A Fast Distributed Implementation of Optimal Power Flow

Ross Baldick, Balho H. Kim,* Craig Chase, Yufeng Luo†
The University of Texas at Austin, TX 78712

ABSTRACT

We describe a distributed implementation of optimal power flow on a network of workstations. High performance is achieved on large real-world systems, including a 2587 line representation of the ERCOT system. The approach illustrates a general framework for parallelizing power system optimization problems.

1 INTRODUCTION

In previous work [8], we described an algorithm that allowed parallel computation of optimal power flow (OPF). In the present paper, we present results from a fast implementation of that algorithm on a network of workstations using an interior point OPF code [16] as the basic OPF engine. We present several cases that we have studied including a full AC OPF solution for essentially the whole of the Electric Reliability Council of Texas (ERCOT) system.

Many traditional approaches to parallelization of power flow problems involve parallelizing the computations associated with factorization [4, 10, 12, 14]. Typically, this requires all the data to be assembled centrally in a unified database. However, any such approach inherently limits the possibilities for distributed implementation because of the difficulties with assembling and maintaining a central database. Particularly if calculations are to span the systems of several companies, then the issue of combining several incompatible databases can be extremely problematic.

In contrast, our approach involves decomposing a large power system geographically into regions and performing computations for each region in parallel. Unlike diakoptics approaches [3, 9], which do not appear to have natural implementations in an optimization context, our decomposition fits naturally into an optimization framework. Communication between regions is modest so that the decomposition is ideally suited to distributed computation where a separate processor is assigned to each regional calculation. Each regional processor only needs to access the generation, demand, and network data for its own region. The regions can be chosen to correspond to existing utility service areas, so that the problem of combining disparate databases is obviated.

Under this arrangement, each region performs a calculation for itself. In the event of a processor failure in one region or of a communication failure between regions, the remaining regions can continue to function, albeit at possibly higher operating costs.

The decomposition we use could be applied to any nonlinear optimization problem that involves power flow equations as equality constraints. This approach may prove essential to the solution of large-scale security assessment by independent system operators (ISOs).

In Section 2 we present the basic decomposition approach. In Section 3 we describe the modifications to the interior point OPF code that we utilized for our fast distributed implementation. In Section 4 we describe the environment for running the distributed computations, while in Sections 5 and 6, respectively, we describe the test systems and the results. We conclude in Section 7.

2 DECOMPOSITION SCHEME

The decomposition approach has been described previously in [7, 8] and is based on the "auxiliary problem principle" [1]. The basic approach is to divide the overall system and corresponding overall OPF problem into geographical regions. Any transmission line that crosses between two adjacent regions is conceptually divided into two lines by adding a "dummy bus" at the border between the two regions. Real and reactive power flow variables and voltage magnitude and angle variables are defined for the dummy bus and these four variables are duplicated, with one copy assigned to each area.

To recover the overall OPF formulation, we must require that each of the two corresponding duplicated variables has the same value. That is, for each dummy bus, we must add four constraints into the problem to enforce equality between duplicated variables. We dualize these equality constraints using an augmented Lagrangian approach and then linearize the cross-terms in the augmented Lagrangian according to the auxiliary problem principle [1].

This approach decomposes the overall problem into a

*Balho Kim is now at Hong-Ik University, Seoul, South Korea.
†Yufeng Luo is now at Synopsys Inc., Mountain View, California.

PE-167-PWRS-0-04-1998 A paper recommended and approved by the IEEE Power System Analysis, Computing and Economics Committee of the IEEE Power Engineering Society for publication in the IEEE Transactions on Power Systems. Manuscript submitted December 5, 1997; made available for printing April 24, 1998.

set of optimization problems, one for each region. The solution of the regional optimization problems is alternated with an update of the Lagrange multipliers on the equality constraints. The objective for each regional optimization problem is very similar to the objective of a standard OPF for the region alone, neglecting the rest of the system. In particular, the cost of each generator within the region is represented in the objective. However, in addition to these costs, there are terms in the objective for "dummy generators" located at each of the dummy buses. Each dummy generator has a cost function that is convex and separable quadratic in the four variables associated with its bus: real and reactive power, voltage magnitude and angle.

For example, for a system consisting of two adjacent regions, a and b, with a single tie-line joining them, the decomposition scheme takes the form ([8]):

$$(x^{k+1}, y_a^{k+1}) = \underset{(x,\, y_a) \in A}{\operatorname{argmin}} \{c_a(x) + \frac{\beta}{2} \left\| y_a - y_a^k \right\|^2 + y_a{}^\dagger[\gamma(y_a^k - y_b^k) + \lambda^k]\}, \quad (1)$$

$$(y_b^{k+1}, z^{k+1}) = \underset{(y_b,\, z) \in B}{\operatorname{argmin}} \{c_b(z) + \frac{\beta}{2} \left\| y_b - y_b^k \right\|^2 + y_b{}^\dagger[\gamma(y_b^k - y_a^k) - \lambda^k]\}, \quad (2)$$

$$\lambda^{k+1} = \lambda^k + \alpha(y_a^{k+1} - y_b^{k+1}), \quad (3)$$

where:

- the superscript k is the iteration index and we are presenting the computations performed for the k-th iteration,
- x and z are vectors of the voltage angles and magnitudes and real and reactive power generations in areas a and b, respectively, excluding the border variables,
- y_a and y_b are, respectively, the duplicate copies of the border variables. That is, they are vectors of real and reactive power flow and voltage angle and magnitude at the dummy bus on the tie-line between the regions,
- A and B are feasible sets for the OPF problems for regions a and b, respectively,
- c_a and c_b are the costs of generation in areas a and b, respectively,
- $\|\bullet\|^2$ is the square of the L_2 norm and evaluates the sum of the squares of the components of its argument,
- λ is the vector of Lagrange multipliers,
- superscript $\dagger$ denotes transpose, and
- α, β, and γ are suitable positive constants.

The updates in (1) and (2) represent the regional optimization problems for a and b. The Lagrange multipliers are updated according to (3). Consider the problem for region a. In addition to the cost of generation, c_a, in region a there is a term in the objective of (1) given by:

$$\frac{\beta}{2} \left\| y_a - y_a^k \right\|^2 + y_a{}^\dagger[\gamma(y_a^k - y_b^k) + \lambda^k].$$

This term is a separable quadratic function of the real and reactive power flow and the voltage magnitude and angle at the dummy bus. It is the objective of the dummy generator for region a. Region b has an analogous term in its objective.

The objectives of the dummy generators represent the interactions of a region with the rest of the system, but can be interpreted as the cost function of generators at the border of the region. For this reason, each region can utilize modified versions of existing OPF software to implement the system-wide distributed OPF.

It is natural to think of solving each regional problem on a separate processor. For example, in the two region case, problem (1) would be solved on the region a processor, while problem (2) would be solved on the region b processor. Inter-processor communication consists solely of transmitting the most recent values of the border variables y_a^k and y_b^k calculated for the dummy generators to the adjacent regions. This information is used to update the cost function of the dummy generators and to update the Lagrange multipliers.

The OPFs solved in successive iterations have the same constraints. The OPFs differ in their objectives only in the terms corresponding to the dummy generators. At each iteration after the first, the OPF algorithm can be started with a feasible initial solution, namely, the solution obtained in the previous iteration. The availability of a previous solution allows rapid convergence to the solution of each iteration subsequent to the first.

In previous work [8], we discussed conditions for convergence. We also demonstrated empirically that for a variety of case study systems the decomposition scheme takes only 3–5 iterations to converge to essentially the same generation dispatch pattern as provided by an undecomposed OPF algorithm applied to the overall system.

If the regions are chosen carefully to have few border variables, then the inter-processor communication requirements are modest. In the case of a multi-utility system, the natural regionalization is to divide the system into its constituent utilities since there are usually relatively few tie-lines *between* utilities compared to the number of transmission lines *within* a utility service area. Further details concerning the decomposition as well as preliminary convergence results based on an inefficient implementation of OPF are described in [7, 8].

3 Interior point OPF

We obtained the interior point OPF (IPOPF) code that was described in [16]. That software implements a non-contingency constrained OPF with an objective considering real power only. In this section we describe the modifications to the source code of the IPOPF to allow it to be used in the decomposition scheme.

As mentioned in Section 2, the decomposition scheme utilizes dummy generator cost functions that are separable and quadratic in real and reactive power and in voltage angle and magnitude. To implement this more comprehensive cost function for the dummy generators required changes to the IPOPF code in two main areas:

1. we augmented the data file specification to accommodate the extra cost information and modified the input data routine, and

2. we altered the code that calculated the terms in the update equations (equation (9) in [16]) to reflect the new objective. In particular, this involved changing the code that calculated the gradient and the Hessian of the Lagrangian corresponding to the OPF problem.

Further details can be found in [6], which can be obtained from the first author of the present paper.

The decomposition approach also requires that the OPF be run repeatedly on OPF problems having, at each iteration:

- the same constraints,
- the same cost function for physical generators, but
- slightly different cost data for the dummy generations.

We altered the IPOPF code so that:

- the IPOPF could be called from a calling program,
- the cost functions could be modified for the successive calls from the calling program, and
- IPOPF could be restarted from a previously solved case.

Further details concerning the modifications are contained in [5], which can be obtained from the first author of the present paper.

4 Distributed environment

We used the Parallel Virtual Machine (PVM) system, version 3.3 [2] to run the distributed OPF algorithm on a network of workstations. PVM is a freely available software library for writing parallel programs, and was developed at the Oak Ridge National Laboratory. PVM has a user community estimated in the tens of thousands. There are many software tools available for developing and tuning PVM programs. Detailed information on the use of PVM can be found in [2] and at the PVM home page [13].

For our purposes, we required only a small subset of the PVM functionality. Specifically, we utilized PVM functions to:

- create a process on a remote machine in the network,
- pack data in a machine-independent binary format, and
- send and receive messages using conventional message passing semantics.

To perform the experiments, we used a network of seven Sun UltraSparc^{tm} workstations each with 128 megabytes of RAM and running the Solaris^{tm} operating system (SunOS 5.5.1). For the full ERCOT system with eight regions, we used one additional computer, a Sun Sparc 20^{tm}.

A common source program was run on every workstation, with one workstation assigned to each region. Different data files were processed by each workstation, with each data file including the generator, demand, and network data for the corresponding region. This methodology is often referred to as single program, multiple data (SPMD) parallel computing.

An execution of the system begins by starting a master program on one machine in the network. This machine uses the PVM facilities to create slave processes on each of the (up to seven) other workstations. The master and each of the slaves then reads its individual data file. From this point onwards the processes interact as peers.

The computation progresses in iterations. At each iteration, processes modeling adjacent regions communicate by sending one message to each neighbor containing the values of real and reactive power flow and voltage angle and magnitude for each tie line between the respective regions. The iterative process is terminated when the stopping criterion to be described in Section 6 is satisfied.

5 Test systems

We used nine test systems, seven of which have been described previously in [8]. Test systems 1 through 5 represent various combinations of IEEE Reliability Test Systems (RTS). Systems 6 through 8 represent parts of the ERCOT system divided into regions according to utility geography boundaries. Test system 9 models virtually the complete ERCOT system divided into the major utilities in ERCOT. To the best of our knowledge, our study is

Buses	Regions	Core Buses	Ties	Lines	Load
50	2	24, 24	2	80	50
78	3	24, 24, 24	6	126	74
108	4	24, 24, 24, 24	12	186	100
238	2	118, 118	2	376	76
360	3	118, 118, 118	6	570	126
376	2	271, 105	3	574	157
753	4	271, 105, 128, 237	12	1100	209
1459	6	271, 105, 128, 237, 365, 325	28	2145	395
1777	8	271, 105, 128, 237, 365, 325, 74, 213	59	2587	462

Table 1: Case study systems.

System type	IEEE RTS					ERCOT Systems			
Buses	50	78	108	238	360	376	753	1459	1777
Regions and processes	2	3	4	2	3	2	4	6	8
Wall-clock time (sec) (Centralized)	1.9	2.4	4.2	7.2	11.7	17.6	37.3	66.2	89.5
Wall-clock time (sec) (Distributed)	4.6	3.3	3.0	5.2	5.1	16.5	15.1	16.9	19.3
Speed-Up	0.4	0.7	1.4	1.4	2.3	1.1	2.5	3.9	4.6
Efficiency (%)	20.7	24.2	35.0	69.2	76.5	53.3	61.7	65.3	57.9

Table 2: Speed-Up and Efficiency.

the first one that models optimal dispatch of the bulk of the ERCOT system.

Table 1 summarizes the test systems. The first column of table 1 shows the total number of buses in each system, while the second and third columns show the number of regions and the number of "core" buses in each region, where a "core" bus is a bus in the original system. The fourth column shows the number of tie-lines that interconnect the regions, while the fifth column shows the total number of lines in each complete system. The last column shows the total per unit loads in the systems.

The objective to be minimized is the total production cost for real and reactive power for the overall system. Real power costs were adapted from [11] and [15, §3 and §4], while the cost of reactive power was set equal to 10^{-6} of the real power cost for each generator. The parameters α, β, and γ were tuned for each system to improve convergence. A detailed theoretical discussion of the choice of these parameters is contained in [7, Appendix].

Transmission constraints were modeled as thermal limits on all individual lines and through voltage constraints on all buses. We recognize that this does not necessarily capture all the limits on transmission operation in ERCOT; however, the results provide a flavor for the capabilities of distributed OPF. Results of a limited study of distributed contingency-constrained OPF are contained in [7, §6.4].

6 Results

Table 2 shows the measured wall-clock time for the systems solved in two ways:

Centralized: solved as a single system on one processor with no decomposition, and

Distributed: solved on multiple processors using the decomposition scheme. Each region is assigned to a separate process.

In both cases, the IPOPF code was used to solve the corresponding OPF problems. Following [8, §3.3.2], we used the maximum mismatch between the border variables as the stopping criterion for the distributed algorithm, with a maximum mismatch tolerance of 0.03 per unit.

The last two rows of table 2 show the speed-up and efficiency based on the stopping criterion of 0.03 per unit mismatch. The efficiencies for the IEEE RTS are between about 20% and 75%, with higher performance for the larger systems. The efficiencies for the ERCOT systems are between about 50% and 65%.

These results were obtained for a network of workstations in close physical proximity where the communication latency accounts for less than 1% of the total wall-clock time. The efficiencies might be somewhat lower if the OPFs were run on geographically separated processors and there were significant communication delays. Nevertheless, the results in table 2 indicate that the decomposition scheme is extremely effective.

To visually display the results of the complete ERCOT study, figure 1 shows a contour plot of the marginal costs from the ERCOT study geographically superimposed on a map of Texas. The contour plot is based on the marginal costs at all of the generator buses and at most of 345kV buses, with contours drawn at 18, 23, and 28 $/MWh.

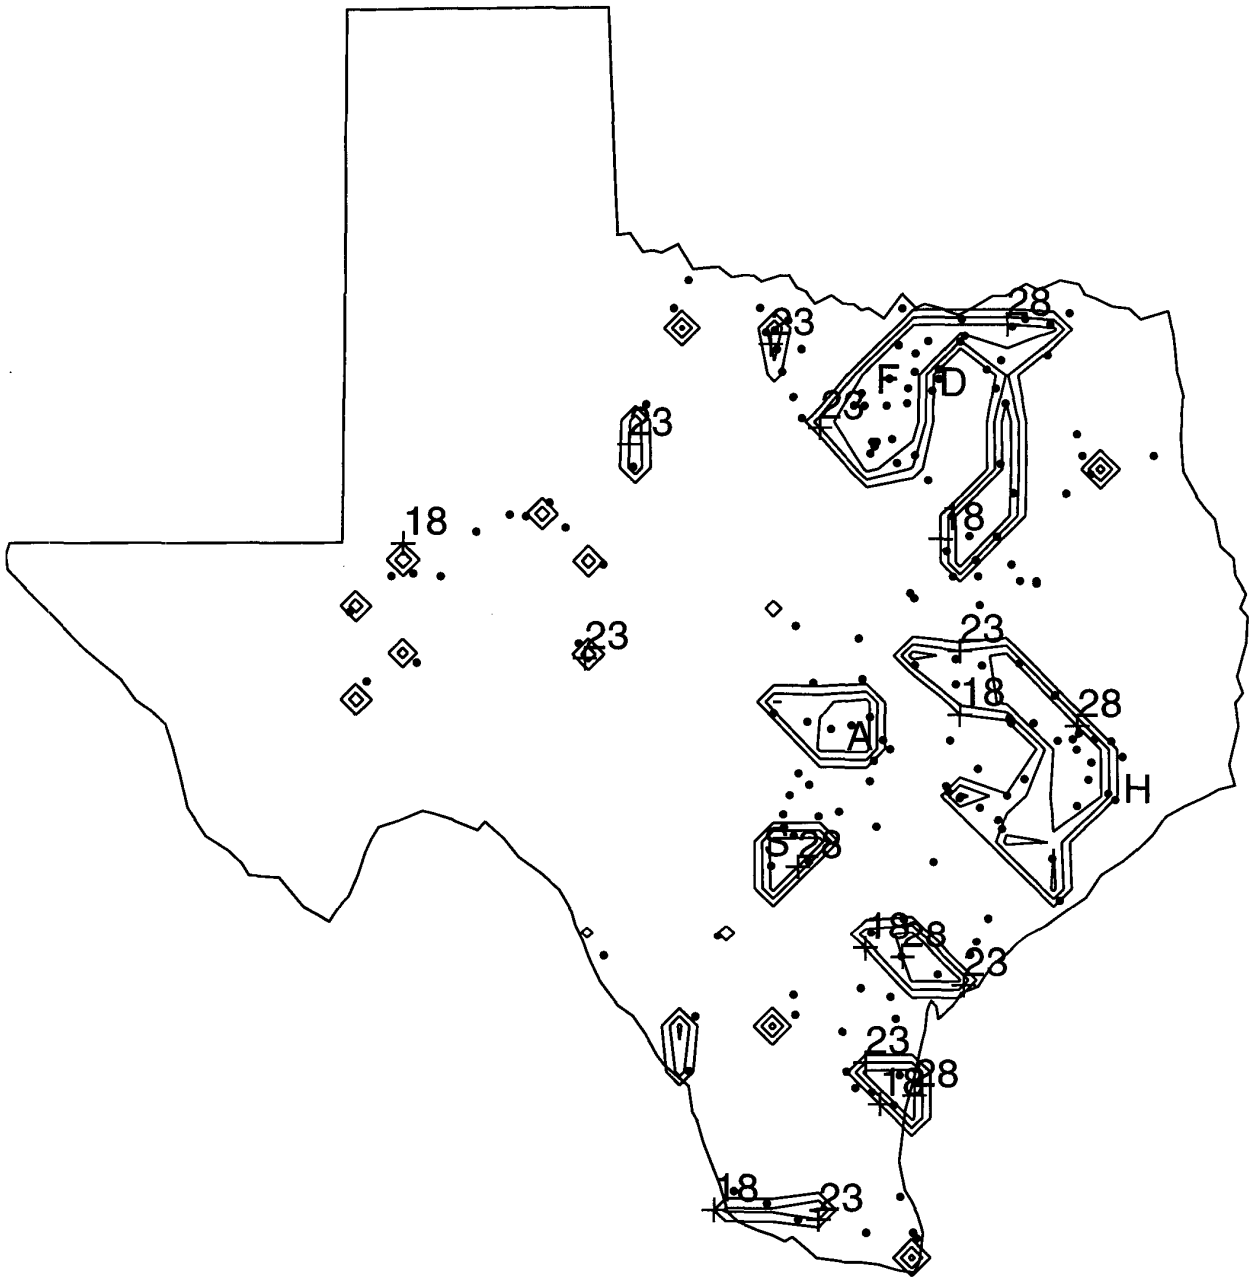

Fig. 1: Marginal costs for ERCOT case study system ($/MWh). Contours are drawn at 18, 23, and 28 $/MWh only. The letters A, D, F, H, S represent the approximate locations of the cities of: Austin, Dallas, Fort Worth, Houston, and San Antonio, respectively.

Some of the contours in figure 1 are marked with crosses. The adjacent numbers indicate the marginal cost of the corresponding contour. The dots in the figure represent the locations of generator buses. The letters A, D, F, H, S represent the approximate locations of the cities of: Austin, Dallas, Fort Worth, Houston, and San Antonio, respectively. Generally, the marginal costs increase towards the centers of these urban areas, as would be expected.

7 Conclusion and future work

We have demonstrated a fast distributed implementation of OPF. The same decomposition scheme can be applied to any power system optimization problem that can treat a separable quadratic objective. For example, we are currently building a distributed state-estimator based on this decomposition.

References

[1] Guy Cohen. Auxiliary problem principle and decomposition of optimization problems. *Journal of Optimization Theory and Applications*, 32(3):277–305, November 1980.

[2] Al Geist, Adam Beguelin, Jack Dongarra, Weicheng Jiang, Robert Manchek, and Vaidy Sunderam. PVM 3 users' guide and reference manual. Technical Report ORNL/TM-12187, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, May 1993.

[3] Harvey H. Happ and John M. Undrill. Multicomputer configurations and diakoptics: Real power flow in power pools. *IEEE Transactions on Power Apparatus and Systems*, 88(6):789–796, June 1969.

[4] G. Huang and W. Ongsakul. Speedup and synchronisation overhead analysis of Gauss-Seidel type algorithms on a Sequent balance machine. *IEE Proceedings, Part C*, 141(5):437–444, September 1994.

[5] Balho Kim. Modifications to interior point OPF code. Technical report, University of Texas at Austin, 1997.

[6] Balho Kim. Multi-area distributed OPF user's guide. Technical report, University of Texas at Austin, 1997.

[7] Balho Kim. *A Study on Distributed Optimal Power Flow.* PhD thesis, University of Texas at Austin, 1997.

[8] Balho Kim and Ross Baldick. Coarse-grained distributed optimal power flow. *IEEE Transactions on Power Systems*, 12(2):932–939, May 1997.

[9] Gabriel Kron. *Diakoptics: The Piecewise Solution of Large-Scale Systems.* Macdonald, London, 1963.

[10] Shyan-Lung Lin and J. E. Van Ness. Parallel solution of sparse algebraic equations. *IEEE Transactions on Power Systems*, 9(2):743–749, May 1994.

[11] Chris Marnay et al. Electric utility resource planning and production costing projects: Final report. Technical report, Center for Energy Studies, The University of Texas, Austin, TX, 1993.

[12] Tsutomu Oyama, Tatsuya Kitahara, and Yasuo Serizawa. Parallel processing for power system analysis using band matrix. *IEEE Transactions on Power Systems*, 5(3):1010–1016, August 1990.

[13] PVM—Parallel Virtual Machine, home page. http://www.epm.ornl.gov/pvm/pvm_home.html.

[14] Task Force of the Computer and Analytical Methods Subcommittee of the Power Systems engineering Committee. Parallel processing in power systems computation. *IEEE Transactions on Power Systems*, 7(2):629–637, May 1992.

[15] Allen J. Wood and Bruce F. Wollenberg. *Power Generation, Operation, and Control.* Wiley, New York, second edition, 1996.

[16] Yu-Chi Wu, Atif S. Debs, and Roy E. Marsten. A direct nonlinear predictor-corrector primal-dual interior point algorithm for optimal power flows. *IEEE Transactions on Power Systems*, 9(2):876–883, May 1994.

Ross Baldick received his B.Sc. and B.E. from the University of Sydney, Australia and his M.S. and Ph.D. from the University of California, Berkeley. He is currently an Assistant Professor in the Department of Electrical and Computer Engineering at the University of Texas at Austin.

Balho H. Kim received his B.S. from Seoul National University, South Korea and his M.S. and Ph.D. from the University of Texas at Austin. Currently, he is an Assistant Professor at Hong-Ik University in Seoul, South Korea. His research interests include Power System Planning and Operation, and Real-Time Pricing.

Craig Chase received his B.S. and Ph.D. degrees from Cornell University and his M.S. degree from Purdue University. He is currently an Assistant Professor in the Electrical and Computer Engineering Department at the University of Texas at Austin. His research interests include Parallel Computing and Distributed Systems.

Yufeng Luo received his B.S. degree from Shanghai Jiaotong University, China, and is a candidate for the M.S. degree from the University of Texas at Austin. He is currently working as a Research and Development Engineer at Synopsys Inc., Mountain View, California.

Acknowledgment This work was funded in part by the National Science Foundation under grant ECS-9496193 and under grant CCR-9409736. We would like to thank Dr. Atif Debs of Decision Systems International and Dr. Yu-Chi Wu for making the interior point OPF code available to us. We would also like to thank Mr. Wan Kyun Ham and Mr. Martin Lin of the University of Texas at Austin for help in creating the contour map.

Discussion

A. Conejo, N. Alguacil (U. de Castilla – La Mancha, Ciudad Real, Spain) **F. J. Prieto, J. Nogales** (U. Carlos III, Madrid, Spain) – We would like to commend the authors for their relevant contribution. A distributed solution technique to the multi-area optimal power flow is of high interest for two main reasons: (i) it allows the computation of an optimum for the global problem while preserving the independent operation of each area, and (ii) it permits a distributed implementation of the optimal power flow computation procedures.

The authors have greatly contributed to the development of algorithms to solve this multi-area optimal power flow problem.

Regarding the contents of the paper and the proposed algorithm, in our experience, an appropriate choice of the values for the updating parameter α and the penalty parameters β and γ is:

(i) not very well documented from theoretical results,

(ii) highly problem dependent and sensitive to problem data,

(iii) the key element to achieve fast convergence.

Furthermore, in some cases it may be of interest to alter the values of these parameters dynamically within the algorithm.

Could the authors comment on their experience in setting up the above parameters?

We congratulate again the authors for their excellent paper.

Ross Baldick (The University of Texas at Austin) and **Balho H. Kim** (Hong-Ik University, Seoul, South Korea): On behalf of all the co-authors, we would like to thank Drs. Conejo, Alguacil, Prieto, and Nogales for their discussion.

The discussers observe correctly that the choice of parameters α, β, and γ is problem dependent and influences the speed of convergence. However, our experience is that the convergence is not unduly sensitive to the choice of parameters. We were able to find a set of parameters for each base-case system studied that worked well for both base-case and corresponding outage cases. These results were reported in reference [8].

The relationship between parameters to assure convergence is sketched in the Appendix of reference [8], which applies the theoretical results of [1] under the assumption that the optimal power flow solution has certain "convexity" properties. Even though these convexity properties may not rigorously hold in practice, we found that the convergence results were a useful guide to the choice of parameters.

Dynamic updates of the parameters at each iteration and the relationship with other decomposition schemes is detailed in reference [7].