2014 ADVANCED GRID MODELING PEER REVIEW PROJECT SUMMARY

Project Title:	**GridPACK™**
Organization:	Pacific Northwest National Laboratory
Presenter:	Bruce Palmer
Collaborating Organization:	Argonne National Laboratory

Project Purpose:

The power grid is considered one of the most complex machines ever devised by human beings but it is still being modeled using programs running primarily on workstations. New models incorporating higher levels of detail, more data from external sources and more complex behaviors will require much higher levels of computing power. While parallel computing holds great promise for expanding the scope of power grid modeling, the complexity of writing programs that can run on advanced architectures has been an impediment to its widespread adoption. A major goal of the GridPACK™ project is to simplify the development of parallel programs for power grid modeling while maintaining high levels of performance. GridPACK™ is designed to enable power grid engineers to develop applications quickly by supplying high level abstractions that remove many of the complexities associated with parallel computing and allowing engineers to focus on the physics and mathematics of their problems.

Technical Approach:

GridPACK™ is a framework that consists of a collection of software modules and libraries that can be used to rapidly construct applications that can run on advanced computer architectures. Furthermore, the framework has been designed to promote reuse of software components so that models that have been developed for one application can be easily made available for other applications. The framework can greatly simplify tasks such as construction of distributed networks, construction of distributed matrices and vectors, developing parallel solution algorithms and I/O of distributed data. At the same time, the GridPACK™ framework is extremely flexible with respect to the type of problems that can be addressed with it. Many different types of models and algorithms can be implemented using the same set of generic software modules. The modular structure of GridPACK™ has also made it possible to leverage the capabilities of several high performance computing libraries without necessarily committing to any one of them. Current libraries that have been incorporated into the GridPACK™

framework include PETSc, which supplies advance parallel solvers and supports distributed matrices and vectors, and Parmetis, which is used for partitioning power grid networks over multiple processors. However, other libraries could potentially be substituted for these without extensively revising the entire GridPACK™ framework. Such modifications would be transparent to any applications that have been built using the GridPACK™ modules.

GridPACK™ capabilities can be divided into three major areas. These include setup and distribution of power grid networks; math libraries that support and distributed matrices and vectors, as well as linear and non-linear solvers; and a mapper capability that maps quantities associated with network elements (buses and branches) to matrices and vectors. Much of the framework is based on software templates, which allow the user to develop arbitrary models for the buses and branches in a network. This enables GridPACK™ to model not only a wide range of different systems, but also a wide range of different analyses on those systems. A schematic diagram of GridPACK™ is shown below.

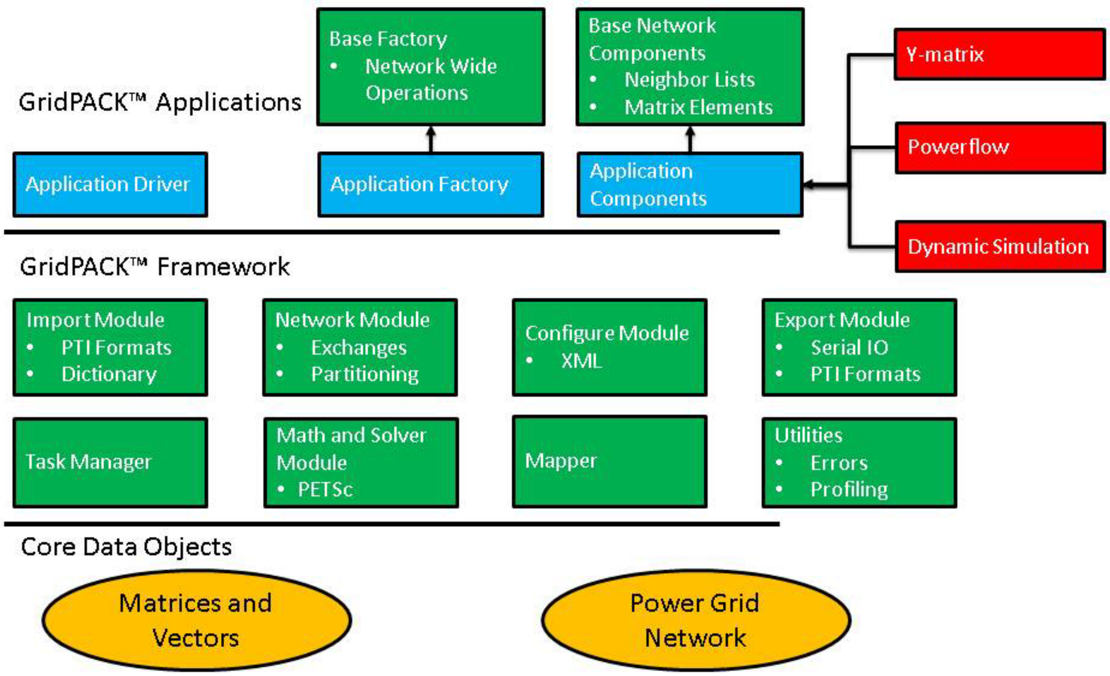

Schematic diagram of GridPACK™ software stack

Network setup and distribution is an important first step in performing parallel computations. The network must be divided by the computer, and then divided in some way between processors. It is also necessary to establish which buses and branches are connected to each other based on the information in the original network configuration file. Distribution of the network is important so that each processor can work on a fraction of the entire problem. An example of a distributed network is shown in the accompanying figure, which illustrates how a large network has been divided between 16 processors. Each color represents a portion of the network that is assigned to an individual process. Note that each process gets a portion of the network that is highly connected and has minimal connections to other processors.

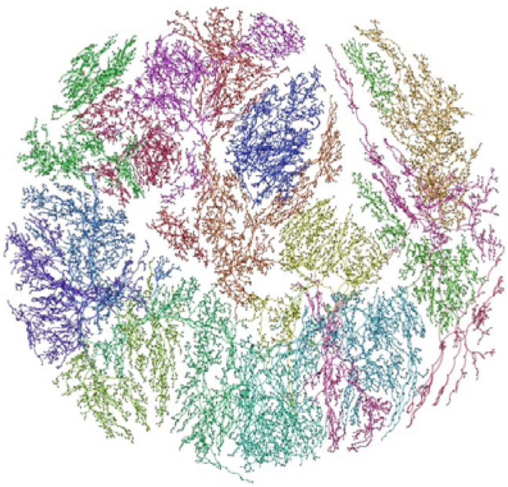

Partition of the WECC (Western Electricity Coordinating Council) network on 16 processors

This kind of partition minimizes communication and other processors, which will improve overall performance. The partition also guarantees that each process ends up with roughly equal sized portions of the network, which again will improve performance. The GridPACK™ modules can accomplish all of this with only a few lines of code.

The mapper functionality in GridPACK™ provides a mechanism for generating the matrices and vectors that appear in most equations modeling the power grid. Matrix elements are usually associated with individual network components and mapping these elements from the network to an appropriate spot in a matrix or vector is both complicated and error prone. The mapper functionality provides a simple mechanism building matrices directly from the network that abstracts the elements from the user. The user is responsible for the calculations to evaluate individual matrix and vector elements but does not need to assemble them into a matrix. This vastly simplifies the generation of algebraic equations from the network models since the evaluations of individual matrix elements are usually quite simple.

The math module allows users to create distributed matrices and vectors that can be used in parallel computations. The mapper functionality eliminates most of the need for explicitly constructing these objects, but the capability is available if needed. The math module also provides basic algebraic and initialization functions, including matrix-vector multiply, norms, addition, transposes, etc. These operations appear frequently in many algorithms. The module also contains linear and non-linear solvers that can be used to solve the algebraic equations generated by power grid models. The math module is designed as a thin wrapper on top of existing parallel solver libraries. The motivation for using a wrapper instead of calling the libraries directly is that it will be possible in the future to wrap other math libraries with the same wrapper. This will enable application developers to switch libraries seamlessly, without modifying their codes. The current math module is built on top of the PETSc solver libraries.

In addition to the modules already described, GridPACK™ offers additional components that are designed to support output, profiling, task management, etc. These can be used to implement portions of the code outside the core algorithms. Although these portions are not necessarily critical to solving the target problem, they are important in terms of getting the results out in a consistent format and benchmarking performance so that performance bottlenecks can be identified. The task manager and the GridPACK™ communicators make it possible to implement multiple levels of parallelism, which is important for applications like contingency analysis.

Several network component models that can be used across multiple applications have also been created. These include a set of Y-matrix components that have served as a base class for all other components, a set of powerflow components and a set of dynamic simulation components. The fact that these components can build on top of each illustrates how the GridPACK™ framework promotes software reuse. It also demonstrates how modifications in one of the core models (in this case, the Y-matrix) are instantly propagated to all applications that use that functionality.

2014 Results:

During FY14 the GridPACK™ project has released the initial version 1.0 of the GridPACK™ framework as well as a subsequent version 1.1 containing two additional power grid applications, as well as some additional examples that illustrate how to use basic features of the framework. The three power grid applications available for download are powerflow, dynamic simulation and static contingency analysis. The contingency analysis application required extension of the GridPACK™ communicator capability and the development of a dynamic task management module. These support both multiple levels of parallelism and a dynamic load balancing capability that reduces efficiency losses due to load imbalance. The applications developed so far demonstrate the ability of GridPACK™ to model qualitatively different problems. Preliminary performance results indicate that parallel applications developed using GridPACK™ are capable of scaling to a larger number of processors for large systems

- Powerflow: Solver scales with increasing numbers of processors
- Dynamic simulation: Overall strong scaling performance is seen to 64 cores
- Static contingency analysis: Significant performance enhancement to 128 cores.

Smaller problems remain a challenge since the overhead of parallel communication and the use of a generalized framework overwhelm the performance gains of using parallel programming for small systems. We are currently working with the solver group at Argonne National Laboratory to try and gain more performance for small systems.

In addition, we have developed a nightly build system for automated testing of GridPACK™, are working to extend this to testing on a parallel system, and have begun work on a state estimation application. Work has also started on a Fortran interface for GridPACK™.

2015 Plans and Expectations:

Plans for 2015 include

- Continuing to refine and enhance performance and robustness of existing GridPACK™ functionality
- Develop interfaces to optimize external libraries for economic analysis of power grid systems, especially mixed integer optimization models.
- Incorporate grid topological analysis into the GridPACK™ framework that identifies islands in grid networks. We will be looking for external packages and/or collaborators for this task. This is necessary for more robust contingency analysis and for dynamic systems such as cascading failures.

- Continue integrating with data management and curation tools
- Investigate incorporation of real-time data into GridPACK™ applications

Published Papers and Presentations:

"The GridPACK™ Toolkit for Developing Power Grid Simulations on High Performance Computing Platforms," Bruce Palmer, William Perkins, Kevin Glass, Yousu Chen, Shuangshuang Jin, David Callahan, Mark Rice, Ruisheng Diao, Zhenyu (Henry) Huang. Presented at 3rd International Workshop on High Performance Computing, Networking and Analytics for the Power Grid, November, 2013, Denver, CO.