Parallel Domain-Decomposition-Based Distributed State Estimation for Large-Scale Power Systems

Hadis Karimipour, *Student Member, IEEE*, and Venkata Dinavahi, *Senior Member, IEEE*

***Abstract*—Growing system sizes and complexity, along with the large amount of data provided by phasor measurement units (PMUs), are the drivers to accurate state estimation algorithms for online monitoring and operation of power systems. In this paper, a distributed weighted-least-square state estimation method using an additive Schwarz domain decomposition technique is proposed to reduce the computational execution time. The proposed approach divides a data set into several subsets to be processed in parallel using a multiprocessor architecture considering data exchange among distributed areas. The slow coherency method and balanced partitioning are utilized to reduce the communication overhead and increase accuracy. Moreover, bad data analysis is also investigated in a distributed manner. The performance of the proposed distributed state estimator, along with the speed-up for several test systems, was compared with the traditional centralized state estimator. The simulation results show a speed-up of 6.5 for a 4992-bus system.**

***Index Terms*—Bad data identification (BDI), distributed state estimation, domain decomposition, large-scale systems, parallel programming, phasor measurement units (PMUs), weighted least square (WLS).**

I. Introduction

DOMAIN decomposition techniques are primarily partitioning methods that try to split a system into several subsystems that can be solved individually [1], [2]. The main advantage of decomposition techniques that makes them an excellent candidate for distributed state estimation is that they are suitable for parallel application on multiprocessors since independent subsystems can be solved simultaneously.

State estimation is a vital part of power system operation and control that tries to roughly estimate voltage magnitudes and phase angles in network buses [3], [4]. With growing system sizes, along with the wide usage of phasor measurement unit (PMU) data, the need for faster state estimation becomes more urgent. Accelerated performance using parallel programming to relieve computational burden on multicore central processing units and many-core graphics processors (graphics processor units) has been previously reported in studies such as transient stability simulation, electromagnetic transient simulation, and dynamic state estimation [5]–[7].

Generally, from a power system point of view, domain decomposition methods reduce the problem size by dividing the network into several subnetworks that results in less computation effort in each individual subnetwork. Most of the approaches in distributed state estimation randomly partitioned a system or are based on the geographical distances by assigning equal numbers of generators and buses among the subsystems, ignoring overlapping and boundary buses. However, this is not an efficient method because the network buses have different connectivity that leads to a load balancing problem and inaccuracy [8]–[10]. Other methods that tried to overcome the load balancing problem based on graph theory approaches are mainly too complicated for online applications [11]–[13]. Another option is to split the computation burden among processors based on the total number of equations. However, this approach will increase both the programming and communication complexities [14], [15].

In this paper, to overcome the aforementioned drawbacks, the measurement sets were divided into several subsets, instead of the network. Using an additive Schwarz method (ASM), the solution of each subsystem was carried out by using the conventional numerical techniques and exchanging the boundary data among subsystems. To increase the accuracy, a slow coherency method [16]–[18] was used to decide the domain decomposition. In addition, load balancing by distributing equal workload among processors was utilized to minimize interprocessor communication.

The proposed approach has the following advantages over existing approaches.

- It reduces the execution time by splitting equal amount of work among several processors.
- It minimizes the effect of boundary buses in accuracy by exchanging data.
- It localizes the effect of bad data on the state estimation result.
- It is applicable for large-scale power systems.
- It does not require major changes in the existing power system state estimation paradigm.

The rest of this paper is organized as follows. Section II describes the state estimation problem formulation. The domain decomposition method and its implementation are presented in Section III. Experimental results are illustrated in Section IV, followed by the conclusion in Section V.

Manuscript received January 5, 2015; accepted June 5, 2015. Date of publication September 29, 2015; date of current version March 17, 2016. Paper 2015-PSEC-0016, presented at the 2015 IEEE/IAS Industrial and Commercial Power Systems Technical Conference, Calgary, AB, Canada, May 5–8, and approved for publication in the IEEE TRANSACTIONS ON INDUSTRY APPLICATIONS by the Power Systems Engineering Committee of the IEEE Industry Applications Society. This work was supported by the Natural Sciences and Engineering Research Council of Canada.

The authors are with the University of Alberta, Edmonton, AB T6G 2V4, Canada (e-mail: hadis@ualberta.ca; dinavahi@ualberta.ca).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIA.2015.2483703

II. WLS State Estimation Formulation

The weighted least square (WLS) algorithm is the most popular method used by majority of the existing state estimators that tries to minimize the weighted sum of the squares of the residuals between the estimated and actual measurements [3]. Consider the measurement set vector $\boldsymbol{Z}$ as

$$\boldsymbol{Z} = \boldsymbol{h}(\boldsymbol{x}) + \boldsymbol{\varepsilon} \tag{1}$$

where the bold notation refers to arrays, and $\boldsymbol{Z}$, $\boldsymbol{h}(\boldsymbol{x})$, and $\boldsymbol{\varepsilon}$ are the vectors of measurements, nonlinear measurement functions, and measurement errors, respectively. For a system with n buses and m lines, there are $2m + 2n + 1$ elements in each vector, i.e., $2m$ power flows, $2n$ power injections, and slack bus measurements. $\boldsymbol{x}$ is a vector of system states comprised of voltage magnitudes and phase angles. Since the phase angle in a slack bus is considered 0, there are $2n - 1$ states to be estimated. For simplicity, error vector $\boldsymbol{\varepsilon}$ is assumed to be uncorrelated Gaussian noise with zero mean. Substituting the first-order Taylor's expansion of $\boldsymbol{h}(\boldsymbol{x})$ around $\boldsymbol{x}_0$ in (1), we obtain

$$\boldsymbol{Z} - \boldsymbol{h}(\boldsymbol{x}_0) = \boldsymbol{H}\boldsymbol{\Delta}(\boldsymbol{x}) + \boldsymbol{\varepsilon} \tag{2}$$

where $\boldsymbol{H} = \partial \boldsymbol{h}/\partial \boldsymbol{x}$ is the $(2m + 2n + 1) \times (2n - 1)$ Jacobian matrix, and $\boldsymbol{\Delta}(\boldsymbol{x}) = \boldsymbol{x} - \boldsymbol{x}_0$ is the $(2n - 1) \times 1$ state mismatch vector. The objective function $\boldsymbol{J}(\boldsymbol{x})$ to be minimized by the WLS formulation can be expressed as

$$\begin{aligned} J(x) &= \sum_{k=1}^{2m+2n+1} \left(Z_k - h_k(x)\right)^2 R_{kk}^{-1} \\ &= \left[\boldsymbol{Z} - \boldsymbol{h}(\boldsymbol{x})\right]^T \boldsymbol{R}^{-1} \left[\boldsymbol{Z} - \boldsymbol{h}(\boldsymbol{x})\right] \end{aligned} \tag{3}$$

where $\boldsymbol{R}$ is the $(2m + 2n + 1) \times (2m + 2n + 1)$ covariance matrix. Index k refers to the kth measurement. The following equation satisfies the first-order optimality condition at the minimum of $\boldsymbol{J}(\boldsymbol{x})$:

$$\boldsymbol{g}(\boldsymbol{x}) = \frac{\partial \boldsymbol{J}}{\partial \boldsymbol{x}} = \boldsymbol{H}^T(\boldsymbol{x})\boldsymbol{R}^{-1}\left[\boldsymbol{Z} - \boldsymbol{h}(\boldsymbol{x})\right] = 0 \tag{4}$$

where $\boldsymbol{g}(\boldsymbol{x})$ is the $(2n - 1) \times 1$ matrix of the gradient of the objective function. Substituting the first-order Taylor's expansion of $\boldsymbol{g}(\boldsymbol{x})$ in (4), the following equation is iteratively solved to find the solution that minimizes $\boldsymbol{J}(\boldsymbol{x})$:

$$\boldsymbol{G}(\boldsymbol{x})\boldsymbol{\Delta}(\boldsymbol{x}) = \boldsymbol{g}(\boldsymbol{x}) \tag{5}$$

$$\boldsymbol{\Delta}(\boldsymbol{x}) = \boldsymbol{G}(\boldsymbol{x})^{-1}\boldsymbol{g}(\boldsymbol{x}) = f(\boldsymbol{x}, \boldsymbol{\varepsilon}) \tag{6}$$

where $\boldsymbol{G}(\boldsymbol{x}) = \partial \boldsymbol{g}/\partial \boldsymbol{x}$ is the $(2n - 1) \times (2n - 1)$ gain matrix. The WLS state estimation algorithm given by (1)–(4) can be iteratively solved until the convergence of $\boldsymbol{\Delta}(\boldsymbol{x})$.

III. Domain Decomposition for Parallel and Distributed WLS State Estimation

There are two main approaches for decomposing a domain, i.e., overlapping and nonoverlapping subdomains. In this paper, nonoverlapping decomposition considering interfaces between subdomains was utilized. An additive Schwarz alternating procedure was then used to iteratively compute the solution via solving subdomain problems. A prerequisite for implementing the ASM on parallel processors is the decomposition of the original system into subsystems in which tightly coupled variables are grouped together. Since the coherency characteristic reflects the level of dependence between generators in an area, coherent generators can be grouped into the same subsystem that can be then independently solved from other subsystems.

A. Coherency Analysis

Following a disturbance in the system, some generators lose their synchronism with the network that causes sudden changes in the buses connected to those generators. State estimation for these buses will take more iterations since it should be repeated after clearing the disturbance. Partitioning the system into several areas, in which generators are in step together or are coherent, will increase the accuracy of the state estimation and save time by localizing the effect of disturbances. In a network, a pair of generators is called coherent if the difference between their rotor angles remains constant over time as follows:

$$\delta_i(t) - \delta_j(t) = \Delta\delta_{ij} \pm \epsilon \tag{7}$$

where $\Delta\delta_{ij}$ is a constant value, and ϵ is a small positive number. Since the coherent groups of generators are independent of the size of the disturbance and the complexity of the generator model, the linearized model of the system and the simple model of the generator can be used for time-domain simulation for coherency determination. Using the partitioning pattern discussed in [18], Case 1, an IEEE 39-bus system has been divided into three subsystems, i.e., {1, 8, 9}, {2, 3, 4, 5, 6, 7}, and {10}. For computational load balancing, another criterion was considered to have almost an equal number of buses in each subdomain. Satisfying both conditions simultaneously resulted in the following four domains that are shown in Fig. 1.

B. ASM

The ASM is a type of domain decomposition method that approximately solves a boundary value problem by splitting it into subproblems on smaller domains [19], [20]. Consider a decomposition of domain Ω into M nonoverlapping subdomains as follows:

$$\Omega = \bigcup_{i=1}^{M} \Omega_i, \quad \Omega_i \cap \Omega_j = \{\emptyset\}_{i \neq j}. \tag{8}$$

Let $\boldsymbol{\Delta}(\boldsymbol{x})_i^{(0)}$ denote the initial condition in subdomain Ω_i. A general ASM algorithm for state estimation can be written as

$$\begin{cases} \boldsymbol{G}(\boldsymbol{x})_i^{(k+1)}\boldsymbol{\Delta}(\boldsymbol{x})_i^{(k+1)} = \boldsymbol{g}(\boldsymbol{x})_i^{(k+1)} & \text{in } \Omega_i \\ \boldsymbol{x}_i^{(k+1)} = \boldsymbol{x}_j^{(k)} & \text{on } \partial\Omega_i \cap \Omega_j \\ \boldsymbol{Z}_i^{(k+1)} = \boldsymbol{Z}_j^{(k+1)} & \text{on } \partial\Omega_i \cap \Omega_j. \end{cases} \tag{9}$$

The process starts from an initial condition for all the subsystems. To achieve convergence, several iterations may be

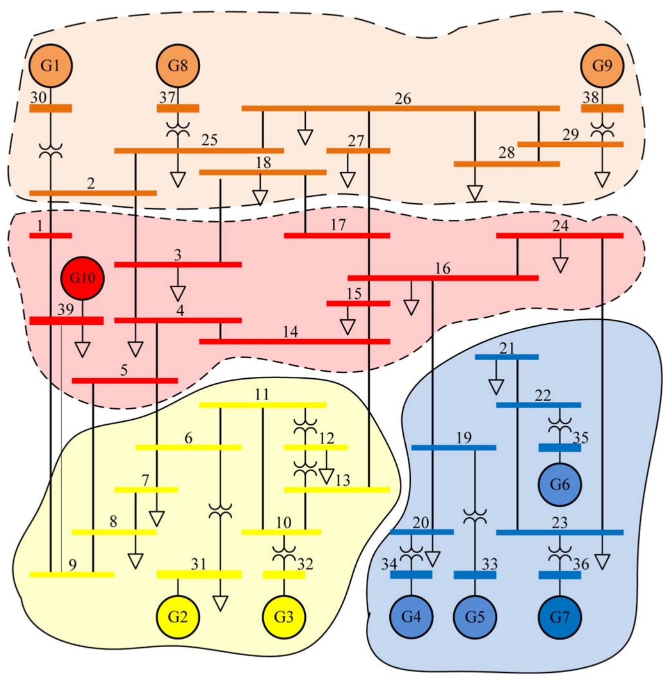

Fig. 1. Decomposing the Case-1 test system into four subsystems to apply the additive Schwarz algorithm.

required, where each of the subsystems exchanges boundary information and is then solved with the updated data collected from other subsystems. This process is repeated using a Jacobi method to iterate among subsystems until all variables converge with the necessary accuracy. The flowchart of the Jacobi ASM-based WLS algorithm for a time interval of $[0\ T]$ is depicted in Fig. 2.

C. Parallel Implementation of ASM-Based WLS

The hardware used in this paper is the hexacore Intel Xeon E5-2620v2 with a 2.1-GHz core clock and 32-GB memory with a 51.2-GB/s memory bandwidth, running the 64-bit Windows 8.1 operating system. Programming was done in Matlab using its Parallel Computing Toolbox. After decomposition, each subdomain is stored and computed by a processor core. All processor cores solve the subdomains in parallel. After each iteration, the estimated states of boundary nodes are transferred among processors to update each subdomain's boundary information. If the stopping criterion is not satisfied, a new iteration is performed. In case of bad measurements, state estimation will only repeat for the subdomain affected by the bad measurement. One of the main problems associated with the parallel ASM is the communication overhead among subsystems. In our paper, since subsystems are fully independent of each other, only necessary boundary bus values are needed to be exchanged. In addition, PMU measurements are considered installed in the boundary buses to accelerate data transfer among subsystems.

IV. Case Studies and Discussion

To evaluate the accuracy and efficiency of the parallel ASM algorithms, the results were compared with traditional central-

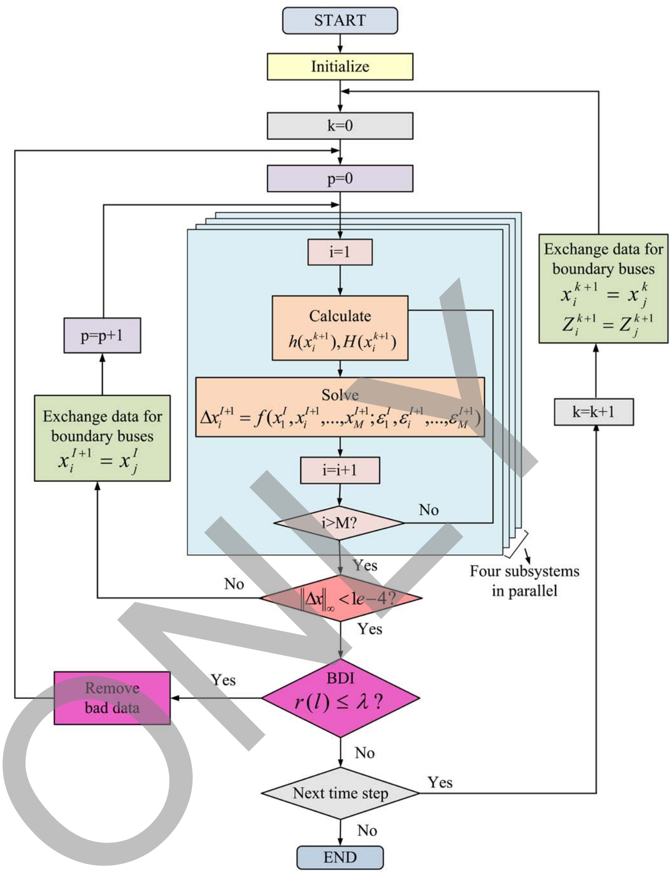

Fig. 2. Jacobi ASM-based WLS algorithm with BDI: k, the time step; i, the number of subsystems; p, the iteration counter; l, the index of the component in the residual vector.

ized state estimation. Large-scale case studies were constructed to explore the efficiency of the proposed method. Case 1 is the IEEE 39-bus system that was duplicated and interconnected to create large-scale systems. It is assumed that PMUs are installed at the boundary buses. The voltage magnitudes and angles of all buses are set to $1\angle 0$ per unit for flat start. The inputs to the parallel ASM-based WLS state estimator are the power flow results from PSS/E corrupted with noise that are used as pseudomeasurements. Therefore, to assess the accuracy of the state estimator, the results were also compared with the original power flow results from PSS/E.

A. Accuracy Analysis and BDI

The role of bad data identification (BDI) in state estimation is necessary since bad measurements easily affect the accuracy of the results. One of the popular methods for BDI that is used in this paper is based on a normalized residual test [21] as follows:

$$r_N(l) = \frac{|r(l)|}{\sigma_r(l)} \le \lambda \tag{10}$$

where $r_N(l)$ is the largest residual among all, and $\sigma_r(l)$ is the standard deviation of the lth component of the residual vector. In this paper, the measurements having the largest normalized

TABLE I
SUMMARY OF RESULTS FOR THE COMPARISON OF ASM-BASED WLS WITH CENTRALIZED WLS

Case	No. of buses	No. of meas.	Jacobian (H(x))	Gain (G(x))	$\mathbf{E}_V^{Cen.}$	$\mathbf{E}_\phi^{Cen.}$	$\mathbf{E}_V^{ASM}$	$\mathbf{E}_\phi^{ASM}$	$\mathbf{T}_{Ex}^{Cen.}$	$\mathbf{T}_{Ex}^{ASM}$	$\mathbf{S}_p$
1	39	171	171×77	77×77	0.009	0.085	0.006	0.46	0.08s	0.11s	0.72
2	78	347	347×155	155×155	0.004	0.06	0.003	0.43	0.26s	0.33s	0.78
3	156	699	699×311	311×311	0.0032	0.055	0.001	0.47	0.51s	0.61s	0.83
4	312	1421	1421×623	623×623	0.0041	0.06	0.0014	0.48	1.49s	0.91s	1.6
5	624	2865	2865×1247	1247×1247	0.0033	0.07	0.0012	0.49	5.2s	1.84s	2.8
6	1248	5825	5825×2495	2495×2495	0.004	0.065	0.0011	0.5	24.5s	8.1s	3.1
7	2496	11553	11553×4991	4991×4991	0.0044	0.06	0.0013	0.49	68.3s	15.5s	4.4
8	4992	23151	23151×9983	9983×9983	0.005	0.065	0.0017	0.49	364.5s	56.3s	6.5

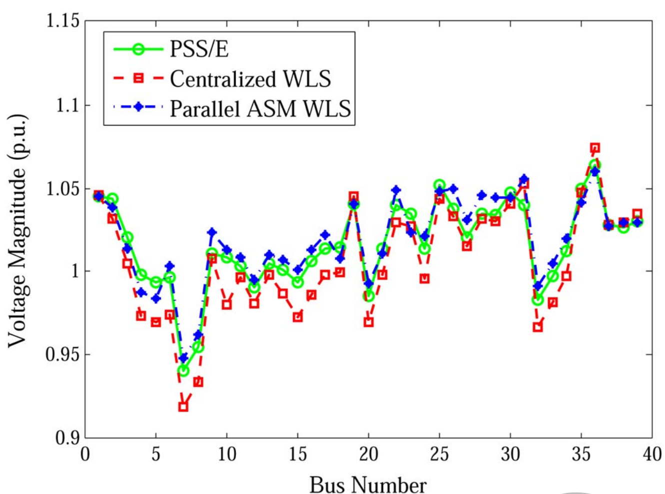

Fig. 3. Voltage magnitudes for Case 1 with respect to the system size.

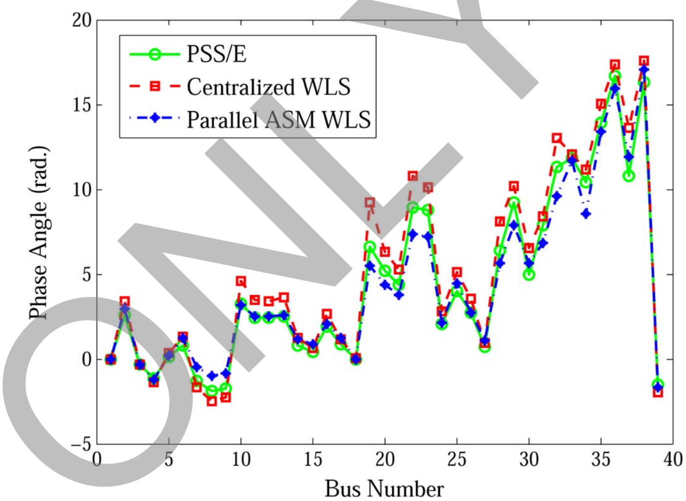

Fig. 4. Phase angles for Case 1 with respect to the system size.

residual and larger than 3 were considered bad data, with a 99.7% confidence level. After removing bad data, state estimation was repeated starting from the most recent estimate as many times as needed after each identification and elimination of bad data. Generally, BDI is a time-consuming process that slows the state estimation; however, since the data set here is divided into several subsets, state estimation was only repeated for the subset with bad data, not for the whole system. For large-scale systems, this localization of bad data can save a lot of time that can, in turn, accelerate the state estimation process.

To demonstrate the accuracy of the ASM-based WLS state estimation, the results of parallel distributed state estimation are compared with a traditional centralized state estimation method. The simulations were done using the test data sets listed in Table I, with a tolerance of 0.0001 for the convergence of the estimated parameters. The performance of the proposed method was evaluated for different case studies. The estimated states for Case 1 are shown in Figs. 3 and 4. It is clear from the results that the proposed approach can accurately estimate the voltage magnitude and the phase angle. The normalized Euclidian norm of the state estimation is also defined as a factor to evaluate the accuracy using

$$E_x = \frac{\|x - \hat{x}\|}{\sqrt{\dim(x)}} \tag{11}$$

where E_x is the normalized Euclidian norm of the estimation error, and $\dim(x)$ is the dimension of vector x. x and $\hat{x}$ are the vectors of true states and estimated states, respectively. Table I shows the accuracy index for both the voltage magnitude (E_V) and the phase angle (E_ϕ) for all case studies that clarifies the performance of the proposed method for large-scale systems.

B. Speed-Up Analysis

To demonstrate the efficiency of the proposed approach, the execution time using the parallel ASM ($T_{\text{Ex}}^{\text{ASM}}$) is compared with the traditional centralized ($T_{\text{Ex}}^{\text{Cen.}}$) state estimation method. As the results show for Cases 1–3, as these cases are small in size, the ASM is slower than the centralized method due to the communication overhead in parallelism. However, as the size of the system increased, the advantage of utilizing the ASM for parallelization was highlighted and resulted in almost 6.5 times speed-up (S_p) for Case 8.

Generally, when a system with α nodes is partitioned into M subsystems, each subdomain approximately has α/M nodes. Assume that solving a linear system with an iterative method has a complexity of $O(\alpha^\beta)$, where $\beta \geq 1$. Using the domain decomposition technique, the complexity of solving each subsystem is $O((\alpha/M)^\beta)$ that results in the complexity of $O((\alpha)^\beta/(M)^{\beta-1})$ for the entire system. It should be noted that this only occurs in the ideal case. However, it still justifies the speed-up reported in Table I and Fig. 5. As can be seen in Fig. 5, the percentage of the required execution time for

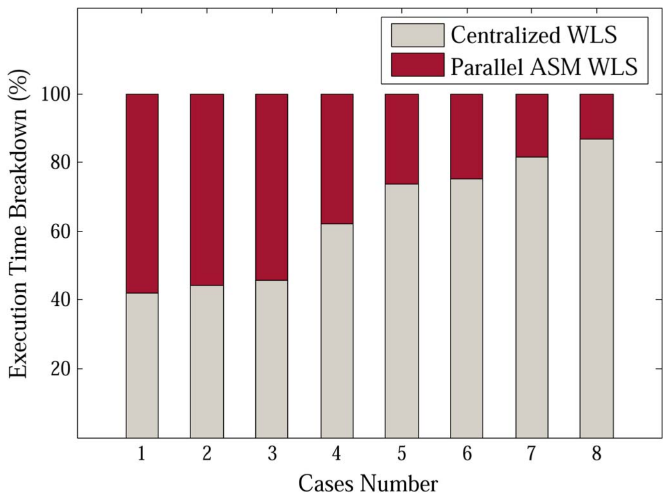

Fig. 5. Percentage of the execution time breakdown with respect to the system size.

the centralized WLS method increases very fast that shows the higher complexity of this method.

V. Conclusion

This paper has developed a parallel additive Schwarz-based WLS state estimator with bad data detection and identification using coherency analysis and balanced partitioning domain decomposition for large-scale power systems. Numerical experiments in this paper have proven that successful parallelization and balanced domain decomposition yield equal and less workload among processor cores, and they minimize the cost of communication overhead, resulting in reduction in the execution time; 6.5 times faster state estimation for a 4992-bus power system was reported. In addition, the accuracy of the proposed method was compared with the centralized WLS state estimator that proved higher accuracy for all case studies. The main objective of this paper has been to investigate the effect of domain decomposition and parallelization on the accuracy and speed of state estimation. Since the programming structure is one of the most important factors that affects the execution time, it is possible to achieve faster results by different multithread programming paradigms. For future work, a multithread programming structure using OpenMP and massive threading on graphic processors can further accelerate the process of state estimation significantly.

References

[1] T. F. Chan and D. Goovaerts, "On the relationship between overlapping and nonoverlapping domain decomposition methods," *IAM J. Matrix Anal. Appl.*, vol. 13, no. 2, pp. 663–670, Apr. 1992.

[2] B. Smith, *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[3] A. Abur and A. G. Exposito, *Power System State Estimation Theory and Implementation*. New York, NY, USA: Marcel Dekker, 2004, pp. 20–25.

[4] H. Karimipour and V. Dinavahi, "Accelerated parallel WLS state estimation for large-scale power systems on GPU," in *Proc. NAPS*, Sep. 2013, pp. 1–6.

[5] V. Jalili-Marandi, Z. Zhou, and V. Dinavahi, "Large-scale transient stability simulation of electrical power systems on parallel GPUs," *IEEE Trans. Parallel Distrib. Syst.*, vol. 23, no. 7, pp. 1255–1266, Jul. 2012.

[6] Z. Zhou and V. Dinavahi, "Parallel massive-thread electromagnetic transient simulation on GPU," *IEEE Trans. Power Del.*, vol. 29, no. 3, pp. 1045–1053, Jun. 2014.

[7] H. Karimipour and V. Dinavahi, "Extended Kalman filter based parallel dynamic state estimation," *IEEE Trans. Smart Grid*, vol. 6, no. 3, pp. 1539–1549, May 2015.

[8] R. Ebrahimian and R. Baldick, "State estimation distributed processing," *IEEE Trans. Power Syst.*, vol. 15, no. 4, pp. 1240–46, Nov. 2000.

[9] W. Jiang, V. Vittal, and G. T. Heydt, "Diakoptic state estimation using phasor measurement units," *IEEE Trans. Power Syst.*, vol. 23, no. 4, pp. 1580–1589, Nov. 2008.

[10] L. Xie, D. H. Choi, S. Kar, and H. V. Poor, "Fully distributed state estimation for wide-area monitoring systems," *IEEE Trans. Smart Grid*, vol. 3, no. 3, pp. 1154–1169, Sep. 2012.

[11] B. A. Carre, "Solution of load-flow problems by partitioning systems into trees," *IEEE Trans. Power Appl. Syst.*, vol. PAS-87, no. 11, pp. 1931–1968, Nov. 1968.

[12] L. Xie, D. H. Choi, and S. Kar, "Cooperative distributed state estimation: Local observability relaxed," in *Proc. IEEE PES*, Jul. 2011, pp. 1–11.

[13] Y. Weng, Q. Li, R. Negi, and M. Ilic, "Distributed algorithm for SDP state estimation," in *Proc. Innov. Smart Grid Tech.*, Feb. 2013, pp. 1–6.

[14] H. Sasaki, K. Aoki, and R. Yokoyama, "A parallel computation algorithm for static state estimation by means of matrix inversion lemma," *IEEE Trans. Power Syst.*, vol. PWRS-2, no. 3, pp. 624–632, Aug. 1987.

[15] D. M. Falcao, F. F. Wu, and L. Murphy, "Parallel and distributed state estimation," *IEEE Trans. Power Syst.*, vol. 10, no. 2, pp. 724–730, May 1995.

[16] S. B. Yusof, G. J. Rogers, and R. T. H. Alden, "Slow coherency based network partitioning including load buses," *IEEE Trans. Power Syst.*, vol. 8, no. 3, pp. 1375–1382, Aug. 1993.

[17] H. You, V. Vittal, and X. Wang, "Slow coherency-based islanding," *IEEE Trans. Power Syst.*, vol. 19, no. 1, pp. 483–491, Feb. 2004.

[18] X. Wang, V. Vittal, and G. T. Heydt, "Tracing generator coherency indices using the continuation method: A novel approach," *IEEE Trans. Power Syst.*, vol. 20, no. 3, pp. 1510–1518, Aug. 2005.

[19] A. Quarteroni and A. Valli, *Domain Decomposition Methods for Partial Differential Equations*. London, U.K.: Oxford Univ. Press, 1999.

[20] A. Toselli and O. Widlund, *Domain Decomposition Methods-Algorithms and Theory*. Berlin, Germany; Springer-Verlag, 2005.

[21] R. R. Nucera, V. Brandwajn, and M. L. Gilles, "Observability analysis and bad data analysis using augmented blocked matrices," *IEEE Trans. Power Syst.*, vol. 8, no. 2, pp. 426–433, May 1993.

Hadis Karimipour (S'11) received the B.Sc. degree in electrical engineering from Ferdowsi University of Mashhad, Mashhad, Iran, in 2007 and the M.Sc. degree from Shahrood University of Technology, Shahrood, Iran, in 2009. She is currently working toward the Ph.D. degree in electrical engineering at the University of Alberta, Edmonton, AB, Canada.

Her research interests include large-scale power system state estimation and parallel and distributed computing.

Venkata Dinavahi (SM'08) received the Ph.D. degree from the University of Toronto, Toronto, ON, Canada, in 2000.

He is currently a Professor with the Department of Electrical and Computer Engineering, University of Alberta, Edmonton, AB, Canada. His research interests include the real-time simulation of power systems, large-scale system simulation, and parallel and distributed computing.