

Computational and Economic Limitations of Dispatch Operations in the Next-Generation Power Grid

Victor M. Zavala, Audun Botterud, Emil Constantinescu, and Jianhui Wang

Mathematics and Computer Science Division and Decision and Information Sciences Division, Argonne National Laboratory
9700 South Cass Avenue, Argonne, IL 60439, vzavala@mcs.anl.gov, abotterud@anl.gov, emconsta@mcs.anl.gov, jianhui.wang@anl.gov

Abstract—We study the interactions between computational and economic performance of dispatch operations under highly dynamic environments. In particular, we discuss the need for extending the forecast horizon of the dispatch formulation in order to anticipate steep variations of renewable power and highly elastic loads. We present computational strategies to solve the increasingly larger optimization problems in real time. To illustrate the developments, we use a detailed dispatch model of the entire Illinois system with out-of-state wind generation.

I. INTRODUCTION

The next-generation power grid will be operated under highly dynamic regimes, including distributed storage and cogeneration, large-scale renewable generation, and highly elastic loads. These resources act as fast disturbances that need to be balanced out in the grid in real time. Wind power ramping events are already demanding more proactive and fast operational systems [14]. This situation is illustrated in Fig. 1, where we present typical profiles for the total load and wind power at different adoption levels. As can be seen, wind power supply can fluctuate by an order of magnitude in a few minutes. Similar trends are expected for the loads under smart grid environments.

Economic dispatch (ED) is one of the most important operational tasks in the power grid. The ED system updates the output levels of the committed generators to match the load demands in a cost-optimal manner. The solution has to satisfy both transmission and generation ramping constraints [19]. This task is of great importance since it clears the real time market and sets the locational marginal prices [16]. The ED system is currently designed by using forecast horizons on the order of a couple of hours with a time resolution (i.e., time steps) of a few minutes. Under stable operations, this horizon might be sufficient to capture load trends. However, in the presence of steep trends such as those observed during wind ramping events, the performance of the ED system might deteriorate if it does not have enough foresight and resolution. This can lead, for instance, to load and wind curtailment.

Increasing the foresight and resolution of the ED problem comes at the expense of additional computational complexity. The problem is usually cast as a large-scale linear or quadratic optimization problem [20]. The main source of complexity is the inherent transmission network that has to be accounted




Fig. 1. Snapshot of total load and increasing levels of wind power.

for at each time step in the horizon. In addition, coupling due to ramping constraints can render the problem intractable for even a few time steps. While state-of-the-art optimization solvers are currently able to solve large ED problems, scalability bottlenecks are still a concern.

In this work, we analyze the interactions between the economic and computational performance of ED. We first analyze the effect of increasing the horizon of the problem and discuss the associated computational implications. To this end, we present a detailed model for the Illinois system and compare the performance of barrier and simplex optimization solvers. In particular, we identify scalability bottlenecks associated with the core linear algebra kernels of the solvers. We exploit the advantages of both barrier and simplex strategies and notions of model predictive control to derive effective warm-starting strategies. The proposed developments can enable the implementation of detailed ED formulations under tight solution time constraints.

II. ECONOMIC DISPATCH PROBLEM

We consider the traditional social welfare formulation where the objective is to minimize generation costs subject to direct-current transmission and generation ramping constraints. The

ED problem has the following form:

$$\min \sum_{k=\ell}^{\ell+N} \sum_{j \in \mathcal{G}} c_j \cdot G_{k,j} \quad (1a)$$

$$\text{s.t. } G_{k+1,j} = G_{k,j} + \Delta G_{k,j}, \quad k = \ell, \dots, \ell + N - 1, j \in \mathcal{G} \quad (1b)$$

$$\begin{aligned} & \sum_{(i,j) \in \mathcal{L}_j} P_{k,i,j} + \sum_{i \in \mathcal{G}_j} G_{k,i} + \sum_{i \in \mathcal{W}_j} (W_{k,i} - \Delta W_{k,i}) \\ & = \sum_{i \in \mathcal{D}_j} D_{k,i}, \quad k = \ell, \dots, \ell + N, j \in \mathcal{B} \end{aligned} \quad (1c)$$

$$P_{k,i,j} = b_{i,j}(\theta_{k,i} - \theta_{k,j}), \quad k = \ell, \dots, \ell + N, (i,j) \in \mathcal{L} \quad (1d)$$

$$G_j^{\min} \leq G_{k,j} \leq G_j^{\max}, \quad k = \ell, \dots, \ell + N, j \in \mathcal{G} \quad (1e)$$

$$\Delta G_j^{\min} \leq \Delta G_{k,j} \leq \Delta G_j^{\max}, \quad k = \ell, \dots, \ell + N - 1, j \in \mathcal{G} \quad (1f)$$

$$P_{i,j}^{\min} \leq P_{k,i,j} \leq P_{i,j}^{\max}, \quad k = \ell, \dots, \ell + N, (i,j) \in \mathcal{L} \quad (1g)$$

$$\theta_j^{\min} \leq \theta_{k,j} \leq \theta_j^{\max}, \quad k = \ell, \dots, \ell + N, j \in \mathcal{B} \quad (1h)$$

$$0 \leq \Delta W_{k,j} \leq W_{k,j}, \quad k = \ell, \dots, \ell + N, j \in \mathcal{W} \quad (1i)$$

Here, ℓ is the real time index, k is the horizon time index, and N is the number of time steps in the horizon. The sets \mathcal{B} , \mathcal{L} , \mathcal{G} , \mathcal{W} , and \mathcal{D} are the buses, lines, thermal generators, wind generators, and load demands, respectively. Subindexed sets in j represent subsets at bus j . The problem *variables* are the thermal generation levels $G_{k,j}$, the ramp increments $\Delta G_{k,j}$, the power flows $P_{k,i,j}$, the bus angles $\theta_{k,j}$, and the wind curtailment flows $\Delta W_{k,j}$. The problem *data* are the load demands $D_{k,j}$ and wind power flows $W_{k,j}$. The multipliers of the network constraint (1c) are the locational marginal prices $\lambda_{k,j}$ at time k and bus j .

III. FORECAST HORIZON AND COST PERFORMANCE

At each time instant ℓ , the ED problem is solved by using the observed and forecasted data for the load demands and wind flows. The observed flows are $W_{\ell,j}$, $D_{\ell,j}$, while $W_{\ell+i,j}$, $D_{\ell+i,j}$, $i = 1, \dots, N - 1$ are forecast. The solution of this problem sets the generator levels for the current time step $G_{\ell,j}$ with associated cost $\varphi_{\ell}^{MH} = c_j \sum_{j \in \mathcal{G}} G_{\ell,j}$, and the locational marginal prices $\lambda_{j,\ell}$, $j \in \mathcal{B}$. At the next time step $\ell + 1$, the *true* loads and wind power flows are observed, and a new ED problem is solved to obtain $G_{\ell+1,j}$ and $\varphi_{\ell+1}^{MH}$. To manage the forecast horizon, one can use either a moving or a shrinking horizon approach. Both approaches have advantages and disadvantages from computational and implementation perspectives.

A. Moving Horizon

In the moving horizon approach, the horizon at time ℓ is $k = \ell, \dots, \ell + N$. At the next step, the horizon is shifted forward in time $k = \ell + 1, \dots, \ell + N + 1$. This approach has the computational advantage that the problem size remains *fixed*. The horizon length is usually constrained by the solution

time, which must match the time resolution (e.g., typically five minutes). From an implementation perspective, the horizon needs to be shrunk toward the end of the *bidding cycle* where the unit commitment decisions are made. The cycle is usually 24 hours. Extending the horizon over the bidding cycle can introduce a significant amount of uncertainty since the minimum power outputs are determined after bidding. Thus, if the bidding cycle contains T time steps and $\ell = 0$ at the beginning of the cycle, the horizon satisfies

$$N = \begin{cases} N & \text{if } \ell + N < T \\ T - \ell & \text{if } \ell + N \geq T \end{cases} \quad \ell = 0, \dots, T - 1. \quad (2)$$

Warm-starting moving horizon problems is complicated because of the horizon *shifting*. Reusing the solution at ℓ to initialize the problem at $\ell + 1$ is beneficial but inconsistent, thus limiting the achievable solution times.

B. Shrinking Horizon

In the shrinking horizon approach, the problem is solved for the entire bidding cycle $N = T$ and this is updated at each step by dropping only the first element of the horizon such that $N = T - \ell$, $\ell = 0, \dots, T - 1$. The advantage of this approach are that it is consistent with the bidding cycle and that it satisfies Bellman's principle of optimality [1]. Bellman's principle states that, *under perfect foresight*, the solution profile obtained with horizon $k = \ell, \dots, T$ is optimal for the problem with shrunk horizon $k = \ell + 1, \dots, T$. A disadvantage of this approach is that the bidding cycle can be extremely large compared to the time resolution of the problem. For instance, if we use a resolution of 5 minutes, a bidding cycle will contain $T = 288$ steps. However, as shown in Section IV, Bellman's principle can be exploited to derive effective warm-starting strategies that can enable the implementation of ED problems with high time resolutions and long horizons.

The perfect foresight problem with $N = T$, gives the *best possible* cost trajectory φ_{ℓ} , $\ell = 0, \dots, T$ over the bidding cycle. For the moving horizon approach, as $N \rightarrow T$, the moving horizon cost approaches the optimal cost. The convergence rate is problem dependent and thus difficult to establish *a priori*. However, this property can be exploited to derive *hybrid* moving-shrinking horizon strategies that do not need to set $N = T$ and can still exploit Bellman's principle to generate warm-starts. This approach has been proposed in the model predictive control literature [6], [21].

IV. COMPUTATIONAL ISSUES

We can write the ED problem (1) in the general form

$$\min c^T x \quad (3a)$$

$$\text{s.t. } Ax = b \quad (3b)$$

$$x \geq 0, \quad (3c)$$

where $x \in \mathbb{R}^n$ is the variable vector, $A \in \mathbb{R}^{m \times n}$ is the Jacobian matrix, $c \in \mathbb{R}^n$ is the cost vector, and $b \in \mathbb{R}^m$ is the data. Highly efficient solvers can be used to solve large-scale problems of this form. Commercial solvers include

Cplex from IBM, Gurobi, Mosek, and Knitro; non commercial solvers include Clp and Ipopt from the COIN-OR repository <http://www.coin-or.org>. Cplex, Gurobi, Mosek, and Clp are targeted to linear and quadratic optimization problems; these solvers use implementations of the primal and dual simplex method and of Mehrotra's predictor-corrector barrier method [13]. Ipopt [18] and Knitro [5] are barrier solvers for general nonlinear optimization problems. The reason for considering these nonlinear solvers in this study is that their linear algebra kernels are highly efficient, making them ideal for large-scale applications. These advantages are explained in the following subsections.

A. Simplex Methods

The simplex method starts by partitioning the variable space into *basic* and *non-basic* variables $x^T = [x_B^T x_N^T]$ with $x_B \in \mathbb{R}^m$, $x_N \in \mathbb{R}^{(n-m)}$. With this, the Jacobian matrix can be partitioned as $A = [A_B \ A_N]$, where $A_B \in \mathbb{R}^{m \times m}$ is a square matrix and $A_N \in \mathbb{R}^{m \times (n-m)}$. Similarly, the cost vector can be partitioned as $c^T = [c_B^T \ c_N^T]$. The optimality conditions of (3) are

$$c - A^T \lambda - \nu = 0 \quad (4a)$$

$$Ax - b = 0 \quad (4b)$$

$$x^T \nu = 0, x \geq 0, c - A^T \lambda \geq 0, \quad (4c)$$

where $\lambda \in \mathbb{R}^m$ and $\nu \in \mathbb{R}^n$ are the constraint and bound multipliers, respectively. In pseudo-code, the basic steps of the simplex method are as follows [2]:

- **At** iteration $k = 0$: *Start* with a non-singular basis A_B^0 , $x_N^0 = 0$, and $x_B^0 \geq 0$. If basis not available, set $A_B^0 \leftarrow \mathbb{I}_{m \times m}$
- **For** iteration $k \geq 0$:
 - 1) *Factorize* basis A_B^k using LU decomposition to obtain L_B^k, U_B^k , or *update* existing factors L_B^{k-1}, U_B^{k-1} .
 - 2) *Compute basic variables* by solving $A_B^k x_B^k = b - A_N^k x_N^k$ and *multipliers* by solving $A_B^{kT} \lambda^k = c_B$ with available factors.
 - 3) *Check* $\nu_N^k = c_N - A_N^{kT} \lambda^k \geq 0$. If it holds, *solution is optimal*; otherwise, choose any variable $x_N^{k,e}$ in x_N^k for which $\nu_N^{k,e} < 0$ as an *entering variable* for the basis.
 - 4) *Compute basis step* by solving $A_B^k \Delta x_B^k = A_N^k(:, e)$ and ratios $\Theta^k = x_B^k / \Delta x_B^k$. Here, $A_N^k(:, e)$ is the e -th column of A_N^k .
 - 5) *Apply ratio test* to find *leaving variable* $x_B^{k,l}$ with $\Theta_l^k \geq 0$ such that $x_B^{k+1} = x_B^k + \Theta_l^k \Delta x_B^k \geq 0$.
 - 6) *Update basis* by setting $A_B^k(:, l) \leftarrow A_N^k(:, e)$, set $x_B^{k,l} \leftarrow 0$, and go to next step $k \leftarrow k + 1$.

In this algorithm, the factorization step (1) is the most computationally intensive step [12], [17]. Efficient LU factorization routines (such as MA48 from Harwell [8]) are used to factorize the basis matrix, which is sparse, unsymmetric, and indefinite. The factorization time of this matrix, will increase with the horizon length and network complexity. Note that if a basis is not originally supplied, the algorithm can take a very

large number of iterations (on the order of m) to obtain a feasible basis. Consequently, a large number of factorizations and long computational times can be expected. Once a good basis matrix has been identified, strategies such as the Forrest-Tomlin and Golub-Bartels can be used to update the basis LU factors inexpensively [9], [17]. In real time applications, it is thus critical to provide the algorithm with a good starting basis.

B. Barrier Methods

Another approach to solve the problem consists of relaxing the complementarity conditions (4c) as $x^T(c - A^T \lambda) = \mu \mathbf{e}$, $\mu^k \geq 0$ and applying Newton's method directly to the nonlinear optimality conditions. Here, $\mathbf{e} \in \mathbb{R}^n$ is a vector of ones. The search step for the variables and multipliers is computed simultaneously by solving the optimality conditions for decreasing values of $\mu^k \rightarrow 0$. For fixed μ^k , the search step at iteration j is computed from the solution of the linear system:

$$\begin{bmatrix} \Sigma^j & A^T \\ A & \end{bmatrix} \begin{bmatrix} \Delta x^j \\ \Delta \lambda^j \end{bmatrix} = - \begin{bmatrix} c - A^T \lambda^j - X^j \mu^k \mathbf{e} \\ Ax^j - b \end{bmatrix}, \quad (5)$$

where $X^j = \text{diag}(x^j)$, $V^j = \text{diag}(\nu^j)$, and $\Sigma^j = X^{j-1} V^j$. The bound multipliers are recovered from $\Delta \nu^j = -X^{j-1}(\mu \mathbf{e} + V^j \Delta x^j) - \nu^j$. In the most basic setting, the Newton iterations $j > 0$ try to converge to the solution $x^*(\mu^k)$, and then μ_k is decreased. Some more advanced μ -updates can be used.

The factorization of the matrix on the left-hand side (Karush-Kuhn-Tucker matrix) is the most computationally intensive step in the algorithm. Note that this matrix is symmetric and indefinite and is much larger than the basis matrix factorized in the simplex method (i.e.; $(n+m) \times (n+m)$ against $m \times m$). In order to solve the linear system, two approaches are normally used. The first one consists of eliminating the step for the multipliers to form the normal equations

$$\left(A \Sigma^{j-1} A^T \right) \Delta \lambda^j = - \left(r_\lambda^j - A^j \Sigma^{j-1} r_x^j \right), \quad (6)$$

where $r_x^j = -(c - A^T \lambda^j - X^j \mu^k \mathbf{e})$ and $r_\lambda^j = -(Ax^j - b)$. The matrix on the left-hand side is known as the normal matrix. The step for the primal variables is recovered from $\Delta x^j = \Sigma^{j-1} (r_x^j - A^T \Delta \lambda^j)$. If the Jacobian matrix A is full rank, then the normal matrix is positive definite. This enables the application of a Cholesky factorization to obtain factors of the form L^j and L^{jT} . Even though the normal matrix is significantly smaller ($m \times m$) than the original KKT matrix, forming the normal system might destroy the sparsity of the original KKT matrix, making the Cholesky factorization inefficient. This is the strategy used in most barrier solvers specialized for linear optimization problems such as Cplex [4], [3] and Clp.

A more efficient approach that can be used to factorize the KKT matrix consists of directly applying a saddle-point solver such as MA57 [7] or Pardiso [15]. Saddle-point solvers have advanced significantly in the last few years and are capable of




Fig. 2. Illinois system. Gray dots are generation buses.

solving very large systems efficiently. The key to this approach is the ability to preserve and exploit the high degree of sparsity of the KKT matrix.

A fundamental problem of interior-point solvers is that they cannot exploit warm-start information efficiently [10]. The reason is that barrier methods proceed progressively from the interior toward the boundary of the feasible region. On the other hand, this strategy also makes the number of iterations insensitive to the problem size and number of variable bounds. Consequently, these solvers are much more efficient than the simplex method when no warm-start information is provided. In the following section, we evaluate the performance of barrier and simplex methods on large-scale ED problems. In addition, we propose strategies to exploit the advantages of these competing approaches to accelerate the solutions.

V. ILLINOIS SYSTEM SET UP

We have built an ED model using real data for the Illinois system. The system comprises 1900 buses, 2538 transmission lines, 870 load nodes, and 261 generators. Our data consists of detailed specifications for the network topology, ramp and generation limits, fuel costs, and transmission lines. The Illinois system is sketched in Fig. 2. We have added artificial wind power data in out-of-state buses to simulate a nominal wind power adoption of 10%.

A. Economic Issues

We first analyze the effect of increasing the forecast horizon in the ED formulation. We run the system using a moving horizon approach for a single bidding cycle. We assume that all the units are committed. In Fig. 3, we plot cost savings as a function of horizon length using a one-hour horizon as the reference. We use a time resolution of one hour. As can be seen, significant savings can be realized by extending the horizon over 8 hours. The optimal cost can be reached with




Fig. 3. Effect of horizon length on cost savings.

an horizon of around 10 hours. The savings over the bidding cycle are around \$100,000 this represents around 1% of the daily generation cost. We have observed that the magnitude of the savings depends on the ramp constraints, the initial conditions for the generators, and the variability of the wind power and loads. Consequently, while the overall trends are realistic, assessing the full potential savings requires a more detailed study.

B. Computations Issues

The previous study suggests that increasing the horizon of the ED formulation can bring increased performance. In Table I and Fig. 4 we present the problem dimensions as the horizon increases. In addition, we present solution times with no warm-start for the Ipopt (version 3.8) and Cplex (version 12.2) solvers. The Harwell subroutine MA57 was used for factorization of the KKT matrix in IPOPT. The best reordering strategy was nested dissection, implemented in the Metis package [11]. The dual simplex method was used in Cplex (Cplex-Dual). All calculations were obtained on a quad-core Intel processor running Linux at 2.4 GHz.

The size of the problem increases linearly with the horizon length. A problem with 24 time steps already contains more than 100,000 variables. Most of the complexity comes from the network constraints. We observe, however, that despite the network complexity, the problems are very sparse and the sparsity increases with the problem size. Ipopt has been found to be significantly more efficient than Cplex-Dual in the case where no warm-start is supplied. In particular, for a problem with 24 time steps, the solution time of Ipopt is less than 3 minutes, while that of Cplex is more than 11 minutes. The largest problem solved with Ipopt contains 48 time steps and 205,000 variables and can be solved in less than 10 minutes. We point out that the barrier method implemented in Cplex was not competitive in solution time and robustness. For instance, the solution of an ED problem with three time steps using Cplex-Barrier takes around two

TABLE I
COMPUTATIONAL PERFORMANCE OF OPTIMIZATION SOLVERS (NO WARM-START).

N	n	m	Nonzeros Jacobian [%]	Ipopt CPUs - Iter	Cplex-Dual CPUs - Iter
1	4272	4009	0.068	0.7 -22	0.3 -1154
3	12816	12027	0.024	4.2 -36	4.5 -5100
5	21360	20045	0.014	9.5 -41	17.6 -10312
10	42720	40090	0.007	36.0 -46	120.1 -25427
12	51264	48108	0.006	42.0 -43	181.7 -31191
16	68352	64144	0.005	94.1 -51	344.7 -46099
20	85440	80180	0.004	110.4 -47	600.5 -63737
24	102528	96216	0.003	163.8 -50	679.3 -68540




Fig. 4. Effect of horizon length on computational performance of optimization solvers (no warm-start).

minutes. This can be attributed mainly to the linear algebra kernel and ill-conditioning of the KKT matrix. Unfortunately, since the Cplex output display is limited, it is difficult to pinpoint performance bottlenecks.

In Table II and Fig. 5 we present the performance of the solvers when warm-start is provided. In this set up, the problems are solved with nominal wind power values to obtain the initial solution. The wind power outputs are then perturbed by 10% of their nominal value. The rationale behind this set up is that the solutions of the previous cycle at the same time can be used. In other words, this warm starting approach exploits the natural periodicity of the load in the bidding cycle can be exploited. When warm-started, Cplex significantly outperforms Ipopt. In particular, note that for a problem with 24 time steps, Cplex takes 28 seconds, while Ipopt takes more than 2 minutes. We observe that Cplex requires only 9 refactorizations of the basis matrix. Another interesting observation is that the factorization times of the KKT matrix are lower than those of the basis matrix, despite the fact that the KKT matrix is twice as large. This clearly illustrates the efficiency of MA57 and the Metis reordering.

In Fig. 6 we analyze the robustness of the warm-starts provided to the solvers. We solve a problem with six time

TABLE II
COMPUTATIONAL PERFORMANCE OF OPTIMIZATION SOLVERS (WARM-START). TOTAL WIND POWER PERTURBATION OF 10%.

N	Ipopt CPUs - No. Refactorizations	Cplex CPUs - No. Refactorizations
1	0.9 - 14	0.2 - 1
6	17.67 - 38	4.85 - 8
12	47.82 - 43	14.95 - 9
18	118.33 - 46	16.06 - 8
24	135.45 - 47	28.28 - 9

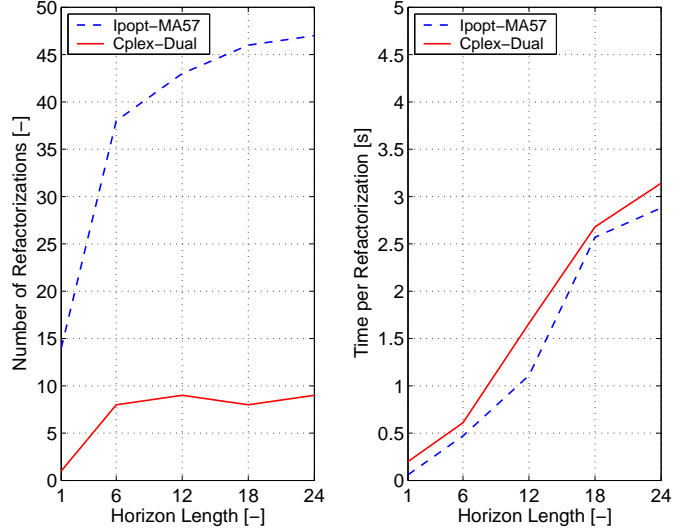


Fig. 5. Effect of horizon length on computational performance of optimization solvers (with warm-start).

steps and perturb the wind power profiles by 10, 20, 30, 40, and 50% of the nominal values. We have found that the basis matrix can keep the solution times of Cplex relatively stable despite the strong perturbations. Similar behavior has been observed for perturbations in the loads. For a problem with 12 time steps and a perturbation of 10% in all the bus loads, the solution time goes down from 442 seconds with no warm-start to 37 seconds with warm-start. In the warm-started case, only 883 dual simplex iterations and four refactorizations are needed. For a problem with 16 time steps the solution time can be reduced from 350 seconds to 50 seconds with only 4 refactorizations needed. Note that a perturbation of 10% in the loads is much larger than the forecast errors observed in real operations (2-5%). This result is important because it suggests that we can construct basis matrices in advance (e.g., one day ahead using the forecasted load) and reuse them in real time to accelerate the solutions. Moreover, the warm-start basis matrix can be constructed with barrier solvers such as Ipopt or Knitro, and this can be fed to the simplex solver to perform fast, real time LU updates.

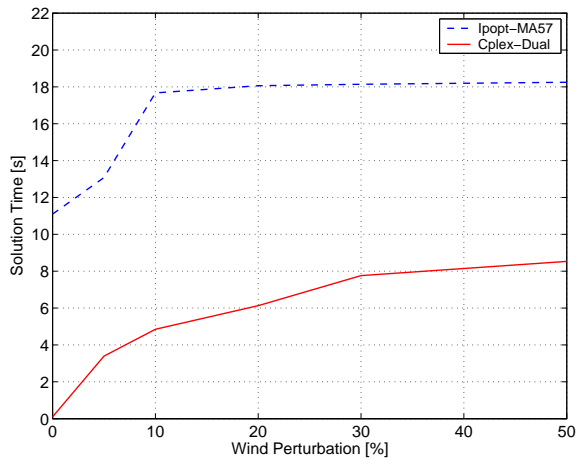


Fig. 6. Effect of data perturbations on solution times. Case with six time steps in horizon.

VI. CONCLUSIONS AND FUTURE WORK

We have presented a preliminary evaluation of the effects of increasing the detail of dispatch formulations. In particular, it is clear that longer horizons are needed in more dynamic operations such as those expected in the next-generation grid. We have tested the performance of two state-of-the-art solvers implementing barrier and simplex methods in a large-scale systems. We have found that the basis matrix in the simplex method is robust to data perturbations. In addition, we have found that barrier solvers that directly factorize the Karush-Kuhn-Tucker matrix scale well in large-scale problems. These complementary advantages can be used to derive warm-starting strategies to avoid computational bottlenecks. For instance, we suggest that warm-start basis matrices should be constructed one bidding cycle in advance using forecast information and reused in real time. The presented computational analysis also sets a reference for the expected performance of state-of-the-art solvers. This is important in moving to more complex dispatch formulations, including real time unit commitment, storage, and transmission switching decisions.

ACKNOWLEDGMENT

This work was supported by the U.S. Department of Energy through contract DE-AC02-06CH11357.

REFERENCES

- [1] R. Bellman. *Dynamic Programming*. Dover, New York, 1972.
- [2] R. E. Bixby. Implementing the simplex method: The initial basis. Technical Report ADA453079, Rice University, <http://handle.dtic.mil/100.2/ADA453079>, 1991.
- [3] R. E. Bixby. Solving real-world linear programs: A decade and more of progress. *Operations Research*, 50:3–15, 2002.
- [4] R. E. Bixby, J. W. Gregory, I. J. Lustig, R. E. Marsten, and D. F. Shanno. Very large-scale linear programming: A case study in combining interior point and simplex methods. *Operations Research*, 40:885–897, 1992.
- [5] R. H. Byrd, J. Gilbert, and J. Nocedal. A trust-region method based on interior-point techniques for nonlinear programming. *Mathematical Programming*, 89:149–185, 2000.

- [6] M. Diehl, H. J. Ferreau, and N. Haverbeke. Efficient numerical methods for nonlinear MPC and moving horizon estimation. In *Nonlinear Model Predictive Control*, pages 391–417, 2009.
- [7] I. S. Duff. Ma57 - a code for the solution of sparse symmetric definite and indefinite systems. *ACM Transactions on Mathematical Software*, 30:118–144, 2004.
- [8] I. S. Duff. and J. K. Reid. The design of ma48: a code for the direct solution of sparse unsymmetric linear systems of equations. *ACM Transactions on Mathematical Software*, 22:187–226, 1996.
- [9] J. J. H. Forrest and J. A. Tomlin. Updated triangular factors of the basis to maintain sparsity in the product form simplex method. *Mathematical Programming*, 2:263–278, 1972.
- [10] J. Gondzio and A. Grothey. Reoptimization with the primal-dual interior point method. *SIAM J. Opt.*, 13:842–864, 2003.
- [11] G. Karypis and V. Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput.*, 20:359–392, 1999.
- [12] R. Luce, J. Duintjer, M. Groetschel, T. Koch, J. Liesen, R. Nabben, and O. Schenk. On the linear algebra kernel of simplex-based LP solvers. Technical report, Technical University of Berlin, Preprint, <http://www.math.tu-berlin.de/~nabben/Publikation/lasimp08.pdf>, 2008.
- [13] S. Mehrotra. On the implementation of a primal-dual interior point method. *SIAM Journal on Optimization*, 2:575–601, 1992.
- [14] C. Monteiro, R. Bessa, V. Miranda, A. Boterrud, J. Wang, and G. Conzelmann. Wind power forecasting: state-of-the-art 2009. Technical report, INESC Porto and Argonne National Laboratory, 2009.
- [15] O. Schenk, A. Wächter, and M. Hagemann. Matching-based preprocessing algorithms to the solution of saddle-point problems in large-scale nonconvex interior-point optimization. *J. Comp. Opt. and App.*, 36:321–341, 2007.
- [16] M. Shahidehpour, H. Yamin, and Z. Li. *Market Operations in Electric Power Systems: Forecasting, Scheduling, and Risk Management*. Wiley, New York, 2002.
- [17] L. M. Suhl and U. H. Suhl. A fast LU update for linear programming. *Annals of Operations Research*, 43:33–47, 1993.
- [18] A. Wächter and L. T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106:25–57, 2006.
- [19] C. Wang and S. M. Shahidehpour. Effects of ramp-rate limits on unit commitment and economic dispatch. *IEEE Transactions on Power Systems*, 8:1341–1350, 1993.
- [20] A. J. Wood and B. F. Wollenberg. *Power Generation, Operation and Control*. Wiley, New York, 1994.
- [21] V. M. Zavala, C. D. Laird, and L. T. Biegler. Fast implementations and rigorous models: Can both be accommodated in NMPC? *Int. J. Robust Nonlinear Control*, 18:800–815, 2008.

(To be removed before publication) The submitted manuscript has been created by the University of Chicago as Operator of Argonne National Laboratory (“Argonne”) under Contract No. DE-AC02-06CH11357 with the U.S. Department of Energy. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.