Guest Editorial

Special Section on High Performance Computing (HPC) Applications for a More Resilient and Efficient Power Grid

THE POWER GRID has been evolving over the last 120 years, but it is seeing more changes in this decade and next than it has seen over the past century. In particular, the widespread deployment of intermittent renewable generation, smart loads and devices, hierarchical and distributed control technologies, phasor measurement units, energy storage, and widespread usage of electric vehicles will require fundamental changes in methods and tools for the operation and planning of the power grid. The resulting new dynamic and stochastic behaviors will demand the inclusion of more complexity in modeling the power grid. Solving such complex models in the traditional computing environment will be a major challenge. Along with the increasing complexity of power system models, the increasing complexity of smart grid data further adds to the prevailing challenges. In this environment, the myriad of smart sensors and meters in the power grid increase by multiple orders of magnitude, so do the volume and speed of the data. The information infrastructure will need to drastically change to support the exchange of enormous amounts of data as smart grid applications will need the capability to collect, assimilate, analyze and process the data, to meet real-time grid functions. High performance computing (HPC) holds the promise to enhance these functions, but it is a great resource that has not been fully explored and adopted for the power grid domain.

Computers have evolved significantly over the past decade, in an even faster pace than the power grid. Multi-core and many-core computers and computer clusters are ubiquitous today, which makes every processor to be a parallel computer. The easy gains in the past when sequential computing applications simply got faster due to the increased performance and clock frequencies of newer processors, are gone. Any performance gains for applications must be realized through the use of parallelism across multi-core processors. It is imperative to adapt or re-develop power grid software tools with explicit parallelization for massive smart grid applications.

In this special section, we cover a wide spectrum of the development of HPC applications for the power grid. This includes adapting commercial tools to run on small- or medium-scale parallel computers and massive scenario analysis on tens of thousands of computer cores. Pilot projects, demonstration or field application experiences, and HPC benefit analyses are addressed. The 17 final accepted papers can be broadly organized into three categories:

- *HPC for Transmission Applications:* At the power system transmission level, the challenges are that the computation has to be analyzed in a shorter time framework with more complex scenarios in order to make timely decisions. The eight papers in this category focus on fundamental power grid functions such as power flow solutions and transient stability simulation, extending to static and dynamic security assessment with a large number of scenarios and to optimal power flow with transient constraints. Graphical Processing Units (GPUs) are explored in addition to classic CPU-based HPC for faster numerical handling.

- *HPC for Distribution Applications:* Majority of the new technologies such as demand response, distributed generation, and electric vehicles are integrated into the distribution systems. Computation for distribution applications is challenging because of the large number of nodes involved and higher uncertainties from human behaviors. The four papers in this category focus on distribution network optimization and probabilistic analysis.

- *HPC Algorithm and Framework Development:* Some features of power systems such as sparsity and multi-spatial-temporal scales can be explored from the perspectives of algorithms and frameworks for further computational improvements. The five papers in this category study model reduction, subspace algorithms, and Benders decomposition as well as distributed and multi-scale frameworks.

The following sections provide brief introductions to all 17 papers in this special section.

I. HPC for Transmission Applications

"Parallel Transient Stability Simulation Based on Multi-Area Thévenin Equivalents": In this paper, Multi-Area Thévenin Equivalents (MATE), a diakoptics-based branch-tearing method is employed for parallelizing transient stability simulations. The equations associated with dynamic and static devices and passive networks are distributed among computing processes. The structure of the MATE algorithm provides a partitioned manner to solve both differential and algebraic equations associated with devices connected to a specific subsystem. This strategy yields a parallel-in-space algorithm for transient stability simulations, where subsystems only interact at the links level.

*"Parallel Dynamics Simulation using Krylov-Schwarz

Linear Solution Scheme": This paper presents a parallel Krylov-Schwarz linear solution scheme that uses the Krylov subspace based iterative linear solver GMRES with an overlapping restricted additive Schwarz preconditioner. Fast dynamics simulation of large-scale power systems is a computational challenge because of the need to solve a large set of stiff, nonlinear differential-algebraic equations at every time step, which involves an iterative process of solving a linear system. Performance tests of the proposed Krylov-Schwarz scheme for several large test cases ranging from 2,000 to 20,000 buses, including a real utility network, show good scalability on different computing architectures.

"Comparative Implementation of High Performance Computing for Power System Dynamic Simulations": This paper presents and compares four parallel dynamic simulation schemes in two state-of-the-art HPC environments: Message Passing Interface (MPI) and Open Multi-Processing (OpenMP). Dynamic simulation for transient stability assessment is one of the most important but computationally intensive tasks for power system planning and operation. This paper aims to speed up a single dynamic simulation by parallelizing its kernel algorithms without compromising computational accuracy. Some general guidance was derived through comparative implementation of the four schemes for maximizing the utilization of computing hardware and the performance of the simulation. The scalability and speedup performance of parallelized dynamic simulation are thoroughly studied to determine the impact of simulation algorithms and computing hardware configurations.

"A Two-layered Parallel Static Security Assessment for Large-Scale Grids Based on GPU": This paper introduces a hierarchical parallel LU decomposition algorithm based on stratified elimination tree and develops a GPU-based parallel linear system solver using this algorithm. An improved hybrid matrix ordering algorithm is proposed to reduce the height of the elimination tree. This solver is implemented for static security assessment through a two-layered coarse-grained task-parallel approach. The performance of the proposed parallel approach is 9.2 times faster, benchmarked on a Tesla C2050 GPU against a sequential implementation on a Xeon E5620 CPU.

"GPU-Accelerated Batch-ACPF Solution for N-1 Static Security Analysis": This paper presents a novel approach for accelerating N-1 stead-state security analysis by using Graphical Processing Units (GPUs). Through a GPU-accelerated QR algorithm and a novel way of processing batch ACPF on GPUs, large-scale case studies show significant time reduction as compared with CPU implementation.

"Implementation of a Massively Parallel Dynamic Security Assessment Platform for Large-Scale Grids": This paper presents a computational platform for dynamic security assessment (DSA) of large electricity power grids, developed as part of the iTesla project. It leverages HPC techniques to analyze large power systems, with many scenarios and possible contingencies, thus paving the way for pan-European operational stability analysis. A case study of the French grid is presented, with over 8000 scenarios and 1980 contingencies. Performance data of the case study (using 10,000 parallel cores) is analyzed, including task timings and data flows.

"High Throughput Computing for Massive Scenario Analysis and Optimization to Minimize Cascading Blackout Risk": This paper describes a simulation-based optimization method that allocates additional capacity to transmission lines in order to minimize the expected value of load shedding due to a cascading blackout. Estimation of the load-shedding distribution is accomplished via the ORNL-PSerc-Alaska (OPA) simulation model. Key to achieving an effective algorithm is the use of a High-Throughput Computing environment that allocates computational resources on a platform of more than 14,000 cores simultaneously among several users. The method is demonstrated through a prototype that reduces the expected load shedding by 76% by allocating only 1.1% of the installed capacity.

"Parallel Transient Stability-Constrained Optimal Power Flow using GPU as Coprocessor": This paper presents a novel GPU acceleration approach for transient stability-constrained optimal power flow (TSCOPF), one of the most computational challenging tasks in large-scale power system applications. Utilizing the two-level decomposition parallelism in the reduced-space interior point method (RIPM), GPUs serve as plug-and-play coprocessors for time-consuming linear algebra operations. Enhanced by multi-GPU processing and mixed-precision iterative refinement techniques, the computational performance is greatly improved without re-designing and re-implementing the existing algorithm framework.

II. HPC for Distribution Applications

"HPC based Intelligent Volt/VAr Control of Unbalanced Distribution Smart Grid in the presence of Noise": This paper presents a Volt/VAr optimization to minimize power loss with realistic distribution system models that include unbalanced loadings and multiphased feeders and the presence of gross errors such as communication errors and device malfunction, as well as random noise. At the core of the optimization process is an intelligent particle swarm optimization based technique that is parallelized using high performance computing techniques. Extensive experiments covering the different aspects of the proposed framework show significant improvement over existing Volt/VAr approaches in terms of both the accuracy and scalability on the IEEE 123-node and a larger IEEE 8500-node benchmark test systems.

"GPU-based Online Optimization of Low Voltage Distribution Network Operation": The parallel programming of a differential evolutionary algorithm is implemented using GPUs. The speed-up in computational efficiency is demonstrated to be suitable for online control of distributed generation and responsive demand within a low voltage distribution network. Voltage profile enhancement,

distribution network loss minimization and optimal demand response are included in the formulation of the optimization problem.

"Distributed Computing Architecture for Optimal Control of Distribution Feeders with Smart Loads": In this paper, a MapReduce based distributed computing architecture is proposed and developed in order to enable the real-time optimization of the control of distribution network feeders. The presented distributed computing approach provides processing parallelism and multi-tasking required for computationally efficient optimization of unbalanced distribution networks.

"HPC-Based Probabilistic Analysis of LV Networks with EVs: Impacts and Control": This paper investigates the adoption of HPC techniques to accelerate the probabilistic analyses of impact and control on residential low voltage (LV) networks with electric vehicles (EVs). The impacts of uncontrolled charging of EVs are quantified through a Monte Carlo-based approach using 1,000 time-series daily simulations per penetration level (i.e., 0-100%). Two real-life residential, underground UK LV networks considering realistic demand and EV load profiles (1-min resolution) are analyzed. Results show that the processing time for the impact analysis is reduced almost proportionally to the number of cores. From the control perspective, it is demonstrated that HPC can be a feasible and implementable alternative in the management of future Smart Grids.

III. HPC Algorithm and Framework Development

"Super-positioning of Voltage Sources for Fast Assessment of Wide-Area Thévenin Equivalents": This paper provides an algorithm for superimposing voltage sources and computing Thevein equivalents seen from all voltage controlled nodes in a network. It also provides a revised algorithm that exploits a close relation between super-positioning of voltage sources and Schur's complement of the network admittance matrix. The revised algorithm shows a polynomial relation between execution time and problem size.

"Fast Parallel Stochastic Subspace Algorithms for Large-Scale Ambient Oscillation Monitoring": This paper presents new formulations and computational strategies for speeding up an ambient oscillation monitoring algorithm, namely, Stochastic Subspace Identification (SSI). Two Singular Value Decomposition (SVD) approaches are applied to the SVD evaluation within the SSI algorithm. Block structures are exploited so that the large-scale dense matrix computations can be processed in parallel. This helps in memory savings as well as in overall computational time. Experimental results from three sets of archived data of the Western Interconnection demonstrate that the new approaches can provide significant speedups while retaining modal estimation accuracy.

"A New Nested Benders Decomposition Strategy for Parallel Processing Applied to the Hydrothermal Scheduling Problem": This paper explores a new nested Benders decomposition strategy suitable for parallel processing, where time-coupled stages are solved simultaneously. An alternative procedure is employed to share initial conditions for the next stages as well as Benders cuts for the previous stages. The methodology is applied to the deterministic short-term hydrothermal scheduling problem for the real large-scale Brazilian system.

"A Distributed Computing Platform Supporting Power System Security Knowledge Discovery Based on Online Simulation": A distributed computing platform, comprising massive sampling and distributed feature selection, is presented in this paper and evaluated on the Guangdong power system in China. Simulation results show that the distributed computing platform can greatly improve computing efficiency and perform better than a centralized platform.

"IGMS: An Integrated ISO-to-Appliance Scale Grid Modeling System": This paper presents a power system modelling platform which integrates transmission and distribution systems on a high performance computing platform – the integrated grid modelling system (IGMS). The system provides high resolution across a wide range from ISO markets down to end use appliances. A hierarchical co-simulation framework is used by the IGMS to connect different sub-domain models for the integrated simulation purpose. This paper shows that open-source tools such as FESTIV, MATPOWER and GridLAB-D, can be effectively integrated for co-simulation using the IGMS.

Zhenyu (Henry) Huang, *Guest Editor-in-Chief*
Pacific Northwest National Laboratory,
Richland, Washington, U.S.A.

Zeb Tate, *Guest Editor-in-Chief*
University of Toronto
Toronto, Ontario, Canada

Shri Abhyankar, *Guest Editor*
Argonne National Laboratory
Argonne, Illinois, U.S.A.

Zhaoyang Dong, *Guest Editor*
The University of New South Wales
Australia

Siddhartha Khaitan, *Guest Editor*
Iowa State University
U.S.A.

Liang Min, *Guest Editor*
Lawrence Livermore National Laboratory
U.S.A.

Gary Taylor, *Guest Editor*
Brunel University
United Kingdom