collection, the contents of a **DataCollection** object can be accessed using the functions

```
bool getValue(const char *name, int *value)
bool getValue(const char *name, long *value)
bool getValue(const char *name, bool *value)
bool getValue(const char *name, std::string *value)
bool getValue(const char *name, float *value)
bool getValue(const char *name, double value)
bool getValue(const char *name, ComplexType *value)
```

These functions return true if a variable of the correct type is stored in the **DataCollection** object with the key “**name**”, otherwise it returns false.

If the variable is stored multiple times in the **DataCollection**, then it can be accessed with the functions

```
bool getValue(const char *name, int *value, const int idx)
bool getValue(const char *name, long *value, const int idx)
bool getValue(const char *name, bool *value, const int idx)
bool getValue(const char *name, std::string *value, const int idx)
bool getValue(const char *name, float *value, const int idx)
bool getValue(const char *name, double value, const int idx)
bool getValue(const char *name, ComplexType *value, const int idx)
```

where **idx** is an index that identifies a particular instance of the key. These functions are used primarily to implement the network component **load** method, described above for the **BaseComponent** class.

Factories

The network component factory is an application-dependent piece of software that is designed to manage interactions between the network and the network component objects. Most operations in the factory run over all buses and all branches and invoke some operation on each bus and each branch. An example is the “**load**” operation. After the network is read in from an external file, it consists of a topology and a set of simple data collection objects containing key-value pairs associated with each bus and branch. The **load** operation then runs over all buses and branches and instantiates the appropriate objects by invoking a local **load** method in each branch and bus object that takes the values from the data collection object and uses it to instantiate the bus or branch. The application network factory is derived from a base network factory class that contains some additional routines that set up indices, assign neighbors (which initially are only known by the network) to individual buses and branches and assign buffers.

The network component factory may also execute other routines that contribute to setting up the network and creating a well-defined state.

Factories can be derived from the **BaseFactory** class, which is a templated class that is based on the network type. It resides in the **gridpack::factory** namespace. The constructor for a **BaseFactory** object has the form

```
BaseFactory<MyNetwork>(boost::shared_ptr<MyNetwork> network)
```

The **BaseFactory** class supplies some basic functions that can be used to help instantiate the components in a network. Others can be added for particular applications by subclassing the **BaseFactory** class. The two most important functions in the **BaseFactory** class are

```
virtual void setComponents()
```

```
virtual void setExchange()
```

The **setComponents** method pushes topology information available from the network into the individual buses and branches using methods in the base component classes. This operation ensures that operations such as **getNeighborBuses**, etc. work correctly. The topology information is originally only available in the network and it requires additional operations to make it available to individual buses and branches. However, having this information imbedded in the buses and branches themselves can simplify application programming substantially.

The **setExchange** function allocates buffers and sets up pointers in the components so that exchange of data between buses and branches can occur and ghost buses and branches can receive updated values of the exchanged parameters. This functions loops over the **getXCBufferSize** and **setXCBuf** commands defined in the network component classes and guarantees that buffers are properly allocated and exposed to the network components.

Two other functions are defined in the **BaseFactory** class as convenience functions. The first is

```
virtual void load()
```

This function loops over all buses and branches and invokes the individual bus and branch **load** methods. This moves information from the **DataCollection** objects that are instantiated when the network is created from a network configuration file to the bus and branch objects themselves. The second convenience function is

```
virtual void setMode(int mode)
```

This function loops over all buses and branches in the network and invokes each bus and branch **setMode** method. It can be used to set the behavior of the entire network in single function call.

Another utility function in the **BaseFactory** class that is useful in some contexts is

bool checkTrue(bool flag)


This function returns true if the variable flag is true on all processors, otherwise it returns false.

Mapper Module


The mapper is a generic capability that can be used to generate a matrix or vector from the network components. This is done by running over all the network components and invoking methods in the matrix-vector interface. The mapper is basically a transformation that converts a set of network components into a matrix or vector based on the behavior of their matrix-vector interfaces. It has no explicit dependencies on either the network components or the type of analyses being performed so this capability is applicable across a wide range of problems. At present there are two types of mapper, the standard mapper described here that is implemented on top of the **MatVecInterface** and a more generalized mapper that utilizes the **GenMatVecInterface**. The generalized mapper and its corresponding interface are described in a later section below. The mapper discussed in this section is used for problems where both dependent and independent variables are associated with buses, which is the case for problems such as power flow calculations and dynamic simulation. Other problems, such as state estimation, have variables associated with both buses and branches and require the more general interface.

The basic matrix-vector interface contains functions that provide two pieces of information about each network component. The first is the size of the matrix block that is contributed by the component and the second is the values in that block. Using this information, the mapper can figure out what the dimensions of the matrix are and where individual elements in the matrix are located. The construction of a matrix by the mapper is illustrated in Figure 5 for a small network. Figure 5(a) shows a hypothetical network. The contributions from each network component are shown in Figure 5(b). Note that some buses and branches do not contribute to the matrix. This could occur in real systems because the transmission line corresponding to the branch has failed or because a bus represents the reference bus. In addition, it is not necessarily true that all buses and branches contribute the same size elements. The mapping of the individual contributions from the network in Figure 5(b) to initial matrix locations is shown in Figure 5(c). This is followed by elimination of gaps in the matrix in Figure 5(d).

(a)



(b)






Figure 5. A schematic diagram of the matrix map function. The bus numbers in (a) and (b) map to approximate row and column locations in (c). (a) a small network (b) matrix blocks associated with branches and buses. Not that not all blocks are the same size and not all buses and branches contribute (c) initial construction of matrix based on network indices (d) final matrix after eliminating gaps

The most complex part of generating matrices and vectors is implementing the functions in the **MatVecInterface**. Once this has been done, actually creating matrices and vectors using the mappers is quite simple. The **MatVecInterface** is associated with two mappers, one that creates matrices from buses and branches and a second that can create vectors from buses. Both mappers are templated objects based on the type of network being used and use the **gridpack::mapper** namespace.. The **FullMatrixMap** object runs over both buses and branches to set up a matrix. The constructor is

```
FullMatrixMap<MyNetwork>(boost::shared_ptr<MyNetwork> network)
```

The network is passed in to the object via the constructor. The constructor sets up a number of internal data structures based on what mode has been set in the network components. This means

that a mapper is created while the mode is set to construct the Y-matrix, then it will be necessary to instantiate a second mapper to create the Jacobian for a power flow calculation.

Once the mapper has been created, a matrix can be generated using the call

```
boost::shared_ptr<gridpack::math::Matrix> mapToMatrix()
```

This function creates a new matrix and returns a pointer to it. If a matrix already exists and it is only necessary to update the values, then the functions

```
void mapToMatrix(  
    boost::shared_ptr<gridpack::math::Matrix &matrix)
```

```
void mapToMatrix(gridpack::math::Matrix &matrix)
```

can be used. These functions use the existing matrix data structures and overwrite the values of individual elements. For these to work, it is necessary to use the same mapper that was used to create the original matrix and to have the same mode set in the network components.

The vector mapper works in an entirely analogous way to the matrix mapper. The constructor for the **BusVectorMap** class is

```
BusVectorMap<MyNetwork>(boost::shared_ptr<MyNetwork>)
```

and the function for building a new vector is

```
boost::share_ptr<gridpack::math::Vector mapToVector()
```

The functions for overwriting the values of an existing vector are

```
void mapToVector(  
    boost::shared_ptr<gridpack::math::Vector &vector)
```

```
void mapToVector(gridpack::math::Vector &vector)
```

The vector map can also be used to write values back to buses using the function

```
void mapToBus(const gridpack::math::Vector &vector)
```

This function will copy values from the vector into the bus using the **setValues** function in the **MatVecInterface**.

Parser Module

The import module is designed to read an external network file, set up the network topology and assign any parameter fields in the file to simple fields. The import module does not partition the network, it is only responsible for reading in the network and distributing the different network

elements in a way that guarantees that not too much data ends up on any one processor. The import module is also not responsible for determining if the input is compatible with the analysis being performed. This can be handled by the network factory. The import module is only responsible for determining if it can read the file.

Currently, GridPACK™ only supports one file format and there is only one parser. Files based on the PSS/E PTI version 23 format can be read in using the class **PTI23_parser**. This is another templated class that uses the network type as a template argument. **PTI23_parser** is located in the **gridpack::parser** namespace. This class has only two important functions. The first is the constructor

```
PTI23_parser<MyNetwork>(boost::shared_ptr<MyNetwork> network)
```

and the second is the function

```
void parse(const std::string &filename)
```

where filename refers to the location of the network configuration file. To use this parser, the network object with the right bus and branch classes is instantiated and then passed to the constructor of the **PTI23_parser** object. The parse method is then invoked with the location of the network configuration file passed in as an argument and the network is filled out with buses and branches. The parameters in the network configuration file are stored as key-value pairs in the **DataCollection** object associated with each bus and branch. Once the partition method has been called on the network the network is fully distributed with ghost buses and branches in place and other operations can be performed.

Another key part of the parsing capability is the **dictionary.hpp** file, which is designed to provide a common nomenclature for parameters associated with power grid components. It is also the key to extracting parameters from the **DataCollection** objects created by the parser. For example, the parameter describing the resistance of a transmission element is given the common name **BRANCH_R**. This string is defined as a macro in the dictionary.hpp file as

```
#define BRANCH_R "BRANCH_R"
```

The macro is used in all function calls that reference this variable by name. The use of a macro provides compile time error checking on the name. The goal of using the dictionary is that all parsers will eventually store the branch resistance parameter in the **DataCollection** object using this common name. Applications can then switch easily between different network configuration file formats by simply exchanging parsers, which will all store corresponding parameters using a common naming convention.

Serial IO Module

The serial IO module is designed to provide a simple mechanism for writing information from selected buses and/or branches to standard output or a file using a consistent ordering scheme. Individual buses and/or branches implement a write method that will write bus/branch information to a single string. This information usually consists of bus or branch identifiers plus some parameters that are desired in the output. The serial IO module then gathers this information, moves it to the head node, and writes it out in a consistent order. An example of this type of output is shown below.

Bus Voltages and Phase Angles

Bus Number	Phase Angle	Voltage Magnitude
1	0.000000	1.060000
2	-4.982589	1.045000
3	-12.725100	1.010000
4	-10.312901	1.017671
5	-8.773854	1.019514
6	-14.220946	1.070000
7	-13.359627	1.061520
8	-13.359627	1.090000
9	-14.938521	1.055932
10	-15.097288	1.050985
11	-14.790622	1.056907
12	-15.075585	1.055189
13	-15.156276	1.050382
14	-16.033645	1.035530

Figure 6. Example output from buses in a 14 bus problem.

Like the mapper, the serial IO classes are relatively easy to use. Most of the complexity is associated with implementing the **serialWrite** methods in the buses and branches. Data can be written out for buses and/or branches using either the **SerialBusIO** class or the **SerialBranchIO** class. These are again templated classes that take the network as an argument in the constructor. Both classes reside in the **gridpack::serial_io** namespace. The **SerialBusIO** constructor has the form

```
SerialBusIO<MyNetwork>(int max_str_len,  
    boost::shared_ptr<MyNetwork> network)
```

The variable **max_str_len** is the length, in bytes, of the maximum size string you would want to write out using this class and **network** is a pointer to the network that is used to generate output. Two additional functions can be used to actually generate output. They are

```
void header(const char *string) const
```

and


```
void write(const char *signal = NULL)
```

The **header** method is a convenience function that will only write the buffer string from the head processor (process 0) and can be used for creating the headings above an output listing. The **write** function traverses all the buses in the network and writes out the strings generated by the **serialWrite** methods in the buses. The **SerialBusIO** object is responsible for reordering these strings in a consistent manner, even if the buses are distributed over many processors. The optional variable “**signal**” is passed to the **serialWrite** methods and can be used to control what output is listed. For example, in one part of a simulation it might be desirable to list the voltage magnitude and phase angle from a powerflow calculation and in another part of the calculation to list the rotor angle for a generator. These two outputs could be distinguished from each other in the **serialWrite** function using the signal variable.

To generate the output in Figure 6, the following calls are used

```
gridpack::serial_io::SerialBusIO<MyNetwork> busIO(128,network);
busIO.header("\n  Bus Voltages and Phase Angles\n");
busIO.header(
    "\n  Bus Number      Phase Angle      Voltage Magnitude\n");
busIO.write();
```

The first call creates the **SerialIOBus** object and specifies the internal buffers size (128 bytes). This buffer must be large enough to incorporate the output from any call to one of the **serialWrite** calls in the bus components. The next two lines print out the header on top of the bus listing and the last line generates the listing itself. The **serialWrite** implementation looks like

```
bool gridpack::myapp::MyBus::serialWrite(char *string,
    const int bufsize, const char *signal)
{
    double pi = 4.0*atan(1.0);
    double angle = p_a*180.0/pi;
    sprintf(string, "      %6d      %12.6f      %12.6f\n",
        getOriginalIndex(),angle,p_v);
}
```

For this simple case, the signal is ignored as well as the variable bufsize. If more than one type of bus listing was desired, additional conditions based on the value of signal could be included.

If you wish to direct the output to a file, then calling the function

```
void open(const char *filename)
```

will direct all output from the serial IO object to the file specified in the variable filename. Similarly, calling the function

```
void close(void)
```

will close the file and all subsequent writes are directed back to standard output. The same **SerialBusIO** object can be used to write data to multiple different files, if desired. Two additional methods can be used to further control where output goes. If a file already exists, you can use the function

```
boost::shared_ptr<std::ofstream> getStream()
```

to recover a pointer to the file stream currently being used by the **SerialBusIO** object. This can then be used to redirect output from some other part of the code to the same file. The function

```
void setStream(boost::shared_ptr<std::ofstream> stream)
```

can be used to redirect the output from the **SerialIOBus** object to an already existing file. The main use of these two functions is to direct the output from both buses and branches to the same file instead of standard output.

The **SerialBranchIO** module is similar to the **SerialBusIO** module but works by creating listings for branches. The constructor is

```
SerialBranchIO<MyNetwork>(int max_str_len,  
    boost::shared_ptr<MyNetwork> network)
```

and the header and write methods are

```
void header(const char *string) const
```

```
void write(const char *signal = NULL)
```

These have exactly the same behavior as in the **SerialBusIO** class. Similarly, the methods

```
void open(const char *filename)  
void close(void)  
boost::shared_ptr<std::ofstream> getStream()  
void setStream(boost::shared_ptr<std::ofstream> stream)
```

can be used to redirect output to a file instead of standard output.

Configuration Module

The configuration module is designed to provide a central mechanism for directing information from the input file to the components making up a given application. For example, information about convergence thresholds and maximum numbers of iterations might need to be picked up by the solver module from an external configuration file. The configuration module is designed to read files using a simple XML format that supports a hierarchical input. This can be used to control which input gets directed to individual objects in the application, even if the object is a framework component and cannot be modified by the application developer.

The **Configuration** class is in the namespace **gridpack::utility**. This class does not have a public constructor. The static method **configuration()** returns a pointer to the shared instance of this class used by all modules in an application. This function is initialized once:

```
gridpack::utility::Configuration * c =  
    gridpack::utility::Configuration::configuration() ;  
c->open(input_file, MPI_COMM_WORLD);
```

The input file uses XML markup syntax. The single top-level element must be named “Configuration” but other elements may have module and application specific names. Refer elsewhere in this document for specifics. For illustration only, an example configuration file might look like:

```
<?xml version="1.0" encoding="utf-8"?>  
<Configuration>  
  <PowerFlow>  
    <networkConfiguration> IEEE118.raw </networkConfiguration>  
  </PowerFlow>  
  <DynamicSimulation>  
    <StartTime> 0.0 </StartTime>  
    <EndTime> 0.1 </EndTime>  
    <TimeStep> 0.001 </TimeStep>  
    <Faults>  
      <Fault>  
        <StartFault> 0.03 </StartFault>  
        <EndFault> 0.06 </EndFault>  
        <Branch> 3 7 </Branch>  
      </Fault>  
      <Fault>  
        <StartFault> 0.07 </StartFault>  
        <EndFault> 0.06=8 </EndFault>  
        <Branch> 4 8 </Branch>  
      </Fault>  
    </Faults>  
  </DynamicSimulation>  
</Configuration>
```

A value in this configuration file is accessed with a call to the overloaded method **get()**. The following line will return the value of the input file corresponding to the XML field “networkConfiguration”

```
std::string s =  
    c->get("Configuration.PowerFlow.networkConfiguration",  
          "IEEE.raw");
```

The first argument has type **Configuration::KeyType** which is a **typedef** of **std::string**. Values are selected by hierarchically named “keys” using “.” as a separator. The second argument is a default value that is used if field corresponding to the key can’t be found. There are overloads of **get()** for accessing C++ types: **bool**, **int**, **long**, **float**, **double**, **ComplexType** and **std::string**. For each type there are two variants. For integers these look like

```
int get(const KeyType &, int default_value) const ;  
bool get(const KeyType &, int *) const;
```

The first variant takes a key name and a default value and returns the value in the configuration file or the default value when none is specified. In the second variant, a Boolean value is returned indicating whether or not the value was in the configuration file and the second argument points to an object that is updated with the configuration value when it is present. For strings, the second argument is passed by reference.

The method **getCursor(KeyType)** returns a pointer to an internal element in the hierarchy. This “cursor” supports the same **get()** methods as above but the names are now relative to the name of the cursor. Thus we might modify the previous example to:

```
Configuration::CursorPtr p =  
    c->getCursor("Configuration.PowerFlow");  
  
std::string s = p->get("networkConfiguration",  
                      "IEEE14.raw");
```

An additional use of such cursors is to access lists of values. The method

```
typedef std::vector<CursorPtr> ChildCursors;  
  
void children(ChildCursors &);
```

can be used to get a vector of all the elements that are children in the name hierarchy of some element. These children need not have unique names as illustrated by the children of the “Faults”

element shown above. In this example, each of the children is a cursor that can be used to access “StartFault”, “EndFault”, and “Branch” parameters for each of the “Fault” blocks.

Developing Applications

The use of these modules in an application such as power flow is outlined in Figure 7. For different power grid problems, the details of the code will be different, but most of these motifs will appear at some point or other. The main differences will probably be in feedback loops as results from one part of the calculation are fed back into other parts of the calculation. For example, an iterative solver will need to update the current values of the network components, which can then be used to generate new matrices and vectors that are fed back into the next iteration of the solver. The diagram is not complete, but gives an overall view of code structure and data movement.

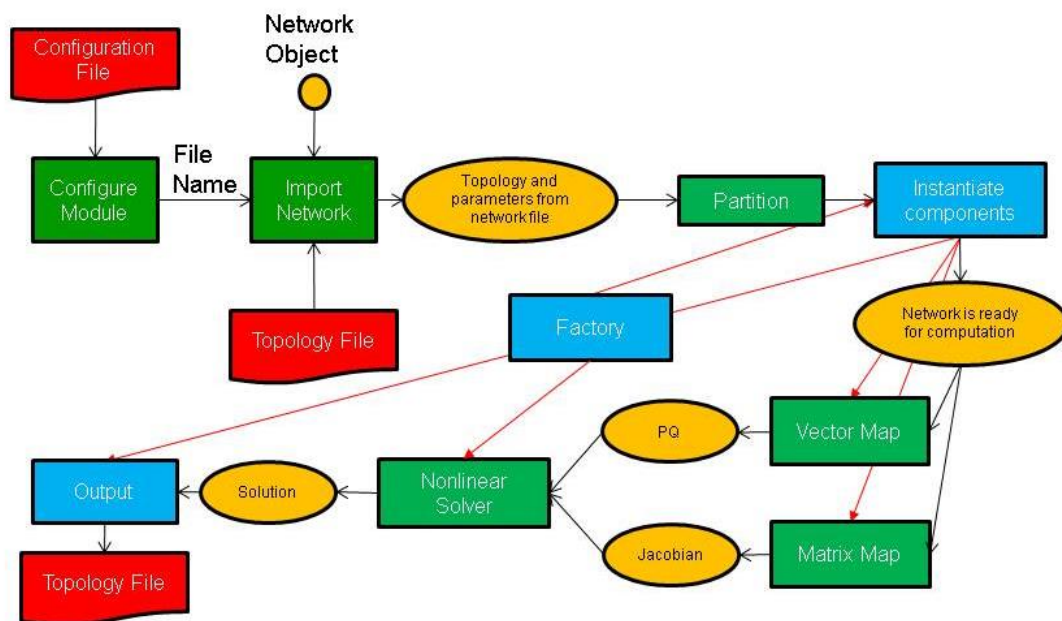


Figure 7. Schematic of program flow for a power flow simulation. The yellow ovals are distributed data objects, the green blocks are GridPACK™ framework components and the blue blocks are application specific code. External files are red.

As shown in the figure, application developers will need to focus on writing two or three sets of modules. The first is the network components. These are the descriptions of the physics and/or measurements that are associated with buses and branches in the power grid network. The network factory is a module that initializes the grid components on the network after the network is originally created by the import module. The power flow problem is simple enough that it can use a non-linear solver directly from the math module but even a straightforward solution such as this requires the developer to overwrite some functions in the factory that are used in the non-linear solver iterations.

Most of the work involved in creating a new application is centered on creating the bus and branch classes for the application. This discussion will focus on creating a code to perform power flow calculations and will describe in some detail the routines that need to be written in order to develop a working simulation. This application has been included as part of the GridPACK™ distribution and users are encouraged to look at the source code. Additional applications for dynamic simulation and contingency analysis have also been included in the distribution and users are encouraged to look at these for additional coding examples on how to use GridPACK™. The discussion below is designed to illustrate how to build an application and for brevity has left out some code compared to the working implementation. The source code contains more comment lines as well as some additional diagnostics that may not appear here. However, the overall design is the same and readers who have a good understanding of the following text should have no difficulty understanding the powerflow source code.

For the power flow calculation, the buses and branches will be represented by the classes **PFBus** and **PFBranch**. **PFBus** inherits from the **BaseBusComponent** class, so it automatically inherits the **BaseComponent** and **MatVecInterface** classes as well. The first thing that must be done in creating the **PFBus** component is to overwrite the load function in the **BaseComponent** class. The original function is just a placeholder that performs no action. The **load** function should take parameters from the **DataCollection** object associated with each bus and use them to initialize the bus component itself. For the **PFBus** component, a simplified **load** function is

```
void gridpack::powerflow::PFBus::load(
    const boost::shared_ptr<gridpack::component
        ::DataCollection> &data)
{
    data->getValue(CASE_SBASE, &p_sbase);
    data->getValue(BUS_VOLTAGE_ANG, &p_angle);
    data->getValue(BUS_VOLTAGE_MAG, &p_voltage);
    p_v = p_voltage;
    double pi = 4.0*atan(1.0); p_angle = p_angle*pi/180.0;
    p_a = p_angle;
    int itype; data->getValue(BUS_TYPE, &itype);
    if (itype == 3) {
        setReferenceBus(true);
    }
    bool lgen;
    int i, ngen, gstatus;
    double pg, qg;
    if (data->getValue(GENERATOR_NUMBER, &ngen)) {
```

```

    for (i=0; i<ngen; i++) {
        lgen = true;
        lgen = lgen && data->getValue(GENERATOR_PG, &pg,i);
        lgen = lgen && data->getValue(GENERATOR_QG, &qg,i);
        lgen = lgen && data->getValue(GENERATOR_STAT, &gstatus,i);
        if (lgen) {
            p_pg.push_back(pg);
            p_qg.push_back(qg);
            p_gstatus.push_back(gstatus);
        }
    }
}

```

This version of the **load** function has left off additional properties, such as shunts and loads and some transmission parameters, but it serves to illustrate how **load** is suppose to work. The **load** method in the base factory class will run over all buses, get the **DataCollection** object associated with that bus and then call the **PFBus::load** method for that bus using the **DataCollection** object as the argument. The parameters **p_sbase**, **p_angle**, **p_voltage** are private members of **PFBus**. The variables corresponding to the keys **CASE_SBASE**, **BUS_VOLTAGE_ANG**, **BUS_VOLTAGE_MAG** were stored in the **DataCollection** object when the network configuration file was parsed. They are retrieved from this object using the **getValue** functions and assigned to **p_sbase**, **p_angle**, **p_voltage**. Additional internal variables are also assigned in a similar manner. The value of the **BUS_TYPE** variable can be used to determine whether the bus is a reference bus. Note that the **CASE_SBASE** etc. are just preprocessor symbols that are defined in the **dictionary.hpp** file, which must be included in the file defining the **load** function.

The variables referring to generators have a different behavior than the other variables. A bus can have multiple generators and these are stored in the **DataCollection** object with an index. The number of generators is also stored in the **DataCollection** object with the key **GENERATOR_NUMBER**. First the number of generators is retrieved and then a loop is set up so that all the generator variables can be accessed. The generator parameters are stored in local private arrays. The loop shows how the return value of the **getValue** function can be used to verify that all three parameters for a generator where found. If they aren't found, then the generator is incomplete and the generator is not added to the local data. This can also be used to determine if the bus has other properties and to set internal flags and parameters accordingly. The load function for the **PFBBranch** is constructed in a similar way, except that the focus is on extracting branch related parameters from the **DataCollection** object.

Both the **PFBus** and **PFBBranch** classes contain an application-specific function called **setYBus** that is used to set up values in the Y-matrix. There is also a function in the powerflow factory class that runs over all buses and branches and calls this function. The **setYBus** function in **PFBus** is

```
void gridpack::powerflow::PFBus::setYBus(void)
{
    gridpack::ComplexType ret(0.0,0.0);
    std::vector<boost::shared_ptr<BaseComponent> > branches;
    getNeighborBranches(branches);
    int size = branches.size();
    int i;
    for (i=0; i<size; i++) {
        gridpack::powerflow::PFBBranch *branch
            = dynamic_cast<gridpack::powerflow::PFBBranch*>
              (branches[i].get());
        ret -= branch->getAdmittance();
        ret -= branch->getTransformer(this);
        ret += branch->getShunt(this);
    }
    if (p_shunt) {
        gridpack::ComplexType shunt(p_shunt_gs,p_shunt_bs);
        ret += shunt;
    }
    p_ybusr = real(ret);
    p_ybusi = imag(ret);
}
```

This function evaluates the contributions to the Y-Matrix associated with buses. The real and imaginary parts of this number are stored in the internal variables **p_ybusr** and **p_ybusi**. The subroutine first creates the local variable **ret** and then gets a list of pointers to neighboring branches from the **BaseBusComponent** function **getNeighborBranches**. The function then loops over each of the branches and casts the **BaseComponent** pointer from the list to a **PFBBranch** pointer. Note that the cast is necessary since the **getNeighborBranches** function only returns a list of **BaseComponent** object pointers. The **BaseComponent** class does not contain application-specific functions such as **getAdmittance**. The **getAdmittance**, **getTransformer** and **getShunt** methods return the contributions from simple transmission elements, transformers and shunts associated with the branch. These are accumulated into the **ret** variable. Note that some parameters, such as **p_shunt**, are set in the full **PFBus::load** method but not in the truncated version discussed above.

The reason that the **getAdmittance** variable has no argument while both **getTransformer** and **getShunt** take the pointer “**this**” as an argument is that the contribution from simple transmission elements is symmetric with respect to whether or not the calling bus is the “from” or “to” buses while the transformer and shunt contributions are not. This can be seen by examining the **getTransformer** function.

```
gridpack::ComplexType
gridpack::powerflow::PFBranch::getTransformer(
    gridpack::powerflow::PFBUS *bus)
{
    gridpack::ComplexType ret(p_resistance,p_reactance);
    if (p_xform) {
        ret = -1.0/ret;
        gridpack::ComplexType a(cos(p_phase_shift),sin(p_phase_shift));
        a = p_tap_ratio*a;
        if (bus == getBus1().get()) {
            ret = ret/(conj(a)*a);
        } else if (bus == getBus2().get()) {
            // ret is unchanged
        }
    } else {
        ret = gridpack::ComplexType(0.0,0.0);
    }
    return ret;
}
```

The variables **p_resistance**, **p_reactance**, **p_phase_shift**, and **p_tap_ratio** are all internal variables that are set based on the variables read in from using the **load** method or are set in other initialization steps. The boolean variable **p_xform** variable is set to true in the **PFBranch::load** method if transformer-related variables are detected in the **DataCollection** objects associated with the branch, otherwise it is false.

The **PFBranch** version of the **setYBus** function is

```
void gridpack::powerflow::PFBranch::setYBus(void)
{
    gridpack::ComplexType ret(p_resistance,p_reactance);
    ret = -1.0/ret;
    gridpack::ComplexType a(cos(p_phase_shift),sin(p_phase_shift));
    a = p_tap_ratio*a;
    if (p_xform) {
```

```

    p_ybusr_frwd = real(ret/conj(a));
    p_ybusi_frwd = imag(ret/conj(a));
    p_ybusr_rvrs = real(ret/a);
    p_ybusi_rvrs = imag(ret/a);
} else {
    p_ybusr_frwd = real(ret);
    p_ybusi_frwd = imag(ret);
    p_ybusr_rvrs = real(ret);
    p_ybusi_rvrs = imag(ret);
}
gridpack::powerflow::PFBUS *bus1 =
    dynamic_cast<gridpack::powerflow::PFBUS*>(getBus1().get());
gridpack::powerflow::PFBUS *bus2 =
    dynamic_cast<gridpack::powerflow::PFBUS*>(getBus2().get());
p_theta = (bus1->getPhase() - bus2->getPhase());
}

```

Note that the branch version of the **setYBus** function calculates different values for the Y-matrix contribution depending on whether the first index in the matrix element corresponds to bus 1 (the forward direction) or bus 2 (the reverse direction). These are stored in the separate variables **p_ybusr_frwd** and **p_ybusi_frwd** for the forward directions and **p_ybusr_rvrs** and **p_ybusi_rvrs** for the reverse direction. This routine also calculates the variable **p_theta** which is equal to the difference in the phase angle variable associated with the two buses at either end of the branch. This last variable provides an example of calculating a branch parameter based on the values of parameters located on the terminal.

The **setYBus** functions are used in the powerflow components to set some basic parameters. These are eventually incorporated into the Jacobian matrix and PQ vector that constitute the matrix and right hand side vector of the powerflow equations. To build the matrix, it is necessary to implement the matrix size and matrix values functions in the **MatVecInterface**. The functions for setting up the matrix are discussed in detail in the following, the vector functions are simpler but follow the same pattern. The mode used for setting up the Jacobian matrix is “**Jacobian**”. The corresponding **matrixDiagSize** routine is

```

bool gridpack::powerflow::PFBUS::matrixDiagSize(int *isize,
    int *jsize) const
{
    if (p_mode == Jacobian) {
        *isize = 2;
        *jsize = 2;
        return true;
    }
}

```

```

    } else if (p_mode == YBus) {
        *isize = 1;
        *jsize = 1;
        return true;
    }
}

```

This function handles two modes, stored in the internal variable **p_mode**. If the mode equals **Jacobian**, then the function returns a contribution to a 2×2 matrix. In the case that the mode is “**YBus**” the function would return a contribution to a 1×1 matrix. (The Jacobian is treated as a real matrix where the real and complex parts of the problem are treated as separate variables, the Y-matrix is handle as a regular complex matrix). The corresponding code for returning the diagonal values is

```

bool gridpack::powerflow::PFBUS::matrixDiagValues(ComplexType *values)
{
    if (p_mode == YBus) {
        gridpack::ComplexType ret(p_ybusr,p_ybusi);
        values[0] = ret;
        return true;
    } else if (p_mode == Jacobian) {
        if (!getReferenceBus()) {
            values[0] = -p_Qinj - p_ybusi * p_v *p_v;
            values[1] = p_Pinj - p_ybusr * p_v *p_v;
            values[2] = p_Pinj / p_v + p_ybusr * p_v;
            values[3] = p_Qinj / p_v - p_ybusi * p_v;
            if (p_isPV) {
                values[1] = 0.0;
                values[2] = 0.0;
                values[3] = 1.0;
            }
            return true;
        } else {
            values[0] = 1.0;
            values[1] = 0.0;
            values[2] = 0.0;
            values[3] = 1.0;
            return true;
        }
    }
}

```

```
}
```

If the mode is “**YBus**”, the function returns a single complex value. If the mode is “**Jacobian**”, the function checks first to see if the bus is a reference bus or not. If the bus is not a reference bus, then the function returns a 2×2 block corresponding to the contributions to the Jacobian matrix coming from a bus element. If the bus is a reference bus, the function returns a 2×2 identity matrix. This is a result of the fact that the variables associated with a reference bus are fixed. In fact, the variables contributed by the reference bus could be eliminated from the matrix entirely by returning false if the mode is “**Jacobian**” and the bus is a reference bus for both the matrix size and matrix values routines. This would also require some adjustments to the off-diagonal routines as well. There is an additional condition for the case where the bus is a “PV” bus. In this case one of the independent variables is eliminated by setting the off-diagonal elements of the block to zero and the second diagonal element equal to 1. The values are returned in row-major order, so **values[0]** corresponds to the (0,0) location in the 2×2 block, **values[1]** is the (1,0) location, **values[2]** is the [0,1] location and **values[3]** is the (1,1) location.

The **matrixForwardSize** and **matrixForwardValues** routines, as well as the corresponding Reverse routines, are implemented in the **PFBranch** class. These functions determine the off-diagonal blocks of the Jacobian and Y-matrix. The **matrixForwardSize** routine is given by

```
bool gridpack::powerflow::PFBranch::matrixForwardSize(int *isize,
    int *jsize) const
{
    if (p_mode == Jacobian) {
        gridpack::powerflow::PFBUS *bus1
            = dynamic_cast<gridpack::powerflow::PFBUS*>(getBus1().get());
        gridpack::powerflow::PFBUS *bus2
            = dynamic_cast<gridpack::powerflow::PFBUS*>(getBus2().get());
        bool ok = !bus1->getReferenceBus();
        ok = ok && !bus2->getReferenceBus();
        if (ok) {
            *isize = 2;
            *jsize = 2;
            return true;
        } else {
            return false;
        }
    }
    else if (p_mode == YBus) {
        *isize = 1;
    }
}
```

```

        *jsize = 1;
        return true;
    }
}

```

If the mode is “**YBus**”, the size function returns a 1×1 block for the off-diagonal matrix block. For the Jacobian, the function first checks to see if either end of the branch is a reference bus by evaluating the Boolean variable “ok”. If neither end is the reference bus then the function returns a 2×2 block, if one end is the reference bus then the function returns false. The false value indicates that this branch does not contribute to the matrix. For this system, the **matrixReverseSize** function is the same, but if the off-diagonal contributions were not square blocks, then the dimensions of the blocks would need to be switched.

The **matrixForwardValues** function is

```

bool gridpack::powerflow::PFBranch::matrixForwardValues(
    ComplexType *values)
{
    if (p_mode == Jacobian) {
        gridpack::powerflow::PFBus *bus1
            = dynamic_cast<gridpack::powerflow::PFBus*>(getBus1().get());
        gridpack::powerflow::PFBus *bus2
            = dynamic_cast<gridpack::powerflow::PFBus*>(getBus2().get());
        bool ok = !bus1->getReferenceBus();
        ok = ok && !bus2->getReferenceBus();
        if (ok) {
            double cs = cos(p_theta);
            double sn = sin(p_theta);
            values[0] = (p_ybusr_frwd*sn - p_ybusi_frwd*cs);
            values[1] = (p_ybusr_frwd*cs + p_ybusi_frwd*sn);
            values[2] = (p_ybusr_frwd*cs + p_ybusi_frwd*sn);
            values[3] = (p_ybusr_frwd*sn - p_ybusi_frwd*cs);
            values[0] *= ((bus1->getVoltage())*(bus2->getVoltage()));
            values[1] *= -((bus1->getVoltage())*(bus2->getVoltage()));
            values[2] *= bus1->getVoltage();
            values[3] *= bus1->getVoltage();
            bool bus1PV = bus1->isPV();
            bool bus2PV = bus2->isPV();
            if (bus1PV & bus2PV) {
                values[1] = 0.0;
                values[2] = 0.0;
            }
        }
    }
}

```

```

        values[3] = 0.0;
    } else if (bus1PV) {
        values[1] = 0.0;
        values[3] = 0.0;
    } else if (bus2PV) {
        values[2] = 0.0;
        values[3] = 0.0;
    }
    return true;
} else {
    return false;
}
} else if (p_mode == YBus) {
    values[0] = gridpack::ComplexType(p_ybusr_frwd,p_ybusi_frwd);
    return true;
}
}

```

For the “**YBus**” mode, the function simply returns the complex contribution to the Y-matrix. For the “**Jacobian**” mode, the function first determines if either end of the branch is connected to the reference bus. If it is, then function returns false and there is no contribution to the Jacobian. If neither end of the branch is the reference bus then the function evaluates the 4 elements of the 2×2 contribution to the Jacobian coming from the branch. To do this, the branch needs to get the current values of the voltages on the buses at either end. It can do this by using the **getVoltage** accessor functions that have been defined in the **PFBUS** class. Finally, if one end or the other of the branch is a PV bus, then some variables need to be eliminated from the equations. This can be done by setting some of the values in the 2×2 block equal to zero.

The **matrixReverseValues** function is similar to the **matrixForwardValues** functions with a few key differences. 1) the variables **p_ybusr_rvrs** and **p_ybusi_rvrs** are used instead of **p_ybusr_frwd** and **p_ybusi_frwd** 2) instead of using **cos(p_theta)** and **sin(p_theta)** the function uses **cos(-p_theta)** and **sin(-p_theta)** since **p_theta** is defined as difference in phase angle on bus 1 minus the difference in phase angle on bus 2 and 3) the values that are set to zero in the conditions for PV buses are transposed. The PV conditions are the same as the forward case if both bus 1 and bus 2 are PV buses, if only bus 1 is a PV bus then **values[2]** and **values[3]** are zero and if only bus2 is a PV bus then **values[1]** and **values[3]** are zero.

The functions for setting up vectors are similar to the corresponding matrix functions, although they are a bit simpler. The vector part of the **MatVecInterface** contains one function that does not have a counterpart in the set of matrix functions and that is the **setValues** function.

This function can be used to push values in a vector object back into the buses that were used to generate the vector. For the Newton-Raphson iterations used to solve the powerflow equations, it is necessary at each iteration to push the current solution back into the buses so they can be used to evaluate new Jacobian and right hand side vectors. The solution vector contains the current increments to the voltage and phase angle. These are written back to the buses using the function

```
void gridpack::powerflow::PFBus::setValues(  
    gridpack::ComplexType *values)  
{  
    p_a -= real(values[0]);  
    p_v -= real(values[1]);  
    *p_vAng_ptr = p_a;  
    *p_vMag_ptr = p_v;  
}
```

This function is paired with a mapper that is used to create a vector with the same pattern of contributions. If for example, the matrix equation $Ax = b$ is being solved, then the mapper used to create the right hand side vector b should be used to push results back onto the buses using the **mapToBus** method. The **setValues** method above takes the contributions from the solution vector and uses them to decrement the internal variables **p_a** (voltage angle) and **p_v** (voltage magnitude). The new values of **p_a** and **p_v** are then assigned to the buffers **p_vAng_ptr** and **p_vMag_ptr** so that they can be exchanged with other buses. This is discussed below.

The two routines that need to be created in the **PFBus** class to copy data to ghost buses are both simple. There is no need to create corresponding routines in the **PFBranch** class since branches do not need to exchange data. Two values need to be exchanged between buses, the current voltage angle and the current voltage magnitude. This requires a buffer that is the size of two doubles so the **getXCBufSize** function is written as

```
int gridpack::powerflow::PFBus::getXCBufSize(void)  
{  
    return 2*sizeof(double);  
}
```

The **setXCBuf** assigns the buffer created in the base factory **setExchange** function to internal variables used within the **PFBus** component. It has the form

```
void gridpack::powerflow::PFBus::setXCBuf(void *buf)  
{  
    p_vAng_ptr = static_cast<double*>(buf);  
    p_vMag_ptr = p_vAng_ptr+1;  
    *p_vAng_ptr = p_a;
```

```

    *p_vMag_ptr = p_v;
}

```

The buffer created in the **setExchange** routine is split between the two internal pointers **p_vAng_ptr** and **p_vMag_ptr**. These are then initialized to the current values of **p_a** and **p_v**. Whenever the **updateBuses** routine is called the pointers on ghost buses are refreshed with the current values of these variables from the processes that own the corresponding buses. Note that both the **getXCBufSize** and the **setXCBuf** routines are only called during the **setExchange** routine. They are not called during the actual bus updates.

One final function in the **PFBus** and **PFBranch** class that is worth taking a brief look at is the set mode function. This function is used to set the internal **p_mode** variable that is defined in both classes. The **PFMode** enumeration, which contains both the “YBus” and “Jacobian” modes, is defined within the `gridpack::powerflow` namespace. The **setMode** function for both buses and branches has the form

```

void gridpack::powerflow::PFBus::setMode(int mode)
{
    p_mode = mode;
}

```

This function is triggered on all buses and branches if the base factory **setMode** method is called.

Once the **PFBus** and **PFBranch** classes have been defined, it is possible to declare a **PFNetwork** as a **typedef**. This can be done using the line

```

typedef network::BaseNetwork<PFBus, PFBranch > PFNetwork;

```

in the header file declaring the **PFBus** and **PFBranch** classes. This type can then be used in other powerflow files that need to create objects from templated classes.

The discussion above summarizes many of the important functions in the **PFBus** and **PFBranch** classes. Additional functions are included in these classes that are not discussed here, but the basic principles involved in implementing the remaining functions have been covered.

The first part of creating a new application is writing the network component classes. The second part is to implementing the application-specific factory. For the powerflow application, this is the **PFFactory** class, which inherits from the **BaseFactory** class. Most of the important functionality in the **PFFactory** is derived from the **BaseFactory** class and is used without modification, but several application-specific functions have been added to **PFFactory** that are

used to set internal parameters in the network components. As an example, consider the **setYBus** function

```
void gridpack::powerflow::PFFactory::setYBus(void)
{
    int numBus = p_network->numBuses();
    int numBranch = p_network->numBranches();
    int i;
    for (i=0; i<numBus; i++) {
        dynamic_cast<PFBus*>(p_network->getBus(i).get())->setYBus();
    }
    for (i=0; i<numBranch; i++) {
        dynamic_cast<PFBranch*>(p_network->getBranch(i).get())->setYBus();
    }
}
```

This function loops over all buses and branches and invokes the **setYBus** method in the individual **PFBus** and **PFBranch** objects. The first two lines in the factory **setYBus** method get the total number of buses and branches on the process. A loop over all buses on the process is initiated and a pointer to the bus object is obtained via the **getBus** bus method in the **BaseNetwork** class. This pointer is returned as a **BaseComponent** object, which doesn't have a **setYBus** method so it must then be cast to a **PFBus** pointer which can then invoke **setYBus**. The same set of steps is then repeated for the branches. The factory can be used to create other methods that invoke functions on buses and/or branches. Most of these functions follow the same general form as the **setYBus** method just described.

The last part of building an application is creating the top level application driver that actually instantiates all the objects used in the application and controls the program flow. Running the code is broken up into two parts. The first is creating a main program and the second is creating the application driver. The main routine is primarily responsible for initializing the communication libraries and creating the application object, the application object then controls the application itself. The main program for the powerflow application is

```
main(int argc, char **argv)
{
    int ierr = MPI_Init(&argc, &argv);
    gridpack::math::Initialize();
    GA_Initialize();
    int stack = 200000, heap = 200000;
    MA_init(C_DBL, stack, heap);
}
```

```

gridpack::powerflow::PFApp app;
app.execute();

GA_Terminate();
gridpack::math::Finalize();
ierr = MPI_Finalize();
}

```

The first block of code in this program initializes the MPI and GA communication libraries and allocates internal memory used by GA. It also initializes the math libraries, which, in turn, calls the initialization routines of whatever library the math module is built on. The code then instantiates a powerflow application object and calls the execute method for this object. The remainder of the powerflow application is contained in the **PFApp::execute** method. Finally, when the application has finished running, the main program cleans up the communication and math libraries. The main reason for breaking the code up in this way instead of including the execute function as part of main is to force the invocation of all the destructors in the GridPACK™ objects used to implement the application. Otherwise, these destructors get called after the **GA_Terminate** and **MPI_Finalize** calls are made and the program will fail to exit cleanly.

The powerflow execute method is where the top level control of the application is embedded. The execute method starts off with the code

```

gridpack::parallel::Communicator world;
boost::shared_ptr<PFNetwork> network(new PFNetwork(world));

gridpack::utility::Configuration *config
    = gridpack::utility::Configuration::configuration();
config->open("input.xml",world);
gridpack::utility::Configuration::Cursor *cursor;
cursor = config->getCursor("Configuration.Powerflow");
std::string filename;
if (!cursor->get("networkConfiguration", &filename)) {
    printf("No network configuration specified\n");
    return;
}
gridpack::parser::PTI23_parser<PFNetwork> parser(network);
parser.parse(filename.c_str());

network->partition();

```

The first two lines create a communicator for this application and use it to instantiate a **PFFNetwork** object (note that this is really a **BaseNetwork** template that is instantiated using the **PFFBus** and **PFFBranch** classes). The network object exists but has no buses or branches associated with it. The next few lines get an instance of the configuration object and use this to open the **input.xml** file. This filename has been hardwired into this implementation but it could be passed in as a runtime argument, if desired. The code then creates a **Cursor** object and initializes this to point into the **Configuration.Powerflow** block of the **input.xml** file. The cursor can then be used to get the contents of the **networkConfiguration** block in **input.xml**, which corresponds to the name of the network configuration file containing the powergrid network. After getting the file name, the code creates a **PTI23_parser** object and passes in the current network object as an argument. When the parse method is called, the parser reads in the file specified in filename and uses that to add buses and branches to the network object. At this point, the network has all the bus and branches from the configuration file, but no ghost buses or branches exist and buses and branches are not distributed in an optimal way. Calling the partition method on the network then distributes the buses and branches and adds appropriate ghost buses and branches.

The next set of calls initialize the network components and prepare the network for computation.

```
gridpack::powerflow::PFFactory factory(network);  
factory.load();  
  
factory.setComponents();  
factory.setExchange();  
  
network->initBusUpdate();  
  
factory.setYBus();
```

The first call creates a **PFFactory** object and instantiates it with a reference to the current network. The next line calls the **BaseFactory load** method which invokes the component **load** method on all buses and branches. These use data from the **DataCollection** objects to initialize the corresponding bus and branch objects. Note that when the partition function creates the ghost bus and branch objects, it copies the associated **DataCollection** objects to these ghosts so the parameters from the network configuration file are available to instantiate all objects in the network.

The next two methods are also implemented as **BaseFactory** methods. The **setComponents** method sets up pointers in the network components that point to neighboring branches and buses (in the case of buses) and terminal buses (in the case of branches). It is also responsible for setting up internal indices that are used by the mapper functions to create

matrices and vectors. The **setExchange** method is responsible for setting up the buffers that are used to exchange data between locally owned buses and branches and their corresponding ghost images on other processors. The call to **initBusUpdate** creates internal data structures that are used to exchange bus data between processors and the final factory call to **setYBus** evaluates the Y-matrix contributions from all network components. The network is fully initialized at this point and ready for computation.

The next calls create the Y-matrix and the matrices used in the Newton-Raphson iteration loop.

```
factory.setMode(YBus);
gridpack::mapper::FullMatrixMap<PFNetwork> mMap(network);
boost::shared_ptr<gridpack::math::Matrix> Y = mMap.mapToMatrix();

factory.setSBus();
factory.setMode(RHS);
gridpack::mapper::BusVectorMap<PFNetwork> vMap(network);
boost::shared_ptr<gridpack::math::Vector> PQ = vMap.mapToVector();

factory.setMode(Jacobian);
gridpack::mapper::FullMatrixMap<PFNetwork> jMap(network);
boost::shared_ptr<gridpack::math::Matrix> J = jMap.mapToMatrix();
boost::shared_ptr<gridpack::math::Vector> X(PQ->clone());
```

The first call sets the internal `p_mode` variable in all network components to “**YBus**”. The second call constructs a **FullMatrixMap** object **mMap** and the third call uses the **mapToMatrix** method to generate a Y-matrix based on the “**YBus**” mode. The factory then calls the **setSBus** method that sets some additional network component parameters (again, by looping over all buses and invoking a **setSBus** method on each bus). The next three lines set the mode to “**RHS**” create a **BusVectorMap** object and create the right hand side vector in the powerflow equations using the **vectorToMap** method. This builds the vector based on the “**RHS**” mode. The next three lines create the Jacobian using the same pattern as for the Y-matrix. The mode gets set to “**Jacobian**”, another **FullMatrixMap** object is created and this is used to create the Jacobian using the **mapToMatrix** method. Two separate mappers are used to create the Y-matrix and the Jacobian. This is required unless there is some reason to believe that the “**YBus**” and “**Jacobian**” modes generate matrices with the same dimensions and the *exact* same fill pattern. This is not generally the case, so different mappers should be created for each matrix in the problem. The last line creates a new vector by cloning the PQ vector. The X vector has the same dimension and data layout as PQ so it could be used with the **vMap** object.

Once the vectors and matrices for the problem have been created and set to their initial values, it is possible to start the Newton-Raphson iterations. The code to set up the first Newton-Raphson iteration is

```
double tolerance = 1.0e-6;
int max_iteration = 100;
ComplexType tol;

gridpack::math::LinearSolver solver(*J);
solver.configure(cursor);

int iter = 0;

X->zero();
solver.solve(*PQ, *X);
tol = PQ->normInfinity();
```

The first three lines define some parameters used in the Newton-Raphson loop. The tolerance and maximum number of iterations are hardwired in this example but could be made configurable via the input deck. The next line creates a linear solver based on the current value of the Jacobian, \mathbf{J} . The call to the configure method allows configuration parameters in the input file to be passed directly to the newly created solver. The iteration counter is set to zero and the value of \mathbf{X} is also set to zero. The linear solver is called with \mathbf{PQ} as the right hand side vector and \mathbf{X} as the solution. An initial value of the tolerance is set by evaluating the infinity norm of \mathbf{PQ} . The calculation can now enter the Newton-Raphson iteration loop

```
while (real(tol) > tolerance && iter < max_iteration) {
    factory.setMode(RHS);
    vMap.mapToBus(X);
    network->updateBuses();

    vMap.mapToVector(PQ);
    factory.setMode(Jacobian);
    jMap.mapToMatrix(J);

    X->zero();
    solver.solve(*PQ, *X);
    tol = PQ->normInfinity();
    iter++;
}
```

This code starts by pushing the values of the solution vector back on to the buses using the same mapper that was used to create **PQ**. The network then calls the **updateBus** routine so that the ghost buses have the new values the parameters from the solution vector. New values of the Jacobian and right hand side vector are created based on the solution values from the previous iteration. Note that since **J** and **PQ** already exist, the mappers are just overwriting the old values instead of creating new data objects. The linear solver is already pointing to the Jacobian matrix so it automatically uses the new Jacobian values when calculating the solution vector **X**. If the norm of the new **PQ** vector is still larger than the tolerance, the loop goes through another iteration. This continues until the tolerance condition is satisfied or the number of iterations reaches the value of **max_iteration**.

If the Newton-Raphson loop converges, then the calculation is essentially done. The last part of the calculation is to write out the results. This can be accomplished using the code

```
gridpack::serial_io::SerialBusIO<PFNetwork> busIO(128,network);
busIO.header("\n  Bus Voltages and Phase Angles\n");
busIO.header("\n  Bus Number      Phase Angle");
busIO.header("      Voltage Magnitude\n");
busIO.write();
```

The first line creates a serial bus IO object that assumes that no line of output will exceed more than 128 characters. The next three lines write out the header for the output data and the last line writes a listing of data from all buses. This completes the execute method and the overview of the powerflow application.

Advanced Functionality

The core operations supported by GridPACK™ have been described above and these can be used in to create many different kinds of power grid applications. This section will describe features that are not used as often but can be extremely useful in certain case. Additional capabilities of the GridPACK™ framework include

- Communicators and task managers that can be used to create multiple levels of parallelism and implement dynamic load balancing schemes
- A generalized matrix-vector interface to handle applications where the dependent and independent variables are associated with both buses and branches. The **MatVecInterface** described above can only be used for systems where the dependent and independent variables are only associated with the buses
- Profiling and error handling capabilities
- A hashed data distribution capability that can be used to direct network data to the processors that own the corresponding network components

This functionality is described in more detail in the following sections.

Communicators

The subject of communicators has already been mentioned in the context of the constructor for the **BaseNetwork** class. This section will describe communicators in more detail and will show how the GridPACK™ communicators can be used to partition a large calculation into separate pieces that all run concurrently. A communicator can crudely be thought of as a communication link between a group of processors. Whenever a process needs to communicate with another process it needs to specify the communicator over which that communication will occur. When a parallel job is started, it creates a “world” communicator to which all processes implicitly belong. Any process can communicate with any other process via the world communicator. Other communicators can be created by an application and it is possible for a process to belong to multiple communicators. The concept of communicators is particularly important for restricting the scope of “global” operations. These are operations that require every process in the communicator to participate. Failure of a process to participate in the operation usually results in the calculation stalling because multiple processors are waiting for a communication from a process that is not part of the global operation. A program can remain in this state indefinitely. Many of the module functions in GridPACK™ represent global operations and contain imbedded calls that act collectively on a communicator. In order for two separate calculations to proceed concurrently, they must be run on disjoint sets of processors using separate communicators.

The use of communicators to create multiple concurrent parallel tasks within an application is usually straightforward to implement but it is frequently much more confusing to understand. A diagram of a set of 16 processes that are divided into 4 groups each containing 4 processes is shown schematically in Figure 8. In this example, each subgroup could potentially execute a separate parallel task within the larger parallel calculation.




Figure 8. Schematic diagram illustrating the use of multiple communicators

Global operations on the world communicator involve all 16 processes, global operations on one of the task communicators just involve the 4 processes on the task communicator. If a network object is created on one of the task communicators, then a global operation such as the bus update only occurs between the 4 processes in the task communicator. The network object is, in a certain sense, “invisible” to the processes outside that communicator. If a network is created on a sub-communicator, then all objects derived from the network, such as factories, parsers, serial IO objects, etc. are also associated with the same sub-communicator.

The communicator supports some basic operations that are commonly used in parallel programming. GridPACK™ has been designed to minimize the amount of explicit communication that must be handled by application developers, but it is occasionally useful to have access to standard communication protocols in applications. In particular, it is useful to be able to divide a given communicator into a set of non-overlapping sub-communicators. The basic operations supported by the GridPACK™ communicator class are described below.

The GridPACK™ **Communicator** class is in the **gridpack::parallel** namespace. The basic constructor for this class creates a copy of the world communicator. The constructor has the form

Communicator(void)

and takes no arguments. Two basic functions associated with communicators are


```
int size(void) const
```

and

```
int rank(void) const
```

The first function returns the number of processors in the communicator and the second returns the index of the processor within the communicator. If the communicator contains N processes, then the rank will be an integer ranging from 0 to $N-1$. The process corresponding to rank 0 is often referred to as the head process or head node for the communicator. Note that if a process belongs to more than one communicator, its rank may differ depending on which communicator is being referred to. Information on size and rank is used extensively when explicitly programming in parallel. GridPACK™ has tried to abstract much of this programming so that developers do not need to pay attention to it, but it is still occasionally useful to be able to access these numbers. For example, the header function in the SerialIO classes is essentially equivalent to the following code fragment

```
Communicator comm;  
char buf[128];  
sprint(buf, "My message\n");  
if (comm.rank() == 0) {  
    printf("%s", buf);  
}
```

This code creates some output. If the conditional was not there, the code would print out the message from all N processors in the world communicator and N copies of “My message” would appear in the output. The conditional restricting the print statement to process 0 guarantees that the message appears only once.

A more important use of communicators is to divide up the world communicator into separate communicators that can be used to run independent parallel calculations. This is known as multi-level parallelism. Two functions can be used to split up an existing communicator into sub-communicators. The first is **split**

```
Communicator split(int color) const
```

This function divides the calling communicator into sub-communicators based on the **color** variable. All processors with the same value of the **color** variable end up in the same communicator. Thus, if 16 processors are divided up such that processes 0-3 are color 0, processes 4-7 are color 1, processes 8-11 are color 2 and processes 12-15 are color 3 then split will generate 4 sub-communicators with all the processes of the same color ending up on the same communicator. Note that this function divides the communicator completely into

complementary pieces with all processes in the old communicator ending up in a new communicator and no process ending up in more than one new communicator.

A second function that can be used to decompose a communicator into sub-communicators is **divide**. This function has the form

Communicator divide(int nsize) const

Each sub-communicator returned by this function contains at most **nsize** processes. The function will try and create as many communicators of size **nsize** as possible. For example, if the calling communicator contains 10 processes and **nsize** is set to 4, then this function will create 3 sub-communicators, two of which contain 4 processors and one containing 2 processors.

An example of how communicators can be used to create multiple levels of parallelism is illustrated in Figure 9. The example has 8 tasks that can be evaluated independently. The first row in Figure 9 shows four processors. Two of the 8 tasks are run on each processor so if each task has been parallelized then it needs to run on a communicator with only 1 processor in it. The second row shows the same calculation running on 8 processors. In this case, each processor only has 1 task associated with it but each task is still running on a single processor. If the tasks have not been parallelized, then this is as far as you can go. However, if the tasks have been parallelized, then you can move on to the configuration shown in the third line using 16 processors. In this case, the system has been divided into 8 blocks, each containing two processors. Each block has its own separate subcommunicator and each task can be run in parallel on two processors. This gives additional speed-up over what can be achieved by simply distributing tasks to separate processors.




Figure 9. Schematic diagram of 8 tasks evaluated using multiple levels of parallelism. The first row represents 8 tasks on 4 processors, the second row is 8 tasks on 8 processors and the third row is 8 tasks running on 16 processors.

Task Manager

The task manager functionality is designed to parcel out tasks on a first come, first serve basis to processes in a parallel application. Each processor can request a task ID from the task manager and based on the value it receives, it will execute a block of work corresponding to the ID. The task manager guarantees that all IDs are sent out once and only once. The unique feature of the task manager is that if the tasks take unequal amounts of time, then processes with longer tasks will make fewer requests to the task manager than processes that have relatively short tasks. This leads to an automatic dynamic load balancing of the application that can substantially improve performance. The task manager also supports multi-level parallelism and can be used in conjunction with the sub-communicators described above to implement parallel tasks within a parallel application. An example of the use of communicators and task managers to create a code that uses multiple levels of parallelism can be found in the contingency analysis application that is part of the GridPACK™ distribution.

Task managers use the `gridpack::parallel` namespace. Task managers can be created either on the world communicator or on a subcommunicator. Two constructors are available.

`TaskManager(void)`

`TaskManager(Communicator comm)`

The first constructor must be called on all processors in the system and creates a task manager on the world communicator, the second is called on all processors within the communicator `comm`. Once the task manager has been created, the number of tasks must be set. This can be done with the function

`void set(int ntask)`

where the variable `ntask` corresponds to the number of tasks to be performed. The task IDs returned by the task manager will range from 0 to `ntask-1`. The `set` function must be called on all processors in the communicator that created the task.

Once the task manager has been created, task IDs can be retrieved from the task manager using one of the functions

`bool nextTask(int *next)`

`bool nextTask(Communicator &comm, int *next)`

The first function is called on a single processor and returns the task ID in the variable **next**. The second is called on the communicator **comm** by all processors in **comm** and returns the same task ID on all processors (note that if all processors in **comm** called the first **nextTask** function, each processor in **comm** would end up with a different task ID). Both functions return true if the task manager has not run out of tasks, otherwise they return false and the value of **next** is set to -1.

The task manager also has a function

void printStats(void)

that can be used to print out information to standard out about how many tasks were assigned to each process.

A simple code fragment shows how communicators and task managers can be combined to create an application exhibiting multi-level parallelism.

```
gridpack::parallel::Communicator world
int grp_size = 4;
gridpack::parallel::Communicator task_comm = world.divide(grp_size);
App app(task_comm);
gridpack::parallel::TaskManager taskmgr;
taskmgr.set(ntasks);
int task_id;
while(taskmgr.nextTask(task_comm, &task_id) {
    app.execute(task_data[task_id]);
}
```

This code divides the world communicator into sub-communicators containing at most 4 processes. An application is created on each task communicator and a task manager is created on the world group. The task manager is set to execute **ntasks** tasks and a while loop is created to execute each task. Each call to **nextTask** returns the same value of **task_id** to the processors in **task_comm**. This ID is used to index into an array **task_data** of data structures containing the input data necessary to execute the task. The size of **task_data** corresponds to the value of **ntasks**. When the task manager runs out of tasks, the loop terminates. Note that this structure does not guarantee that tasks are mapped to processors in any fixed order. There is no guarantee that task 0 is executed on process 0 or that some process will execute a given number of tasks. If one task takes significantly longer than other tasks then it is likely that other processors will pick up work from the processors executing the longer task. This balances the workload if each process is involved in multiple tasks. Once the workload drops to 1 task per process, this advantage is lost.

Timers

Profiling applications is an important part of characterizing performance, identifying bottlenecks and proposing remedies. Profiling in a parallel context is also extremely tricky. Unbalanced applications can lead to incorrect conclusions about performance when load imbalance in one part of the application appears as poor performance in another part of the application. This occurs because the part of the application that appears slow has a global operation that acts as an accumulation point for load imbalance. Nevertheless, the first step in analyzing performance is to be able to time different parts of the code. GridPACK™ provides a timer functionality that can help users do this. These modules are designed to do relatively coarse-grained profiling, it should not be used to time the inside of computationally intensive loops.

GridPACK™ contains two different types of timers. The first is a global timer that can be used anywhere in the code and accumulates all results back to the same place for eventual display. Users can get a copy of this timer from any point in the code. The second timer is created locally and is designed to only time portions of the code. These were created to support task based parallelism where there was an interest in profiling individual tasks instead of getting timing results averaged over all tasks. Both timers can be found in the **gridpack::utility** namespace.

The **CoarseTimer** class represents a timer that is globally accessible from any point in the code. A pointer to this timer can be obtained by calling the function

```
static CoarseTimer *instance()
```

A category within the timer corresponds to a set of things that are to be timed. A new category in the timer can be created using the command

```
int createCategory(const std::string title)
```

This command creates a category that is labeled by the name in the string “**title**”. The function returns an integer handle that can be used in subsequent timing calls. For example, suppose that all calls to **function1** within a code need to be timed. The first step is to get an instance of the timer and create the category “**function1**”

```
gridpack::utility::CoarseTimer *timer =  
    gridpack::utility::CoarseTimer::instance();  
int t_func1 = timer->createCategory("Function1");
```

This code gets a copy of the timer and returns an integer handle **t_func1** corresponding to this category. If the category has already been created, then **createCategory** returns a handle to the existing category, otherwise it adds the new category to the timer.

Time increments can be accumulated to this category using the functions

```
void start(const int idx)  
void stop(const int idx)
```

The **start** command begins the timer for the category represented by the handle **idx** and **stop** turns the timer off and accumulates the increment.

At the end of the program, the timing results for all categories can be printed out using the command

```
void dump(void) const
```

The results for each category are printed to standard out. An example of a portion of the output from **dump** for the example power flow code included in the GridPACK™ distribution is shown below.

```
Timing statistics for: Total Application  
  Average time:          14.7864  
  Maximum time:          14.7864  
  Minimum time:          14.7863  
  RMS deviation:         0.0000  
Timing statistics for: PTI Parser  
  Average time:          0.1553  
  Maximum time:          1.2420  
  Minimum time:          0.0000  
  RMS deviation:         0.4391  
Timing statistics for: Partition  
  Average time:          2.8026  
  Maximum time:          2.9668  
  Minimum time:          1.7142  
  RMS deviation:         0.4398  
Timing statistics for: Factory  
  Average time:          1.2424  
  Maximum time:          1.2540  
  Minimum time:          1.2336  
  RMS deviation:         0.0056  
Timing statistics for: Bus Update  
  Average time:          0.0019  
  Maximum time:          0.0025  
  Minimum time:          0.0016  
  RMS deviation:         0.0003
```

For each category, the dump command prints out the average time spent in that category across all processors, the minimum and maximum times spent on a single processor and the RMS standard deviation from the mean across all processors. It is also possible to get more detailed output from a single category. The commands

```
void dumpProfile(const int idx) const
```

void dumpProfile(const std::string title)

can both be used to print out how much time was spent in a single category across all processors. The first command identifies the category through its integer handle, the second via its name.

Some other timer commands also can be useful. The function

double currrentTime()

returns the current time in seconds (if you want to do timing on your own). If you want to turn off profiling in a section of the code the command

void configureTimer(bool flag)

can be used to turn timing off (**flag = false**) or on (**flag = true**). This can be used to restrict timing to a particular section of code.

In addition to the **CoarseTimer** class, there is a second class of timers called **LocalTimer**. **LocalTimer** supports the same functionality as **CoarseTimer** but differs from the **CoarseTimer** class in that **LocalTimer** has a conventional constructor. When an instance of a local timer goes out of scope, the information associated with it is destroyed. Apart from this, all functionality in **LocalTimer** is the same as **CoarseTimer**. The **LocalTimer** class was created to profile individual tasks in applications such as contingency analysis. Each contingency can be profiled separately and the results printed to a separate file. The only functions that are different from the **CoarseTimer** functions are the functions that print out results. The **dumpProfile** functions are not currently supported and the **dump** command has been modified to

void dump(boost::shared_ptr<ofstream> stream) const

This function requires a stream pointer that signifies which file the data is written to.

Exceptions

The math module has been implemented so that failures throw exceptions. These can be caught by other parts of code and managed so that code does something more graceful than simply crash. For example, a calculation that fails because the solver throws an exception might try to run again using a different solver. In a contingency analysis calculation, a contingency that fails because the solver did not converge can be marked as a failed calculation and the code can proceed to the next contingency. This allows the code to evaluate all contingencies even if some don't complete because the solver fails.

A solver exception can be handled using the following construct

LinearSolver solver(*A);

```

:
try {
    solver.solve(*B,*X);
} catch (const gridpack::Exception e) {
    // Do something to manage exception
}

```

If the solve command fails, it throws a **gridpack::Exception** that can then be managed by the code. This could consist of simply exiting cleanly or the code could try and take corrective action by using a different algorithm.

Exceptions can also be added to error conditions that are detected in user written code so that the error can be picked up in some other part of the application and managed there. Exceptions have two constructors that can be used in applications

```

Exception(const std::string msg)
Exception(const char* msg)

```

where message is a text string describing the error that was encountered. This message can be read later using the function

```

const char* what() const throw()

```

The syntax for using this in an exception would then be

```

try {
    // Some action
} catch (const gridpack::Exception e) {
    std::cout << e.what() << std::endl;
    // Print error message and then take some action
}

```

Hash Distribution Module

The hash distribution functionality provides a simple mechanism for quickly distributing data associated with individual buses and branches to the processors that own those buses and branches. This situation can come up in several contexts, particularly when network data is distributed across multiple files. For example, the information on measurements in the state estimation calculation is contained in a file that is distinct from the file that holds the network configuration. The program starts by reading in the network configuration and partitioning it. The program next reads in the measurements, but there is no simple map between the measurements, each of which is associated with either a branch or a bus, and the distributed network. Even if the measurements are read in before the network is distributed, there is still no simple map between measurements and their corresponding buses and branches, since some

components may have no measurements associated with them and other components may have multiple measurements. Moving this data to the right processor and providing a simple way of mapping it to the correct bus or branch on that processor is a non-trivial task.

The **HashDistribution** module is a templated class that assumes that the data that is to be sent to the buses and branches are held in a user-defined structs. The structs used for branches can be different from the structs used for buses. If we designate the bus and branch structs by the names **BusData** and **BranchData** then the constructor for the **HashDistribution** class has the form

```
HashDistribution<MyNetwork, BusData, BranchData>  
    (const boost::shared_ptr<MyNetwork> network)
```

Both the **BusData** and **BranchData** structs must be specified when creating a new **HashDistribution** object, even if only bus or branch data is actually being used. If you are just using bus data you can simply repeat the **BusData** type in the branch slot without causing any problems. You can also use **BranchData** in both slots if you are only interested in moving data to branches.

The following command can be used to move bus data to the processors that actually hold the corresponding buses

```
void distributeBusValues(std::vector<int> &keys,  
                        std::vector<BusData> &values)
```

The integer array “**keys**” holds the original indices of the buses that the data in the vector “**values**” is supposed to map to. The keys and values vectors should be the same length and the data in the values array at index n should be mapped to the bus indicated by the original index stored at the same location in the **keys** array. This function must be called on all processors and all processors can have some initial data that needs to be mapped. The amount of data on each processor does not need to be the same and some processors, or even most of them, can have no data (it is still necessary to call the **distributeBusValues** function even if the processor contains no data). It is also possible that the same original index can appear multiple times in the keys array, i.e. multiple pieces of data can map to the same bus. On output, the values array contains all the data objects that map to buses on that processor and the keys array contains the local indices of the bus. This will include data that maps to ghost buses so a piece of data may map to more than one processor in a distributed system.

An analogous command can be used to distribute data to branches. It has the form

```
void distributeBranchValues(std::vector<std::pair<int,int> > &keys,  
                           std::vector<int> &branch_ids,  
                           std::vector<BranchData> &values)
```

Branches are uniquely identified by the buses at each end of the branch, so the **keys** array in this case is a set of index pairs consisting of the original indices of these buses. The **values** array contains the data to be distributed to the branches and the **branch_ids** array contains the local index of the branch on output. Similar to buses, multiple data items can be mapped to the same branch.

Generalize Matrix-Vector Interface

The matrix-vector interface described above is suitable for problems where the independent and dependent variables are both associated with buses but it won't work for systems where some variables are associated with branches. This can occur in optimization problems such as state estimation, where measurements are made on both buses and branches. Every measurement contributes an equation to the state-estimation optimization, which results in dependent variables associated with branches. To handle these types of problems, a more general approach to creating matrices and vectors is required. This is implemented via the GenMatVecInterface class. The BaseComponent class directly inherits from this interface, along with the MatVecInterface. The complete inheritance diagram for the components now looks like

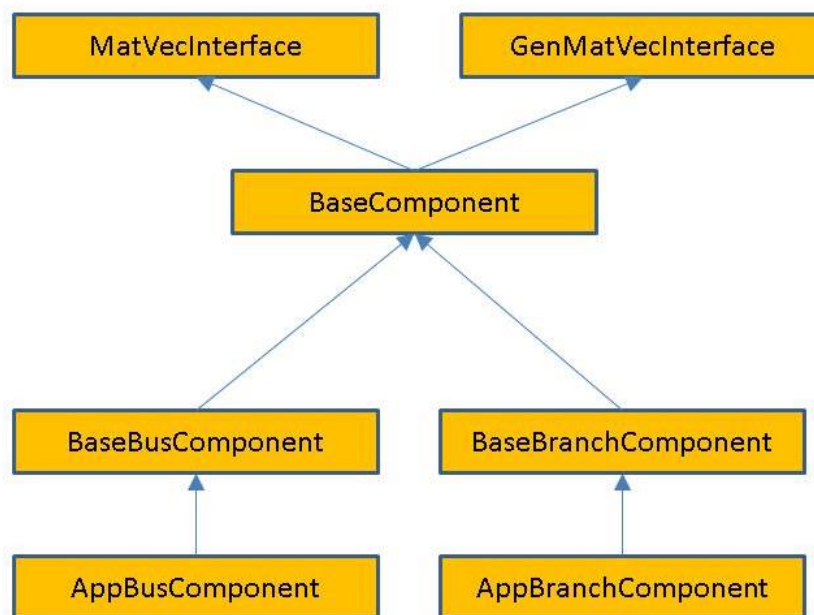


Figure 10. Inheritance diagram for network components, including the GenMatVecInterface class.

Unlike the MatVecInterface class, there is no definitive way to map which elements are contributed by a branch or bus, and the number of elements contributed by a branch or bus does not reduce to simple blocks. Thus, the idea that buses and branches contribute simple blocks of data must be abandoned. The GenMatVecInterface just assumes that buses and branches

contribute some number of equations (dependent variables) to the matrix and that they also contribute some number of independent variables to the matrix. This information is embedded in the function calls

```
virtual int matrixNumRows(void)
```

```
virtual int matrixNumCols(void)
```

These two functions specify how many dependent variables (rows) and how many independent variables (columns) are associated with a bus or branch. For the state estimation application that is currently available in the GridPACK™ release, the dependent variables are the number of measurements that are associated with the bus or branch and the dependent variables are the voltage magnitude and phase angle, which are only associated with buses. Thus, if the state estimation Jacobian is being built, the matrixNumRows function returns the number of measurements on each bus and branch. The matrixNumCols only returns a non-zero value for buses since the branches have no independent variables. This value is generally 2, if the bus has any measurements associated with it or is attached to a bus or branch that has measurements, otherwise the value is 0. If the bus has measurements and is the reference bus, then the function returns 1. These functions allow the generalized mappers to determine the dimensions of the matrix (for state estimation, the Jacobian is not necessarily square).

Unlike the original matrix-vector interface, the user has to assign the row and column indices to each matrix element. The actual values of these indices are evaluated by the mapper but it is up to the user to take the row index for a particular dependent variable (measurement) and the column index for a particular dependent variable (voltage magnitude or phase angle) and pair them with a matrix element (contribution to the Jacobian). The functions that are used for this purpose are

```
virtual void matrixSetRowIndex(int irow, int idx)
```

```
virtual void matrixSetColIndex(int icol, int idx)
```

```
virtual int matrixGetRowIndex(int irow)
```

```
virtual int matrixGetColIndex(int icol)
```

The first two functions are used by the mapper to assign indices for each of the rows and columns contributed by a component. The values of the indices need to be stored in the component so that they can be accessed by other components when evaluating matrix elements. These functions also need to be implemented by the user, since multiple matrices may be generated by the application. For example, state estimation calculation needs to be able to build the Jacobian matrix plus a diagonal matrix that represents the inverse of the uncertainties in all the measurements. The state estimation components have two modes, **Jacobian_H** and **R_inv** for each of these calculations. The matrixSetRowIndex method for the buses has the form

```
void SEBus::matrixSetRowIndex(int irow, int idx)
```

```

{
    if (p_mode == Jacobian_H) {
        if (irow < p_rowJidx.size()) {
            p_rowJidx[irow] = idx;
        } else {
            p_rowJidx.push_back(idx);
        }
    } else if (p_mode == R_inv) {
        if (irow < p_rowRidx.size()) {
            p_rowRidx[irow] = idx;
        } else {
            p_rowRidx.push_back(idx);
        }
    }
}
}

```

The row indices for the Jacobian and R^{-1} are stored in two separate STL arrays `p_rowJidx` and `p_rowRidx`. For the state estimation example, the number of rows (for both the Jacobian and R^{-1}) is equal to the number of measurements associated with the component. These measurements are held in an internal list in some order. If the number of measurements is N then the `irow` index will run from 0,... $N-1$, with the `irow` index corresponding to the `irow` element in the list. This ordering is preserved between different components. The independent variables are also assumed to be ordered in some fashion. Again, for the state estimation example, the phase angle is indexed by 0 and the voltage magnitude is indexed by 1.

The function for accessing the row indices is implemented as

```

int gridpack::state_estimation::SEBus::matrixGetRowIndex(int idx)
{
    if (p_mode == Jacobian_H) {
        return p_rowJidx[idx];
    } else if (p_mode == R_inv) {
        return p_rowRidx[idx];
    }
}

```

Again, depending on the mode, this function will return different values.

The functions that are used to actually evaluate matrix elements are

```

virtual int matrixNumValues(void) const
virtual void matrixGetValues(ComplexType *values,

```

int *rows, int *cols)

The first function returns the total number of matrix elements that will be evaluated by the component. This is used inside the mapper to allocate arrays that can actually hold all matrix elements. The second function is used to evaluate actual matrix elements along with their row and column indices. This function is the one that will make use of the **matrixGetRowIndex** and **matrixGetColIndex** functions. The evaluation of the **matrixNumValues** function can be quite complicated. For the state estimation Jacobian matrix, the number of matrix elements contributed by a component depends on the number of measurements associated with that component and the number of variables that couple to that measurement. A measurement on a bus will usually contribute two values for the independent variables on the bus, plus and additional two values for each bus that is attached to the center bus via a branch. This number will be modified slightly if one of the buses in this group is a reference bus. For branches, the number of matrix elements contributed by each measurement is approximately four, two elements for each bus at either end of the branch. This number may drop if one of the buses is a reference bus.

The **matrixGetValues** function is used to evaluate each of the matrix elements. It also gets the matrix indices for this element from the appropriate network component. The number of matrix elements returned by this function should correspond to the number returned by the **matrixNumValues** function. To see how the assignment of the indices works, we can look at the matrix element of the Jacobian corresponding to the gradient of a real power injection measurement P_i on bus i with respect to the phase angle on another bus j that is connected to i via a single branch. The contribution to the Jacobian from this measurement is given by the formula

$$\frac{\partial P_i}{\partial \theta_j} = V_i V_j (G_{ij} \sin(\theta_i - \theta_j) - B_{ij} \cos(\theta_i - \theta_j))$$

Suppose P_i is measurement k on the bus. Then the row index **im** for this matrix element can be evaluated by calling the function

im = matrixGetRowIndex(k);

The column index is associated with the phase angle measurement on the remote bus j . Assuming that a pointer (**bus_j**) to the remote bus is already available, then the column index **jm** for this matrix element could be obtained by calling

jm = bus_j->matrixGetColIndex(0);

This function is called with the argument 0 since the dependent variables are always ordered as phase angle (0) followed by voltage magnitude (1). The full list of Jacobian matrix elements can be obtained by looping over all measurements. For each bus measurement, there are

contributions from the dependent variables on each connected bus plus two contributions from the calling bus. Similarly, for each branch measurement there are approximately four contributions coming from the independent variables associated with the buses at each end of the branch. A simple counter variable can be used to make sure that the matrix element value and the corresponding row and column indices stored in the same location of the **values**, **rows** and **cols** arrays that are returned by the **getMatrixValues** function.

The **GenMatVecInterface** also includes functions for setting up vectors. These work in a very similar way to the generalized matrix functions, so they will only be described briefly. The two functions

```
virtual void vectorSetElementIndex(int ielem, int idx)  
virtual void vectorGetElementIndices(int *idx)
```

can be used to set and retrieve vector indices. In this case it is usually more convenient to get all indices associate with a component at once, so the **vectorGetElementIndices** returns an array instead of a single value. The function

```
virtual int vectorNumElements() const
```

returns the number of vector elements contributed by a component and the function

```
virtual void vectorGetElementValues(ComplexType *values, int *idx)
```

returns a list of the values along with their vector indices. Again, the index value can be obtained by first calling the **vectorGetElementIndices** function and using this to obtain the correct index for each element.

The vector interface includes one additional function that does not have a counterpart in the matrix interface. This is the function

```
virtual void vectorSetElementValues(ComplexType *values)
```

This function can be used to push values from a solution vector back into the network components. The values are ordered in the same way as the values in the corresponding **vectorGetElementValues** call, so it is possible to unpack them and assign them to the correct internal variables for each component. This function is analogous to the **setValues** call in the regular **MatVecInterface**.

The functions in the **GenMatVecInterface** are invoked in the generalized mappers. These reside in the **GenMatrixMap** and **GenVectorMap** classes. Like the standard mappers, these classes are relatively simple and contain only a few methods. The **GenMatrixMap** class consists of the constructor

```
GenMatrixMap<MyNetwork>(boost::shared_ptr<MyNetwork> network)
```

and the methods

```
boost::shared_ptr<gridpack::math::Matrix> mapToMatrix(void)  
void mapToMatrix(boost::shared_ptr<gridpack::math::Matrix> matrix)  
void mapToMatrix(gridpack::math::Matrix &matrix)  
void overWriteMatrix(boost::shared_ptr<gridpack::math::Matrix> matrix)  
void overWriteMatrix(gridpack::math::Matrix &matrix)  
void incrementMatrix(boost::shared_ptr<gridpack::math::Matrix> matrix)  
void incrementMatrix(gridpack::math::Matrix &matrix)
```

These functions all have the same behaviors as the analogous functions in the standard **FullMatrixMap**. The **GenVectorMap** class has the constructor

```
GenVectorMap<MyNetwork>(boost::shared_ptr<MyNetwork> network)
```

and supports the methods

```
boost::shared_ptr<gridpack::math::Vector> mapToVector(void)  
void mapToVector(boost::shared_ptr<gridpack::math::Vector> &vector)  
void mapToVector(gridpack::math::Vector &vector)
```

These functions have the same interpretations as the analogous functions in the **BusVectorMap** class. A new function is

```
mapToNetwork(boost::shared_ptr<gridpack::math::Vector> &vector)
```

which can be used to push data from a vector back into the network components (both buses and branches).

Fortran 2003 Interface

Most of the functionality in GridPACK™ can also be accessed through Fortran calls. The Fortran interface makes extensive use of the object-oriented features in Fortran 2003, so a compiler that supports Fortran 2003 must be used if creating Fortran applications with GridPACK™. The Fortran compiler must also support the `iso_c_binding` module, but this will usually be available if the compiler supports Fortran 2003. Most recent compilers support Fortran 2003. A working powerflow application written entirely in Fortran has been included in the current release and demonstrates how to use the current Fortran interface. The Fortran implementation is very similar to the C++ interface and most of the C++ documentation applies to the corresponding Fortran functionality. The remainder of this section will highlight the important differences between the C++ and Fortran interfaces.

Because Fortran does not have any support for templates (that we know of) the Fortran interface cannot support multiple different kinds of networks within a single application. This means that only one bus and one branch class can be present in an application, so the bus and branch classes must support all possible types of behavior. It is still possible to have more than one network in an application, but all networks must be of the same type.

The bus and branch classes in the Fortran interface are represented by the Fortran derived types **application_bus** and **application_branch**. These types have procedures bound to them, as well as internal data elements. These types are defined in the Fortran file **component_template.F90** file that is located in the **fortran/component** directory. The application bus and branch classes can be created by modifying a copy of **component_template.F90**. The functions in the math-vector interface and the component base classes are all defined in this file along with default implementations for these functions. Additional data elements and procedures can be added to the bus and branch data types to create appropriate functionality for specific problems.

A brief overview of the **application_bus** type in the **component_template.F90** file is provided here. Similar considerations apply to the **application_branch** type. To use the **component_template.F90** file it should first be copied to the directory where the application source code resides and renamed to something appropriate. We will use the name **app_component.F90**. Inside the component file, the Fortran types **bus_xc_data**, **branch_xc_data**, **application_bus**, **application_branch** are defined as part of the **application_components** module. These are the only types that need concern the application developer. There are also two types defined in this file called **application_bus_wrapper** and **application_branch_wrapper**. These are only used internally but must be defined in this file. They should not be modified. There is a line at the bottom of the **app_component.F90** file that includes an external file **component_inc.F90**. This file contains many functions that are required by the interface and must appear in the **application_components** module. However, these functions should not be modified by the user so to avoid possible errors and to simplify the file somewhat, these functions are put in the include file.

The **application_bus** type has four parts. These consist of 1) application-specific data elements, 2) data elements that must be defined in order for the component to interact with rest of the framework, 3) application-specific functions that are defined by the user and 4) framework functions that must be included in the component. The framework functions all have base implementations can be modified to suit the application. The only data elements that must be included in the **application_bus** type is a variable of type **bus_xc_data** and a pointer to this variable. The **bus_xc_data** type will be discussed further below and represents all data that might need to be exchanged in a bus update.

The framework functions are directly analogous to the functions defined for the C++ implementation and users should refer to the documentation above to find out how these functions work. This section will primarily discuss differences between the Fortran and C++ interfaces. The Fortran compilers do not have the same name-mangling capabilities as C++ so all function names are preceded by either a **bus_** or **branch_** to distinguish between bus and branch versions of the functions. A few functions only appear in the bus class or the branch class and do not necessarily need this prefix, but to be consistent, this convention is used for all functions.

Functions that are bound to the **application_bus** type are already listed in the **component_template.F90**. These functions consist of both a declaration within the **application_bus** type and a function or subroutine implementation within the **application_components** module. The declarations within the **application_bus** type (after the **contains** keyword) have the form

```
procedure::bus_matrix_diag_size
procedure::bus_matrix_diag_values
procedure::bus_matrix_forward_size
procedure::bus_matrix_reverse_size
:
```

The **procedure** keyword distinguishes a function or subroutine bound to the Fortran type from a piece of data (which is declared as a data type using one of the intrinsic Fortran data types or a Fortran type declaration).

After the type declarations within the **applications_components** module, there is a **contains** keyword followed by the subroutine and function implementations for all the procedures declared in the **application_bus** and **application_branch** types. The original implementations in the **component_template.F90** file are just stubs for these functions and typically don't do much. An example is the **bus_matrix_diag_size** function which originally has the implementation

```
logical function bus_matrix_diag_size(bus, isize, jsize)
  implicit none
  class(application_bus), intent(in) :: bus
  integer, intent(out) :: isize, jsize
  bus_matrix_diag_size = .false.
  return
end function bus_matrix_diag_size
```

The initial implementation just returns false if this function is invoked and doesn't set the variables **isize** or **jsize**. Note the argument the first item in the argument list. This is

declared as being of type **class(application_bus)** with intent in. All functions and subroutines that are bound to the **application_bus** type must have this argument, even if they do not have any other arguments. This argument provides a mechanism for accessing data items or functions that are related to a particular **application_bus** instance.

To see how the bus argument is used in actual practice, an implementation of this function in a power flow application is shown below

```
logical function bus_matrix_diag_size(bus, isize, jsize)
  implicit none
  class(application_bus), intent(in) :: bus
  integer, intent(out) :: isize, jsize
  isize = 1
  jsize = 1
  bus_matrix_diag_size = .true.
  if (bus%p_mode.eq.JACOBIAN) then
    if (.not.bus%bus_is_isolated()) then
      isize = 2
      jsize = 2
      bus_matrix_diag_size = .true.
    else
      bus_matrix_diag_size = .false.
    endif
  else if (bus%p_mode.eq.YBUS) then
    if (.not.bus%bus_is_isolated()) then
      bus_matrix_diag_size = .true.
      isize = 1
      jsize = 1
    else
      bus_matrix_diag_size = .false.
    endif
  return
endif
return
end function bus_matrix_diag_size
```

The **application_bus** implementation for power flow contains the variable **p_mode** and a user-specified function **bus_is_isolated** (this is declared as a type-bound procedure). To access this data and this function inside a type-bound procedure, use the Fortran “%” symbol. The **bus** variable in the argument list is acting in a similar way to the “**this**” pointer in C++ and refers back to the **application_bus** instance that made the original call to

bus_matrix_diag_size. Although the **bus_is_isolated** function implementation has the variable **bus** in its argument list, it doesn't need to explicitly pass this argument when making a call from an **application_bus** instance. The **bus** argument is assumed in this case. For comparison, a call to the **bus_matrix_diag_size** function, which has additional arguments, would have the form

```
ok = bus%bus_matrix_diag_size(isize,jsize)
```

Following this syntax, it is possible to construct a complete set of functions for an arbitrary application. Additional application-specific functions can be added to the component types by declaring them as procedures within the type and adding their implementations to the **application_components** module.

There are a few procedures in both the bus and branch types that should not be modified. No stubs for these appear in the component_template.F90 file. For the application_bus type, these procedures are

```
procedure::bus_get_neighbor_branch  
procedure::bus_get_neighbor_bus  
procedure::bus_get_xc_buf_size  
procedure::bus_get_xc_buf
```

For the application_branch type, the procedures are

```
procedure::branch_get_bus1  
procedure::branch_get_bus2  
procedure::branch_get_xc_buf_size  
procedure::branch_get_xc_buf
```

These procedures are required by other parts of the framework, but should not be modified by the user. Some other procedures are defined in the base class and do not appear as procedure declarations in application_bus and application_branch types. These procedures include

```
procedure::bus_get_num_neighbors  
procedure::bus_set_reference_bus  
procedure::bus_get_reference_bus  
procedure::bus_get_original_index  
procedure::bus_compare
```

for buses and

```
procedure::branch_get_bus1_original_index  
procedure::branch_get_bus2_original_index  
procedure::branch_compare
```

for branches. The `bus` and `branch` compare functions are used to determine if a bus or branch is equal to another bus or branch. An example of how to use this function can be found in the function that evaluates transformer contributions on branches for the power flow application. The syntax for calling this function is

```
double complex function branch_get_transformer(branch, bus)
:
  class(application_branch), intent(in) :: branch
  class(application_bus), intent(in) :: bus
  class(application_bus), pointer :: bus1, bus2
:
  if (bus%bus_compare(bus1)) then
:

```

In this fragment, the `bus_compare` function is being used to check if `bus1` is equivalent to `bus`. The `branch_compare` function is used in a similar way.

The only remaining issue in implementing the Fortran application bus and branch classes is understanding the exchange buffers. These buffers are declared at the top of the `component_template.F90` file as the `bus_xc_data` and `branch_xc_data` data types. Although the underlying Fortran interface implementation makes extensive use of the `iso_c_binding` module, we have worked very hard to keep the `iso_c_binding` data types out of the Fortran interface itself. However, the one place where this is not possible is in the exchange buffers, so it is important to use these data type declarations for any variables that are included in the exchange buffers. The exchange buffers are declared as follows in the top of the `component_template.F90` file

```
type, bind(c), public :: bus_xc_data
!
! Example data types. Replace with application-specific values
!
  integer(C_INT) int_reg
  integer(C_LONG) int_long
  real(C_FLOAT) real_s
  real(C_DOUBLE) real_d
  complex(C_FLOAT_COMPLEX) complex_s
  complex(C_DOUBLE_COMPLEX) complex_d
  logical(C_BOOL) log_reg
end type
```

The variables `int_reg`, `int_long`, `real_s`, `real_d`, `complex_s`, `complex_d` and `log_reg` are just examples and should be replaced with the variables used in the actual

application. Not all data types will be used in an application. Any buffer variables used in an application should use the **iso_c_binding** type declarations (**C_INT**, **C_LONG**, **C_FLOAT**, **C_FLOAT_COMPLEX**, **C_DOUBLE_COMPLEX**, **C_BOOL**). Variables declared with the **iso_c_binding** types can be cast to regular Fortran variables by relying on the compiler to automatically cast an assignment to the right sized variable. For example

```
integer f_var
integer(C_INT) c_var
:
f_var = c_var
```

If **f_var** is an 8 byte integer and **c_var** is a 4 byte integer, the compiler can be relied on to do the cast. This also works in the opposite direction, assuming that **f_var** does not exceed the capacity of a 4 byte variable.

The functions that access neighboring branches or buses also work differently than the corresponding C++ functions. Fortran does not support anything that looks like an STL vector so neighbors are accessed from buses using a two step process. The first step is to get the total number of neighbors attached to the bus using the **bus_get_num_neighbors** procedure. This allows users to set up a loop that can be used to run over either the neighboring branches or the neighboring buses that are attached to the calling bus via a single branch. The neighboring branches can then be accessed by using the **bus_get_neighbor_branch** function which returns a Fortran pointer to the neighboring branch. The syntax for using this function is

```
integer i, nbranch
type(application_branch), pointer :: branch
nbranch = bus%bus_get_num_neighbors()
do i = 1, nbranch
  branch => bus%bus_get_neighbor_branch(i)
:
```

The **bus_get_neighbor_bus** function works in a similar way and returns a pointer to the bus at the other end of branch **i**. To get pointers to the buses at either end of a branch, use the functions **branch_get_bus1** and **branch_get_bus2** procedures. Because the Fortran interface only supports one type of bus or branch per application, these functions return pointers of the correct type and there is no need to cast them to something else.

Most of the remaining differences between the Fortran and C++ interfaces are associated with the GridPACK™ factory module. As with the component classes, the Fortran interface only supports one kind of factory. This is the **app_factory** type and it can be created by copying the **factory_template.F90** file in the **fortran/factory** directory and making application-specific changes to it. The factory base class contains the functions

```
procedure::set_components
procedure::load
procedure::set_exchange
procedure::set_mode
procedure::check_true
```

These functions behave the same way as the equivalent C++ functions. In addition, the **app_factory** type contains the two functions

```
procedure::create
procedure::destroy
```

Because Fortran does not support constructors and destructors in the same way as C++, it is necessary to create explicit functions that implement whatever behaviors are imbedded in the C++ constructors and destructors. This is accomplished in the Fortran interface by adding **create** and **destroy** functions (or **initialize** and **finalize** functions) to most of the Fortran implementations of the GridPACK™ modules.

Additional methods can be added to the **app_factory** type to support application-specific functionality. An example of how to do this is the **set_y_bus** procedure for the power flow application. This subroutine is declared as a procedure in the **app_factory** type. The implementation is written as

```
subroutine set_y_bus(factory)
  class(app_factory), intent(in) :: factory
  class(application_bus), pointer :: bus
  class(application_branch), pointer :: branch
  class(network), pointer :: grid
  integer nbus, nbranch, i
  grid => factory%p_network_int
  nbus = grid%num_buses()
  nbranch = grid%num_branches()
  do i = 1, nbus
    bus => bus_cast(grid%get_bus(i))
    call bus%bus_set_y_matrix()
  end do
  do i = 1, nbranch
    branch => branch_cast(grid%get_branch(i))
    call branch%branch_set_y_matrix()
  end do
  return
end subroutine set_y_bus
```

The functions for accessing the bus and branch objects work differently from the functions that get neighboring branches or buses in the component classes. The neighbor bus and branch functions return a pointer to the appropriate bus or branch directly to the calling application. The **get_bus** and **get_branch** functions in the Fortran network class return an opaque object that cannot be directly used in a Fortran code. To convert this to a bus or branch pointer it is necessary to call the **bus_cast** or **branch_cast** functions which return an a pointer that can be called in Fortran.

The last remaining point is to provide a list of the existing Fortran modules that need to be used in a GridPACK application using the Fortran interface. These modules need to be included in any subroutine or function that is using the associated Fortran types. The existing modules are

```

gridpack_network ! type or class network
application_factory ! type or class app_factory
application_components ! type or class application_bus and
                        ! application_branch
gridpack_configuration ! type or class cursor
gridpack_full_matrix_map ! type or class full_matrix_map
gridpack_bus_vector_map ! type or class bus_vector_map
gridpack_gen_matrix_map ! type or class gen_matrix_map
gridpack_gen_vector_map ! type or class gen_vector_map
gridpack_math ! access to math initialization and
               ! finalization routines
gridpack_matrix ! type or class matrix
gridpack_vector ! type or class vector
gridpack_linear_solver ! type or class linear_solver
gridpack_nonlinear_solver ! type or class funcbuilder
                       ! and nonlinear_solver
gridpack_communicator ! type or class communicator
gridpack_parallel ! access to parallel initialization
                  ! and finalization routines
gridpack_parser ! class or type pti23_parser
gridpack_serial_io ! class or type bus_serial_io
                  ! and branch_serial_io

```

The appropriate module should be included in any function or subroutine that uses objects defined in the module. Modules can be included using the standard Fortran “**use**” statement.