



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

High-Performance Computing for Electric Grid Planning and Operations

T. G. W. Epperly, T. Edmunds, A. Lamont, C.
Meyers, S. Smith, Y. Yao, G. Drayton

February 9, 2012

IEEE Power & Energy Society 2012 General Meeting
San Diego, CA, United States
July 22, 2012 through July 26, 2012

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

High-Performance Computing for Electric Grid Planning and Operations

Thomas Epperly, Thomas Edmunds, Alan Lamont, Carol Meyers, Steven Smith, Yiming Yao, Glenn Drayton

Abstract—High-performance computing is having a profound impact on scientific discovery and engineering in a number of areas, and researchers are beginning to study how HPC can impact energy grid planning and operations problems. Contemporary supercomputers can perform over 10^{15} floating point operations per second and have more than 300 terabytes of memory — roughly 5 orders of magnitude more than a commodity PC workstation. This level of computer power changes what is possible.

Researchers at LLNL have already used HPC systems to accelerate a multi-scenario planning study by solving a thousand unit commitment models in parallel which enabled a more detailed study than would have been achievable otherwise in the time given. LLNL is developing and testing new parallel algorithms for the unit commitment problems, including the stochastic unit commitment. These algorithms will enable higher network and time resolution and provide better solutions because of the increase model fidelity. solution of unit commitment model

Index Terms—Supercomputers, parallel machines, power generation planning, photovoltaic systems, wind power generation.

I. INTRODUCTION

SUPERCOMPUTING is having a profound impact in science discovery and engineering. Supercomputers are being used to perform climate modeling, weather prediction, astrophysical simulation, drug discovery, fusion reactor modeling, and many other important simulations. For example, researchers at Lawrence Livermore National Laboratory (LLNL) have used computational fluid dynamics (CFD) models running on supercomputers to study the aerodynamics of trucks. These studies identified ways to retrofit trucks to improve a truck's contours to result in a 17 percent efficiency gain. These changes could save the shipping industry 6 billion gallons of fuel and \$24 billion a year [1]. Similarly, Boeing has used CFD to dramatically reduce the number of wing prototypes they need to construct for testing in the development of new planes. Figure 1 summarizes decades of research and development at Boeing, and on the bottom it shows the number of prototype wings tested starting with 77 for the 767 going to 5 for the latest aircraft [2].

T. Epperly, T. Edmunds, A. Lamont, C. Meyer, S. Smith and Y. Yao are from Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94551 email: epperly2@llnl.gov

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

G. Drayton is with Energy Exemplar Pty Ltd, 3/154-160 Prospect Road, Prospect, Adelaide, SA 5082 Australia.

Manuscript received February 9, 2012.




Fig. 1. Boeing code advances showing the reduction in the number of wing prototypes that required testing. [2]

Contemporary leading supercomputers are able to perform over a quadrillion floating point operations per second (petaFLOPS) and have over 1.4 million gigabytes of memory [3]. To put this in perspective, a top-of-the-line desktop workstation achieves roughly 66.9 gigaFLOPS on the LINPACK benchmark. Thus today's fastest supercomputer has roughly 5 orders of magnitude more computing power and memory than a top-of-the-line desktop system¹.

Orders of magnitude increases in computing power changes what's possible. For example, using LLNL's BG/L supercomputer, which has 212,992 CPUs, LLNL scientists were able to perform the first micron-scale simulation of a Kelvin-Helmholtz instability using molecular dynamics. This simulation required 2.8 CPU-millennia of run time with over 1.5 CPU-millennia of continuous runtime. [5] A simulation of this magnitude is only conceivable on the largest supercomputers.

The power of supercomputers is advancing at an exponential rate. Figure 2 shows performance on the high-performance LINPACK benchmark for the fastest supercomputers for the past 19 years. The dashed line shows the trend for the top ranked machine and extrapolates the results through 2019. Current U.S. Department of Energy plans call for an exaFLOP (10^{18} FLOPS) computer to be fielded before 2020. Areas of science and engineering that are able to utilize supercomputers will be able to solve larger problems with higher fidelity models.

Research in supercomputing is also valuable because supercomputers and desktop systems are looking more and more alike. Due to the difficulty of increasing clock speeds, hardware manufacturers are expanding performance by adding more processing cores to commodity desktop systems. Supercomputer developers already have decades of experience

¹Based on a Dell Dimension T5500 with dual quad-core Intel Xeon processors running at 2.64GHz Using Intel's optimized high performance LINPACK download [4]

Performance of Top500 Supercomputing Sites Since 1993

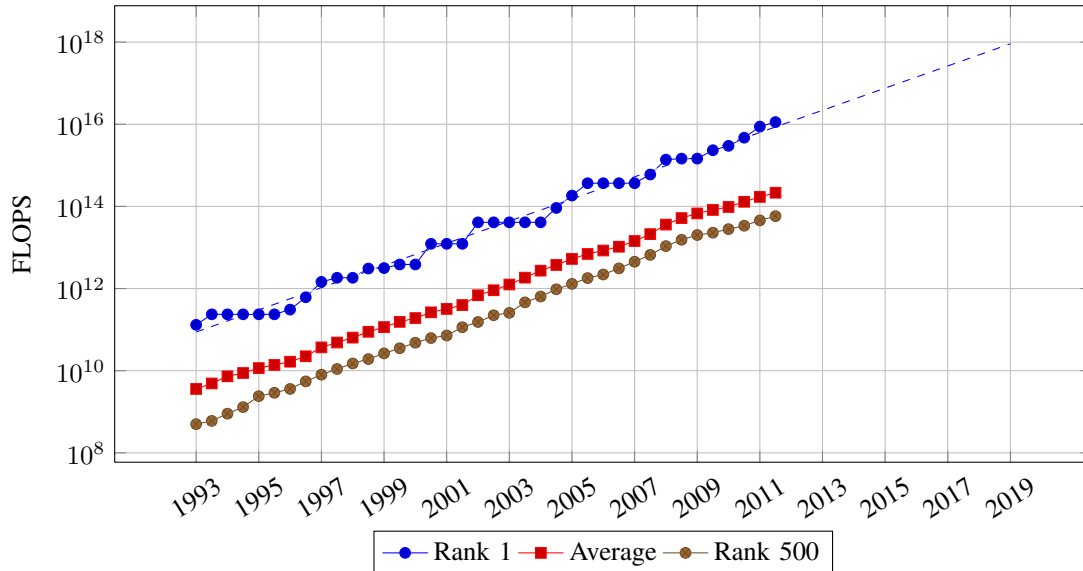


Fig. 2. Performance of top ranked, bottom rank, and average of the top 500 supercomputers based on the LINPACK benchmark. [3]

writing algorithms that can take advantage of multiple concurrent processors. Supercomputer developers are investigating technologies like OpenMP, CUDA, and OpenCL to address multi-core and GPU programming, and these tools are also available on desktop workstation systems [6]–[8].

Researchers at LLNL are researching application of supercomputers to problems in the energy grid. Section II covers a case study in partnership with Energy Exemplar where LLNL supercomputers were used to accelerate a study of the 33% renewable portfolio standard for the California Public Utilities Commission (CPUC). Section III covers our conclusions and discusses future work.

II. SUPERCOMPUTERS FOR GRID PLANNING

In late 2008, the CPUC initiated the CAISO/PLEXOS 33% Renewable Portfolio Standard (RPS) study to investigate the impact of California’s 33% renewables target. Initially, the study was a collaboration between the California Independent Systems Operator (CAISO), PLEXOS Solutions who built the model, and Southern California Edison (SCE) who was tasked with running the models. Mid-project, the CPUC contacted LLNL to investigate using supercomputers to accelerate the execution of the ensemble of simulations required for the study. LLNL partnered with Energy Exemplar, makers of PLEXOS for Power Systems®, to accelerate the solution by running the large computational simulations on LLNL’s Hyperion supercomputer and by modifying the problem formulation [9].

The CAISO/PLEXOS 33% RPS model is based on data maintained and provided by CAISO. It incorporates resource data, hourly load data for wind, import, hydro, solar, geothermal, and qualifying facility generation. The model covers 12 regions total across the WECC grid (shown in Figure 3), and it includes 42 utilities, 2157 generators, and 104 transmission lines [9].

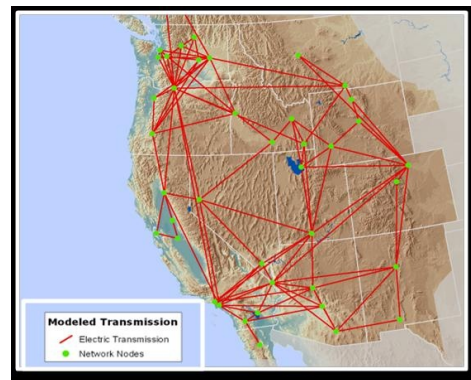


Fig. 3. Transmission grid use in the CAISO/PLEXOS 33% RPS study.

The basic form of the model is an ensemble of deterministic unit commitment models. PLEXOS reads the model definition and data, and it generates a large collection of individual unit commitment models that must be solved independently. Each unit commitment model instance is formulated as a mixed-integer program (MIP), and PLEXOS invokes a MIP solver such as FICO’s Xpress-MP or IBM’s CPLEX. MIP solution requires approximately 99% of the overall execution time.

LLNL focused on running the ensemble of unit commitment models in parallel. To accomplish this, parts of the PLEXOS system, written to run on Windows using the .Net framework, were ported to Linux using mono, a Linux/Unix implementation of the .Net framework [10]. The system was extended to utilize the a parallel computer to solve a number of unit commitment models in parallel — allocating one core to the solution of each unit commitment model. Figure 4 shows LLNL progress in improving the execution speed measured in the number of scenarios per day that could be run. The first improvement (Sept to Oct) came from porting the system to

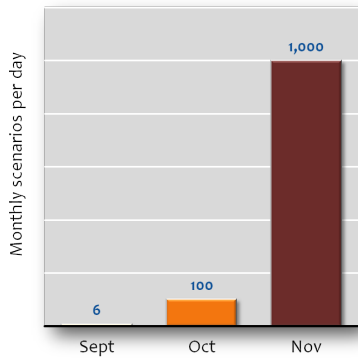


Fig. 4. Improvements in the number of simultaneous scenarios able to be executed, September–November 2010 [9].




Fig. 6. LLNL and Energy Exemplar researchers were also able to improve solution time by improving the formulations used in the PLEXOS model.

Linux, and the improvement from Oct. to Nov. was by utilizing the Hyperion supercomputer at LLNL.

The ability to run thousands of scenarios in parallel translates into the ability to perform large-scale sensitivity analyses on the data and results, a capability the study previously lacked. LLNL performed a massively parallel suite of runs varying the random seed used to generate the scenarios thus creating numerous patterns of generator outages. The results of this sensitivity study were enlightening; for some output parameters of interest, the impact of varying the generator outages caused a normally distributed pattern of ensuing results; for others, this distribution was exponential in shape; and for yet others, the distribution appeared bimodal (see Figure 5 a–c). This implies that constraining events such as generator outages to a fixed pattern can have a potentially large impact on the results. For instance, if the fixed pattern was chosen from one of those on the tail of an exponential distribution, this would produce markedly different results than a pattern from the mean. Altogether, this suggests that such sensitivity analyses may be very important in building confidence in the results obtained overall.

In addition to parallelizing the scenario modeling, Energy Exemplar and LLNL staff worked together to improve the MIP formulation underlying the PLEXOS model. Figure 6 shows the improvement in the solution times for each month of a one year simulation. The overall average is 4x speedup in execution time.

III. CONCLUSIONS AND ONGOING WORK

Overall, LLNL’s collaboration with Energy Exemplar yielded over a thousand fold increase in execution speed for the CAISO/PLEXOS 33% RPS study. It is likely that even greater increases are possible by utilizing a bigger supercomputer. By bringing more computational power to bear on the problem, we were able to provide a more thorough analysis in significantly less time than would have been required to run the simulations on standard desktop computers.

The collaboration between Energy Exemplar and LLNL is ongoing. LLNL is developing a system to enable modelers running PLEXOS on their Windows desktop to automatically launch simulations in parallel on supercomputers. The network-based system manages transferring the model details from the desktop to the supercomputer and transferring individual simulation results back to PLEXOS on the desktop. LLNL researchers have also investigated the impact on MIP solver tolerances on solution time and results. Lowering the tolerance can improve solution times with only small impacts on the simulation results.

LLNL is investigating new parallel algorithms to solve a single unit commitment model rather than an ensemble of runs. Our initial approach uses Lagrangian relaxation to decouple the MIP into a collection of independent MIP that can be run in parallel. LLNL is also beginning work on semi-definite programming to solve the optimal power flow problem. Work in both these areas is in the early phase.

REFERENCES

- [1] J. Tozzi, “Kambiz Salari and the aerodynamics of trucks,” *Bloomberg Businessweek*, July 2011.
- [2] F. T. Johnson, E. N. Tinoco, and N. J. Yu, “Thirty years of development and application of CFD at Boeing Commercial Airplanes, Seattle,” *Computers and Fluids*, vol. 34, pp. 1115–1151, 2005.
- [3] “Top500 supercomputing sites,” <http://top500.org/>.
- [4] *Intel Math Kernel Library*, Intel Corporation, 2011, <http://software.intel.com/en-us/articles/intel-math-kernel-library-documentation/>.
- [5] J. N. Glosli, K. J. Caspersen, J. A. Gunnels, D. F. Richards, R. E. Rudd, and F. H. Streitz, “Extending stability beyond CPU millenium,” in *Proceeding of Supercomputing 2007 (SC07)*, Reno, NV, November 2007.
- [6] “The OpenMP API specification for parallel programming,” <http://openmp.org/>.
- [7] “CUDA: Parallel programming made easy,” http://www.nvidia.com/object/cuda_home_new.html.
- [8] “OpenCL — the open standard for parallel programming of heterogeneous systems,” <http://www.khronos.org/opencl/>.
- [9] C. Meyers, F. Streitz, Y. Yao, S. Smith, and A. Lamont, “Using supercomputers to speed execution of the caiso/plexos 33National Laboratory, Tech. Rep. LLNL-TR-500257, 2011.
- [10] “mono,” <http://www.mono-project.com/>.




Fig. 5. Probability distribution for different output parameters in the sensitivity analysis: (a) an approximately normal distribution for the load following down price, (b) an exponential distribution for the MUNI regulation down price, and (c) a bimodal distribution for the Southern CA Edison's net export.



Thomas Epperly is the Computer Science Group Leader in the Center for Applied Scientific Computing (CASC) at the Lawrence Livermore National Laboratory (LLNL). He has worked as co-architect of Babel, the R&D 100 award winning language interoperability technology. Dr. Epperly, a Department of Energy Computational Science Graduate Fellow, earned his Ph.D. in Chemical Engineering from the University of Wisconsin and his B.S. in Chemical Engineering from Carnegie-Mellon University. His Ph.D. research focused on parallel global optimization

algorithms for nonconvex nonlinear programs. As a post-doc in the Centre for Process Systems Engineering at Imperial College in London, he worked on global optimization algorithms for reactor design under uncertainty. His undergraduate research project contributed to ASCEND, an object-oriented language for building large mathematical models. After his post-doc, he worked for Aspen Technology, Inc. on an equation-oriented modeling framework for real time optimization of chemical plant operations.



Carol Meyers has worked as a mathematician at Lawrence Livermore National Laboratory since receiving her PhD in operations research from MIT in 2006. Her expertise is in the areas of integer and linear programming optimization, agent-based modeling, discrete event simulation, and decision theory. She has been involved in energy modeling since 2009, leading an effort to port the PLEXOS power market modeling software to run on high-performance computers, in conjunction with developers at Energy Exemplar (owners of the PLEXOS code) and optimization assistance from IBM. When she is not working in the energy arena, she spends her time trying to help optimize planning of operations in the nuclear arena.



Dr. Thomas Edmunds has over thirty years of experience in designing and analyzing power systems. He has power plant design experience with Westinghouse and Bechtel, and power system operations analysis experience with the Texas Public Utilities Commission, independent power producers, and Lawrence Livermore National Laboratory, where he is currently Chief Scientist of the Systems and Decision Sciences Section. He is President of the Northern California Chapter of the Institute for Operations Research and the Management Sciences.

Dr. Edmunds earned B.S. and M.S. degrees in nuclear engineering at the University of Virginia, and an M.B.A. and Ph.D. in operations research at University of Texas.

Steven Smith received the B.S. degree in Computer Science from the University of Illinois Urbana/Champaign in 1991 with highest honors and was a recipient of the Bronze Table award. Since then he has been with Lawrence Livermore National Laboratory where he has worked on a verity of high performance computing applications and frameworks. He has served as software project lead for several projects, including Parflow, the Structured Adaptive Mesh Refinement Application Infrastructure (SAMRAI) and the Multi-Scale Epidemiologic Simulation and Analysis Decision Support System (MESA).

Yiming Yao (BSEE and PhD, Operations Research) is experienced in modeling, simulation and optimization of complex systems including communication networks and power grids. Recently he has been working on power grid planning and operations. He's interested in developing parallel optimization algorithms for solving unit commitment and economic dispatch problems that explicitly model stochastic demands and supplies. He has developed optimization software on various platforms, from desktops to supercomputers.



Alan Lamont graduated from Stanford University in 1970 with Bachelors and Masters degrees in Civil Engineering. After a term with the Peace Corps in Venezuela designing small dams, he worked for Woodward-Clyde Consultants in earthquake analysis for nuclear power plants and pipeline construction in Alaska. He returned to Stanford, and completed a PhD in Engineering Economic Systems in 1983. From 1983 to 1987 he was part of the Decision Analysis Group at Woodward-Clyde. Since 1987 he has been with Lawrence Livermore National

Laboratory working in analysis of security systems, decision analysis, risk analysis, and the economics of energy systems.

Glenn Drayton has experience in power system modelling and economics that spans 15 years. He was involved in the design of the New Zealand, Australian and Singapore power markets, in particular the energy-ancillary services co-optimisation. Glenn started Drayton Analytics (now called Energy Exemplar) in 1999 to produce a next-generation power market model based entirely on optimisation, the result being PLEXOS for Power Systems®. Glenn obtained his B.Sc.(Hons) in Operations Research, and Ph.D. in Management Science from the University of Canterbury, New Zealand. Glenn manages the development team for PLEXOS in the Adelaide, Australia office. He specialises in mathematical optimisation applied to power system planning.